

System Identification and Optimal Control for Mixed-Mode Cooling

by

Henry C. Spindler

A.B. Chemistry
Dartmouth College, 1992

M.S. Building Technology
Massachusetts Institute of Technology, 1998

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

September 2004

© Massachusetts Institute of Technology
All rights reserved



Signature of Author _____

Department of Mechanical Engineering
August 8, 2004

Certified by _____

Leslie K. Norford
Professor of Architecture
Thesis Supervisor

Certified by _____

Leon R. Glicksman
Professor of Mechanical Engineering and Building Technology
Thesis Committee Chairman

Accepted by _____

Ain A. Sonin
Chairman, Department Committee on Graduate Students

BARKER

Thesis committee:

Leslie K. Norford, Advisor, Professor of Architecture

Leon R. Glicksman, Chairman, Professor of Mechanical Engineering and Building Technology

Anuradha Annaswamy, Principal Research Scientist, Mechanical Engineering

Olivier L. de Weck, Assistant Professor of Aeronautics & Astronautics and Engineering Systems

System Identification and Optimal Control for Mixed-Mode Cooling

by

Henry C. Spindler

Submitted to the Department of Mechanical Engineering
on August 8, 2004 in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy in
Mechanical Engineering

ABSTRACT

The majority of commercial buildings today are designed to be mechanically cooled. To make the task of air conditioning buildings simpler, and in some cases more energy efficient, windows are sealed shut, eliminating occupants' direct access to fresh air. Implementation of an alternative cooling strategy—mixed-mode cooling—is demonstrated in this thesis to yield substantial savings in cooling energy consumption in many U.S. locations.

A mixed-mode cooling strategy is one that relies on several different means of delivering cooling to the occupied space. These different means, or modes, of cooling could include: different forms of natural ventilation through operable windows, ventilation assisted by low-power fans, and mechanical air conditioning.

Three significant contributions are presented in this thesis. A flexible system identification framework was developed that is well-suited to accommodate the unique features of mixed-mode buildings. Further, the effectiveness of this framework was demonstrated on an actual multi-zone, mixed-mode building, with model prediction accuracy shown to exceed that published for other naturally ventilated or mixed-mode buildings, none of which exhibited the complexity of this building. Finally, an efficient algorithm was constructed to optimize control strategies over extended planning horizons using a model-based approach. The algorithm minimizes energy consumption subject to the constraint that indoor temperatures satisfy comfort requirements.

The system identification framework was applied to another mixed-mode building, where it was found that the aspects integral to the modeling framework led to prediction improvements relative to a simple model. Lack of data regarding building apertures precluded the use of the model for control purposes.

An additional contribution was the development of a procedure for extracting building time constants from experimental data in such a way that they are constrained to be physically meaningful.

Thesis Supervisor: Leslie K. Norford
Title: Professor of Architecture

Table of Contents

- Table of Contents* _____ 5
- Acknowledgments* _____ 10
- 1 Introduction** _____ 11
 - 1.1 Mixed-mode Cooling: Definitions** _____ 11
 - 1.1.1 Mixed-mode Cooling versus Mixed-mode Ventilation _____ 12
 - 1.1.2 Mixed-mode Systems: Challenges and Rewards _____ 13
 - 1.2 Mixed-mode Cooling Control System: Ideal Characteristics** _____ 13
 - 1.3 Mixed-mode Cooling Control System: Thesis Contributions** _____ 14
- 2 The Energy-savings Potential of Mixed-mode Cooling and the Feasibility of Ventilation Cooling in the United States** _____ 19
 - 2.1 Review of Related Literature** _____ 19
 - 2.2 Model and Simulations** _____ 20
 - 2.3 Mechanical Ventilation Strategy** _____ 21
 - 2.4 Mixed-mode Cooling Savings Potential** _____ 21
 - 2.5 Mechanical Ventilation Effectiveness** _____ 22
 - 2.6 Results and Discussion** _____ 22
 - 2.6.1 Assessment of Savings due to MM _____ 22
 - 2.6.2 Assessment of Cooling Effectiveness of MV _____ 28
 - 2.7 Conclusions** _____ 29
- 3 Modeling Natural and Mixed-mode Ventilation in Buildings** _____ 31
 - 3.1 Challenges of Modeling Naturally Ventilated Buildings** _____ 31
 - 3.1.1 Introduction _____ 31
 - 3.1.2 Buoyancy-induced Multi-zone Airflow _____ 32
 - 3.1.3 Data-driven Models of Buoyancy-induced Flow in a Simplified Two-zone Building _____ 37
 - 3.1.3.1 Problem Statement _____ 37
 - 3.1.3.2 Models for “Heat Input” _____ 37
 - 3.1.3.2.1 Linear Regression Model _____ 37
 - 3.1.3.2.2 Kernel Recursive Least Squares Model _____ 39
 - 3.1.3.2.3 Model Assessment _____ 41
 - 3.2 Challenges of Modeling Complex Naturally Ventilated or Mixed-mode Buildings** _____ 42
 - 3.3 Proposed Solution** _____ 43
 - 3.3.1 Survey of Methods _____ 43
 - 3.3.2 PHDRT Theory _____ 45
 - 3.3.3 Proposed Changes to the PHDRT Algorithm _____ 47
 - 3.3.3.1 Modified p-value Significance Criterion _____ 47
 - 3.3.3.2 Automatic Dimension Reduction _____ 47
 - 3.3.3.3 Alternative Cut Locator _____ 47
 - 3.3.3.4 Alternative Error Criterion _____ 48

3.3.3.5	Additional Proposed Improvements	48
3.3.4	PHDRT in Practice	48
3.3.4.1	Demo 1	48
3.3.4.2	Demo 2	54
3.3.4.3	Demo 3: Modeling of Two-zone Data	58
3.3.4.4	Further Comments on the PHDRT Algorithm	59
3.4	Summary and Proposed Approach	60
4	<i>Broadmoor: Experimental Description</i>	61
4.1	Introduction	61
4.2	Building Description	61
4.2.1	Physical Description	61
4.2.1.1	Photographs of Building Exterior	62
4.2.1.2	Architectural Plans: Floor Plan and Building Section	65
4.2.1.3	Photographs of Building Interior	66
4.2.2	Design Intent	70
4.2.3	Building Usage and Operation	70
4.2.4	Experimental Setup	71
4.2.4.1	Modification of Openings to the Attic	71
4.2.4.2	Fan Installation and Properties	73
4.2.4.3	Cupola Improvements	76
4.2.4.4	Controlled Elements	77
4.2.4.5	Temperature Measurements	80
4.2.4.6	Weather Station Measurements	87
4.2.4.6.1	Technical Details: Wind Vane and Anemometer	87
4.2.4.6.2	Technical Details: Pyranometers	88
4.2.4.7	Campbell Scientific 21X Data Logger	90
4.2.4.8	Electrical Measurements	91
4.2.4.9	Weather Forecasts	96
4.2.4.10	PC Data Acquisition and Control	97
4.2.4.10.1	Startup Processes	97
4.2.4.10.2	Processes Occurring Every Hour	98
4.2.4.10.3	Processes Occurring Every Minute	98
4.2.4.11	Graphical User Interface	100
4.2.4.12	Summary of Data Collected	101
4.2.5	Comparison with Prior Measurements	103
4.3	Summary	105
5	<i>Broadmoor: Thermal Modeling</i>	107
5.1	Literature Review	107
5.1.1	Physically-based Models	107
5.1.2	Data-driven Models	109
5.1.3	Data-driven Modeling of Conventional Buildings and Building Systems	111
5.1.4	Literature Review Summary	113
5.2	Definition of Modeling Terminology	113
5.3	Normalization	116
5.4	Development of Potential Features	117

5.5	Definition of Control Modes	118
5.6	Application of PHDRT	121
5.6.1	Sunspace	121
5.7	Linear Modeling Approach	124
5.7.1	Definitions for Model Assessment	125
5.7.1.1	Error Measures	125
5.7.1.2	Simulation Types	125
5.7.1.3	Data Sets: Training, Test and Validation	126
5.7.1.4	Regression with Singular Value Decomposition	126
5.7.2	Initial Modeling Example: Sunspace	127
5.7.3	Multi-zone Modeling and the Performance Benefit of Implementing Mode Models	134
5.7.3.1	Data Description	135
5.7.3.2	Multi-zone Model	141
5.7.3.3	Implementation of Multi-mode Models	151
5.8	Automated Feature Selection	167
5.9	Application of Automated Feature Selection	169
5.10	Automated Mode Selection	173
5.11	Implementation of Automated Mode Selection	175
5.11.1	Illustration of Automated Mode Selection	175
5.11.2	Model Formation, Results and Discussion	181
5.12	Model Testing	183
5.13	Performance Comparison with a Physically Based Model	186
5.14	Performance Comparison with Nonlinear Models	188
5.15	The Effect of Wind	190
5.16	On-line Modeling	193
5.17	Summary	193
6	<i>Broadmoor: Optimal Control Strategy Development</i>	197
6.1	Introduction	197
6.2	Literature Review	197
6.2.1	Thermal Mass Management: Naturally Ventilated and Mixed-mode Buildings	197
6.2.2	Local Control: Naturally Ventilated and Mixed-mode Buildings	199
6.2.3	Thermal Mass Management: Conventional Buildings	201
6.2.4	Local Control: Conventional Buildings	203
6.2.5	Predictions of Exogenous Inputs	205
6.2.6	Literature Review Summary	205
6.3	Control Limitations and Statement of Goals	206
6.4	Solution Approach	208
6.4.1	Further Definition and Reduction of the Control Problem	209
6.4.2	Two Solution Methods	210
6.4.3	Sensitivity Analysis and Pareto Optimality	218
6.5	Sub-cooling	219

6.6	Multi-day Optimization	221
6.6.1	Two-day Optimization	221
6.6.2	Three-day Optimization	223
6.7	Practical Implementation of Optimal Control and Extension to the Supervisory Control/Local Control Architecture	225
6.8	Incorporation of More Complex Mechanical Cooling	228
6.9	Summary	228
7	<i>Modeling of Houghton Hall</i>	229
7.1	Introduction	229
7.2	Building Description	229
7.3	Building Measurements	231
7.3.1	Temperature Measurements	231
7.3.2	Solar Radiation	233
7.3.3	Wind Speed	235
7.3.4	Electrical Usage	235
7.3.5	Data Summary	235
7.4	Model Development	239
7.4.1	PHDRT	239
7.4.2	Simple Model	239
7.5	Model Using Modes, but no Feature Optimization	240
7.6	Model Using Modes and Feature Optimization	241
7.7	Application of the Model	246
7.8	Summary	249
8	<i>Building Dynamic Modes, Time Constants and Time Shifts</i>	251
8.1	Introduction	251
8.2	Discussion of Dynamic Modes	251
8.2.1	Example of Modal Analysis	252
8.3	Extraction of Time Constants from Data	257
8.3.1	Selection of Model Type	257
8.3.2	Constraints Imposed by Dynamic Thermal Systems	260
8.3.2.1	Constraints on Poles	260
8.3.2.2	Constraints on Zeros	262
8.3.3	Algorithm for Obtaining a Physically Plausible Model	262
8.3.3.1	The Algorithm	262
8.3.3.2	Demonstration of Algorithm	264
8.3.3.3	A Caveat	266
8.4	Houghton Hall Time Constants	266
8.4.1	Summertime	266
8.4.2	Wintertime	270
8.4.3	First-order Models	271
8.5	Discussion of Time Constants and Time Shifts	272
8.5.1	Calculating Time Shifts with the System Transfer Function	272

8.5.2	Frequency Response for Houghton Hall	272
8.6	Broadmoor Time Constants	275
8.7	Summary and Possible Future Extension of Concepts	277
9	<i>Summary, Conclusions and Future Research</i>	279
9.1	Summary	279
9.2	Conclusions	282
9.3	Future Research	284
10	<i>References</i>	285
<i>Appendix 1 Electrical Logger Program Files</i>		295
A1.1	C180E (P1630.001)	295
A1.2	K20 (P10046.001)	297
<i>Appendix 2 Campbell Scientific, Inc. Data Logger Usage Details</i>		299
A2.1	Wiring	299
A2.2	Code Listing of Program Loaded onto 21X	300
A2.2.1	BRDMR.CSI	300
A2.2.2	BRDMR.DLD	305
<i>Appendix 3 Data Retrieval from Loggers</i>		307
A3.1	FIRSTDLD.BAS	307
A3.2	UPDATE.BAS	311
A3.3	C180.BAS	315
A3.4	SC1630.DAT	316
A3.5	K20.BAS	316
A3.6	S10046.DAT	317
<i>Appendix 4 Description of Manual Monitoring and Control Spreadsheet</i>		319

Acknowledgments

Financial support for this research was provided by several sources: the d'Arbeloff Foundation, the Kann-Rasmussen Foundation, the MIT Mechanical Engineering Department, the Cambridge-MIT Institute, the Martin Family Society of Fellows for Sustainability, and the Dean for Graduate Students. This generous support was greatly appreciated.

Despite the requisite stressful periods preceding the milestones on the path to a Ph.D., these four years have been extremely enjoyable. This has been possible thanks to the caliber and character of the people with whom I have been fortunate to work.

I am grateful for the many discussions with colleagues in the Building Technology lab as well as for the moral support throughout this process. In particular, I would like to thank Acha Chutarat, Yi Jiang, Helen Xing, Jonathan Smith, Lara Greden, Roger Chang, Rogelio Palomera-Arias, Christine Walker and Peter Armstrong. To Roger, thanks for sharing your experiences and warm welcome at Broadmoor, as well as all of your experimental data files. To Christine and Rogelio, thanks especially for your help with the defense presentation and to Christine for all the ponderings regarding Luton. Peter, without your assistance, I'd probably still be setting things up at Broadmoor! I'm grateful for your help with all the experimental issues, and for the interesting discussions for which I wish we had had more time.

To the staff in the Building Technology and Mechanical Engineering headquarters, thank you for helping smooth the many bumps in the process and for making the administrative details as pleasant as possible. Thank you Dorrit Schuchter, Kathleen Ross, Sara MacDonald, Nancy Dalrymple, Leslie Regan and Joan Kravit!

I'm especially grateful to Elissa Landre at the Massachusetts Audubon Society's Broadmoor Wildlife Sanctuary, who quite literally opened the doors of her building for me, allowing an extended experimental period over the summer and early fall of 2003. Her eagerness for us all to learn as much as possible from her unique building was refreshing. Elissa, thank you for making this research possible, and thanks to you, Dan, Nancy, Chuck, Mike, Melissa and the entire volunteer staff for making my stay in your building so enjoyable.

Of course, without the confidence and guidance of my committee members this effort would have been much more arduous and ultimately less fruitful. I'm grateful for the suggestions, high expectations and encouragement in all our meetings—individual, or as a group. To my advisor, Les Norford, the hours (weeks!) of your time you shared with me in discussions of core and tangential matters were invaluable and always enjoyable to me. I'm thankful the great flexibility I was granted to pursue research questions and for your steady interest throughout this long and tortuous (though not torturous!) project. Thank you!

Finally, my greatest thanks are due to Carol, without whose blessing, support, infinite patience and ceaseless effort, none of this work would have been possible. Danke Dir!

1 Introduction

As awareness grows in the United States of the environmental, economic and national security implications of energy consumption, it is natural that attention is increasingly focused on opportunities for improving energy efficiency in buildings. Since buildings are responsible for greater than one third of the energy consumption in the U.S. (DOE), any increases in the efficiency with which they are heated, cooled and ventilated could have an important impact on national energy consumption.

The majority of commercial buildings today are designed to be mechanically cooled. To make the task of air conditioning buildings simpler, and in some cases more energy efficient, windows are sealed shut, eliminating occupants' direct access to fresh air. Although such an approach is prudent in some places, in many U.S. locations, substantial savings in cooling energy consumption are attainable through the implementation of an alternative cooling strategy: mixed-mode cooling.

1.1 Mixed-mode Cooling: Definitions

Mixed-mode cooling takes many forms. In essence, a mixed-mode cooling strategy is one that relies on several different means of delivering cooling to an occupied space. These different means (or modes) may include: cooling via natural ventilation (NV), cooling via fan-driven, or mechanical ventilation (MV), and cooling via mechanical air conditioning (AC). As an illustration, consider the simple office conditioned by mixed-mode cooling that is pictured below in Figure 1-1.




Figure 1-1 Above Left: Natural ventilation mode. Above Right: Mechanical ventilation mode. Left: Air conditioning mode.

When the outdoor temperature is cooler than the indoor temperature, the office windows may be opened to let in cool air to the office. In the case pictured, cool air enters through the lower awning window and warm air exits through the upper window. This constitutes the NV cooling

mode (upper left pane). When the natural airflow rates and the temperature difference between inside and outside air are not sufficient to remove heat gains from the office, an exhaust fan may be turned on to draw air into the space at an increased rate. This constitutes the MV cooling mode (upper right pane). If the outside air temperature exceeds the inside air temperature and cooling is needed, then the windows may be shut and the air conditioning may be turned on (lower left pane).

Such a mixed-mode cooling strategy is considered to be a “changeover” strategy (BRE 1999a) since one cooling mode is used at a time. Mode switching can occur over the course of a single day to provide comfortable conditions to the space. An “alternate operation” strategy switches between modes on a longer time scale, perhaps seasonally. A “concurrent” strategy employs multiple modes simultaneously.

The definition of a mode is generalized in this thesis in recognition of the fact that there may be more than one type of NV cooling possible. For example, consider a single floor of an office building with windows on two opposing sides. If the windows are opened on one side only, then a certain airflow pattern will develop, resulting in a particular cooling rate and distribution. If the windows on the opposing side are opened at the same time, an entirely different airflow pattern may arise due to the possibility of cross ventilation. The associated cooling rate and distribution are characteristic of this particular configuration of open windows. We will consider these different operating conditions to correspond to different NV cooling modes.

1.1.1 Mixed-mode Cooling versus Mixed-mode Ventilation

Building systems employ outdoor air to achieve several goals. In one case, the primary purpose is to provide fresh air to the building occupants. A year-round strategy must be developed to supply adequate, but not excessive, levels of fresh air in all seasons. Such a strategy might employ trickle ventilators under windows to allow a controlled amount of air ingress in the wintertime without creating drafts (BSE 1996). For entirely passive systems, a ventilation strategy designed to meet such goals would be called natural ventilation. When mechanical systems are used as well, the strategy would be called mixed-mode ventilation (the term hybrid ventilation is also used).

In the other case, outdoor air is used for cooling the building. The inside temperature is controlled passively through natural ventilation (by adjusting building openings) or by mixed-mode cooling (by adjusting building openings and utilizing mechanical systems). The focus in this thesis is on such systems. (To confuse matters, the term hybrid ventilation is used to describe these systems as well.)

Mixed-mode cooling could also refer to the joint control of indoor temperature and humidity. This important topic has not been addressed in the existing research on mixed-mode cooling or in the thesis¹.

¹ The interested reader may wish to consult the literature on the operation of air-side economizers for conventional air-conditioning systems.

1.1.2 Mixed-mode Systems: Challenges and Rewards

As a consequence of its potential benefits, mixed-mode cooling has received increasing attention over the past few years. A multi-year project was organized by the International Energy Agency (Annex 35: Hybrid Ventilation in New and Retrofitted Office Buildings) to study buildings employing this type of cooling. The purpose of the Annex was to survey existing buildings using hybrid ventilation, establish the key barriers to its implementation and success, survey and improve existing control strategies for this type of building, and describe relevant analysis tools (Delsante and Vik 2001). A similar study had been completed for naturally ventilated buildings through the JOULE/NatVent project (BRE 1999; NatVent 1999).

Mixed-mode cooling presents challenges not encountered in the design and implementation of traditional cooling systems. The primary issue to contend with is complexity. Not only must a set of different cooling systems be designed, but they also must be designed to work together. Smooth mode-to-mode transitions must be ensured for both comfort and equipment durability reasons. A significant element of planning must contribute to the control strategy.

Many sources exist to aid in the design of systems with natural ventilation (BS 1991; BRE 1994; Irving *et al.* 1995; Etheridge and Sandberg 1996; CIBSE 1997; Kolokotroni *et al.* 1997; Allard 1998; CIBSE 2000; Allocca 2001; Awbi 2003), and several are available for mixed-mode systems (EDSL 1996; Heiselberg and Tjelflaat 1999; BRE 1999a; Heiselberg 2002). These sources focus primarily on ensuring adequate cooling capacity and comfort, rather than on details of the control systems. One exception is (Martin 1995).

The rewards of employing a mixed-mode strategy may be significant in terms of energy savings, as will be shown in the following chapter. However, the benefits may extend further to occupant satisfaction and productivity due to the occupants' perceived link between comfort and the presence of operable windows. Critical to achieving these benefits, however, is the occupants' clear understanding of the system's operation. The occupants' full cooperation and at least a basic system understanding were identified in the Annex 35 reports to be essential components in the success of these innovative designs.

1.2 Mixed-mode Cooling Control System: Ideal Characteristics

With one possible exception, the ideal elements of a mixed-mode cooling control system are not unique to this particular application. These elements include the capability of:

1. Providing satisfactory temperature control as soon as the building is occupied
2. Learning over time the building's response to a variety of environmental and control inputs
3. Incorporating system knowledge into a control strategy that provides comfortable conditions, uses as little energy as possible, and preserves actuator longevity
4. Reacts smoothly to user overrides and integrates user preferences into future strategies

The first element is desirable since control-system commissioning can be time-consuming and costly. The second element is useful not so much for its own sake, but rather because it enables the development of optimal control strategies in the third element. It is the fourth element that

sets the mixed-mode cooling control problem apart from many others². It would be highly desirable to incorporate knowledge of user behavior into optimal control strategies.

For example, if the zone to be controlled has a single office, the occupant's typical work hours could be learned. Furthermore, it could be noted that the occupant prefers warmer temperatures at the start of the workday than colleagues in other offices. Since nighttime cooling need not be so aggressive for this occupant relative to that for others, an individually tailored control strategy may require less energy.

To provide the reader with an example of the control strategies currently in use, a selection was made from the strategies presented by Levermore (2000). The following strategy for cooling the building during nighttime hours was implemented at the PowerGen building in Coventry, UK.

Initiate and maintain cooling by opening all controlled apertures when:

1. The average room/zone temperature at the end of the day exceeded 23°C
2. The maximum outside temperature during the day exceeded 21°C
3. The room/zone temperature exceeded 18°C

If condition 3 had been violated, night cooling was not resumed until the room/zone temperature rose above 20°C.

It not evident that this night-cooling strategy meets the ideals cited above. Satisfactory control may have been provided, but probably at a significant commissioning expense. The simplicity and heuristic nature of the strategy provides no evidence of model-based control optimization (with either physically-based or data-driven models).

1.3 Mixed-mode Cooling Control System: Thesis Contributions

The emphasis of this thesis research was on the fundamental elements of the control system: system identification and optimal control. The particular system investigated included multiple NV modes in addition to an MV mode. No AC mode was available.

The first ideal cited above remains an intriguing topic of research. Existing research on the control of an unknown plants may yield fruit in this application (Cui and Shin 1993; Sheen and Kumara 1993; Sebald and Schlenzig 1994; Chan and Rad 2000).

The fourth ideal element has been addressed to some degree by others. Seginer (1996) demonstrated how expert knowledge (in growing greenhouse crops) could be assimilated into a neural network that could then emulate the expert's behavior. Occupancy patterns were learned by a neural network and used to advantage in optimizing the operation of a residential heating system (Mozer *et al.* 1996).

The primary contributions of this thesis center on the second and third ideal elements cited above. Specifically, these are:

² Although this capability is now integral to many automotive control strategies.

1. A novel system identification framework was developed to accommodate the unique features of mixed-mode systems.
2. The modeling framework was applied to an existing multi-zone mixed-mode building and was demonstrated to yield highly accurate predictions of the building's thermal response.
3. An efficient algorithm was developed to use system knowledge encoded in the model to generate optimal control sequences over a 24-hour period.

Several underlying principles guided the work:

1. The modeling and control frameworks developed must be portable and not building specific.
2. A machine-learning perspective was taken with the modeling. If a certain approach yielded accurate predictions, it was used even if it possessed physically unrealistic attributes (such as oscillatory behavior in a thermal system).

Implicit in the first principle was the understanding that the actual model or control strategy developed for one building would be of no use for another. It is the modeling and control optimization *process*, or *framework*, that may be applied in other situations to yield building- and climate-specific control strategies.

The second principle is clearly at odds with a physically based modeling approach in which a collection of physical relationships are assembled to form a predictive model of a building. Two factors influenced the decision to pursue a data-driven rather than a physics-driven modeling approach.

The first factor was the difficulty of creating a reliable physical model of the thermal processes in the building. Of foremost concern was the influence of natural ventilation, whose impact is a function of the window and door configuration, their opening sizes and geometries, the inside-outside temperature difference, the wind speed and direction, and interior obstacles to airflow. One may make certain assumptions, such as the air within the building being well-mixed and characterized by a single temperature, and then proceed to construct an analytical model. Unfortunately, such an approach is tractable for only the simplest of configurations. Furthermore, it will be demonstrated in Chapter 5 that even the most basic *a priori* assumptions regarding the nature of the influence of natural ventilation can differ widely from experimental observations.

One could also resort to the use of a zonal or computational fluid dynamics (CFD) airflow modeling program, but such an approach requires the user to establish (and later calibrate) appropriate inputs to the programs, such as window discharge coefficients, convection heat transfer coefficients, *etc.* Furthermore, the computational cost of running simulations (particularly CFD) is prohibitive in a real-time application. Ultimately, the effort spent modeling a single zone using analytical models or CFD would have to be repeated for each new zone configuration or new building.

The structure of a building's thermal model (excluding airflow) is more straightforward. A simple nodal model can readily capture the thermal dynamics. However, *a priori* identification of appropriate parameters is both time-consuming and error-prone, due, in large part, to the great variability in installation quality.

To summarize the first reason for selecting a data-driven modeling approach, the overall model structure and parameters are difficult to establish *a priori*. If it were possible to measure flow through all building apertures in real time, one could bypass the difficulties in predicting flow with a physically based model by simply using measurement data. (Such an approach was taken by Linker *et al.* (1999) in a greenhouse by injecting CO₂ and associating the concentration decay rate with the rate of airflow through the greenhouse.) To date, however, no practical methods exist for measuring bulk airflow through multiple openings of an occupied building.

The second factor influencing the decision to pursue a data-driven rather than a physics-driven modeling approach was the desirability of optimizing the operation of a building as it currently exists. The passage of time and, certainly, renovations can have an impact on the thermal performance of a building. A data-driven model can capture the actual impact of such changes in a way that would be difficult for a physical model to reproduce.

The decision to approach the modeling problem from a data-driven perspective necessarily excludes the possibility of using the modeling framework developed in this thesis for design purposes. Without data, there can be no predictions. However, if a physically-based model were developed as part of the design process (*e.g.*, using TAS (EDSL 1996)), it could be used in conjunction with the control optimization framework developed in the thesis to assess the types of optimal control strategies that could arise from different environmental conditions and different set-point requirements. Furthermore, such a model could serve to inform control decisions in the actual building until the data-driven model was trained. This approach is analogous to that proposed by Curtiss (1996), where a data-driven neural-network model relied on the control inputs generated by a PID controller for the period when the neural network was training.

The remaining chapters of the thesis are organized as follows:

- Chapter 2 investigates the potential energy savings attainable through the use of mixed-mode cooling in different regions of the United States. The feasibility of using outdoor air alone for cooling is also explored.
- Chapter 3 illustrates some of the modeling challenges associated with natural ventilation and proposes a modeling approach to handle the inherent nonlinearities.
- Chapter 4 introduces the building used for experimental purposes—The Saltonstall Nature Center at the Massachusetts Audubon Society's Broadmoor Wildlife Sanctuary—and details the experimental setup and methods.
- Chapter 5 presents the development of the proposed modeling framework as implemented using measurements from Chapter 4.
- Chapter 6 proposes and demonstrates a procedure for control strategy optimization.
- Chapter 7 investigates the performance of the proposed modeling framework on a different building: Houghton Hall in Houghton Regis, U.K.

- Chapter 8 presents a discussion of the extraction of physically realistic time constants from experimental data.
- Chapter 9 concludes the thesis and provides suggestions for future work.
- References and Appendices follow Chapter 9.

Throughout the document, references are made to MATLAB scripts used by the author to generate figures or perform certain calculations. These are included primarily for the author's benefit. Should the reader have questions about the scripts, please direct them to the author via email: spindler@alum.mit.edu.

2 The Energy-savings Potential of Mixed-mode Cooling and the Feasibility of Ventilation Cooling in the United States³

The use of natural ventilation to condition office buildings is much less prevalent in the United States than in many European countries. The purpose of this study was to determine whether there is a climatic rationalization for this observation, or whether energy savings and comfortable conditions are achievable while reducing or even eliminating the role of the standard mechanical cooling system. To address this question, the cooling energy consumption and comfort conditions of an office space in 40 U.S. cities were examined under several cooling strategies: standard air conditioning (AC), mixed-mode cooling (MM), and mechanical ventilation cooling (MV).

2.1 Review of Related Literature

Kammerud *et al.* (1984) showed that a significant potential exists in a variety of locations in the U.S. to reduce mechanical cooling loads in residential buildings via ventilation cooling with fans. The authors found that cooling loads could be reduced as much as 60-80%, depending on the location, night cooling strategy and ventilation rates. However, the energy consumption of the fans used to reduce the cooling load was not considered in that study.

In a more recent study, Axley and Emmerich (2002) performed statistical analyses on a set of U.S. weather files to develop design guidelines providing estimates of the effectiveness of MM for a variety of internal load conditions and building locations. Using the concept of a building's balance-point temperature, the authors estimated the airflow rates that would be required to remove internal loads of 10, 20, 40, and 80 W/m² (at those times when outdoor air may be used for cooling). Also estimated was the equivalent daytime internal load per flow rate that could be removed by ventilating at night. No assessment was made of the cost of using a fan to provide the required flow rates. More complete reports of the authors' work are cited in this paper, as is a similar study on buildings in Canadian cities.

Daly (2002) used TAS software (EDSL 1996) to investigate the potential energy savings associated with a variety of MM control strategies. The thermal response of a single office in Merced, California was simulated on a particular day and over an entire year. No night cooling was used; windows could be opened during the daytime only. Daly presented a MM strategy that could provide identical thermal conditions to those provided by a standard air-conditioning system while using 10% less energy annually. Presumably, one could anticipate greater savings if night cooling were considered.

Several related studies have been performed with a focus on European applications.

Kolokotroni *et al.* and Kolokotroni (1997; 2001) performed parametric studies to assess the effect of different rates of night cooling on peak temperatures and cooling energy consumption on the following day. Overall cooling energy consumption (including fan energy) was

³ Much of the material presented in this chapter was previously published in (Spindler *et al.* 2002).

considered as well. The simulation tool, NiteCool, was developed specifically for the U.K. climate.

Geros *et al.* (1997) presented a study of the impact of night cooling on three Athens buildings. A parametric study of cooling load reduction versus flow rate was performed, as well as investigations of the impact of night cooling on the following day’s maximum inside temperature.

Blondeau *et al.* (1997), used a building in La Rochelle, France to investigate the impact of fan-driven night cooling on both the following day’s temperature and cooling load. The authors introduced the concept of “potential energy efficiency”, which is the ratio of the time integral of the heat exhausted by the fan divided by the total fan energy consumed over the same period. Naturally, this measure varied considerably from day to day since the inside-outside temperature difference also varied. The reported range for this metric was 2.5 – 7.5. Via simulations of the building, the researchers demonstrated the extent and the timing of the cooling loads offset by using MV at night.

2.2 Model and Simulations

The building load analysis software developed by Spindler (1998) was extended for the purposes of this project. The tool was selected because access to the source code enabled experimentation with a variety of cooling techniques and straightforward implementation of parametric studies. The hourly finite difference code, which accounts for variable equipment and occupant loads, solar loads (derived from TMY2 weather data (NREL)), variable ventilation rates and varying levels of thermal mass, provided estimates of annual sensible and latent cooling loads. Output from the simulation was used to determine overall cooling energy consumption, comprised of MV and AC fan energy, and AC chiller energy (sensible and latent).

The office space modeled for this project represents a single floor in a small office building with a carpeted slab floor and frame walls. The four exterior walls were glazed evenly, with windows covering 40% of the wall area. A summary of building specifications is provided in Table 2-1 and the associated load schedule is provided in Table 2-2.

Dimensions	9.64 x 9.64 x 3.05 m ³ (1000sf x 10')
Glazing	40% of wall area on all sides; U=1.42 W/m ² °C
SHGC (Solar Heat Gain Coefficient)	0.4 (reference case); 0.2 (with shading)
Ceiling	No mass; adiabatic
Floor	Carpeted (R=0.18 m ² °C/W) 10.2 cm concrete slab (adiabatic below)
Exterior Walls	12.7 cm sheetrock, R=1.76 m ² °C/W foam insulation, 12.7cm sheathing
Temperature Set Point	21 - 25°C (7am - 6pm)
Temperature Set Point	18 - 25°C (6pm - 7am)
Humidity Ratio Set Point	0.012 kg H2O/kg dry air (no minimum required)
Maximum Occupancy Load	10 occupants @ 70 W/person sensible; 40 W/person latent (includes some plants)
Maximum Equipment Load	32.3 W/m ² (3 W/sf)

Table 2-1 Building specifications.

Hours	Fraction of Maximum Load		
	Occupant Sensible	Equipment	Latent
12am - 7am	0.0	0.1	0.1
7am - 8am	0.3	0.3	0.3
8am - 9am	0.4	0.4	0.4
9am - 5pm	1.0	1.0	1.0
5pm - 6pm	0.4	0.4	0.4
6pm - 12am	0.0	0.1	0.1

Table 2-2 Building load schedule.

Mechanical equipment was modeled in a simplified manner. A chiller with air-side economizer and constant COP of 4 was used, as were constant specific fan powers (SFP) of 2.1 W/(l/s) (1 W/cfm) and 0.53 W/(l/s) (0.25 W/cfm) for the chiller fan and the MV fan, respectively. The lower SFP is a conservative estimate; for comparison, the fan installed in the building used for experiments for this research had an estimated SFP of 0.1 W/cfm.

As the goal of the project was to provide an initial indication of the potential of MM and MV in a large number of locations—each with different cooling system design requirements—modeling simplifications such as those mentioned above were deemed justifiable.

2.3 Mechanical Ventilation Strategy

The minimum airflow required for indoor air quality was provided at all times, including, for simplicity, unoccupied hours⁴. This rate was 1.2 air changes per hour (ACH). Whenever outdoor conditions were appropriate ($T_{out} < T_{in}$, and outside humidity ratio (W_{out}) \leq 0.012), the air exchange rate was increased to 10 ACH. The control strategy prevented ventilation cooling below the minimum temperature set point.

It was assumed that the ventilation system was designed to minimize pressure drops, perhaps using open windows as air intakes. With such a design, the low fan power cited above for the MV fan could be realized. The SFP for the chiller fan assumes a standard, ducted supply and return system. In a building designed to allow strictly natural ventilation (NV) in addition to MV, the effective SFP could be even lower than indicated above for the MV fan.

2.4 Mixed-mode Cooling Savings Potential

The savings potential of MM was assessed as follows. The MV strategy described above was used to maintain the inside temperature below the upper set point. When MV could not satisfy the sensible load, the windows were shut, the MV turned off, and 12.8°C air was supplied from the chiller at a flow rate commensurate with the load. These AC airflow rates and the constant AC fan SFP were used to estimate AC fan energy consumption. Chiller energy consumption was calculated using the total (sensible plus latent) load and the constant COP. It was also assumed that latent loads were removed during periods of MV as well as during periods of AC. The MM strategy was examined throughout the entire year. As noted by Kammerud *et al.* (1984), the effect of ventilation cooling on heat loads should be carefully examined. To minimize total

⁴ This wasteful strategy increases the total cooling energy consumption, which has the effect of reducing estimates of percentage savings due to MM.

(cooling and heating) energy consumption, the simple control strategy employed in this study to determine when to use MV should be re-examined.

2.5 Mechanical Ventilation Effectiveness

To investigate the effectiveness of MV alone, MV was used to maintain the room temperature as low as possible (above minimum specified set points)⁵. MV was selected for study rather than NV to avoid the difficulty and uncertainty of estimating natural airflow rates. The cooling effectiveness of MV represents an upper bound on that of NV.

A simple measure was used to assess comfort; a “comfortable” hour was defined to be an hour when $T_{in} \leq 27^{\circ}\text{C}$ and $W_{in} \leq 0.012$. The fraction of “comfortable” occupied hours during the cooling season was determined.

The cooling season for the MV-only study was defined for simplicity to begin at the beginning of the first week with average daily temperatures exceeding 15°C . The end was the beginning of the first week with average daily temperatures lower than 15°C .

2.6 Results and Discussion

Two sets of simulations were performed for 40 locations throughout the United States. The first set of simulations was designed to investigate the impact of MM on cooling energy consumption. The second set was designed to probe the effectiveness of using exclusively MV for cooling in the U.S.

2.6.1 Assessment of Savings due to MM

To illustrate the different components of the analysis, details have been provided for the case of the office space in Boston, MA.

As a reference for comparison, annual cooling energy consumption was predicted for the office space cooled by AC. Estimated annual cooling energy consumption was 42 kWh/m^2 . Next, the MM strategy was implemented for the same building. The annual cooling energy consumption under MM was 38 kWh/m^2 , reflecting savings of 9%. To investigate an upper limit of savings using MM, the SFP was set to zero (equivalent to NV). In this scenario, the savings due to MM rose to 31%, for an annual cooling energy consumption of 29 kWh/m^2 . Shading was then added to the reference building ($\text{SHGC} = 0.2$) to assess the relative benefit of reducing solar heat gains versus adding MM. The annual cooling energy consumption was 29 kWh/m^2 , 31% lower than the reference. Finally, the reference building was modified by adding both shading and MM. Total cooling energy savings were 34%, and the annual cooling energy consumption 28 kWh/m^2 . The origin of the MM savings is the availability of low-cost cooling via MV (or NV) during favorable periods. While the typical AC system with economizer maintains the space at the maximum set point, the MM system may decrease room temperature below the maximum set point with little energy penalty.

⁵ Energy consumption due to MV was not considered in this section.

Figure 2-1 shows the results of these same calculations for offices in Miami, St. Louis, Boston and San Francisco. Note that the use of MM may be inappropriate, as in Miami. In general, the effect of reducing the SHGC has a greater impact on cooling energy consumption than does the addition of MM. The benefit obtained from adding shading to the MM building is shown in the figure as well. Not shown is that further savings of approximately 10% (depending on climate) are attainable if additional thermal mass is integrated into the shaded, MM office space.

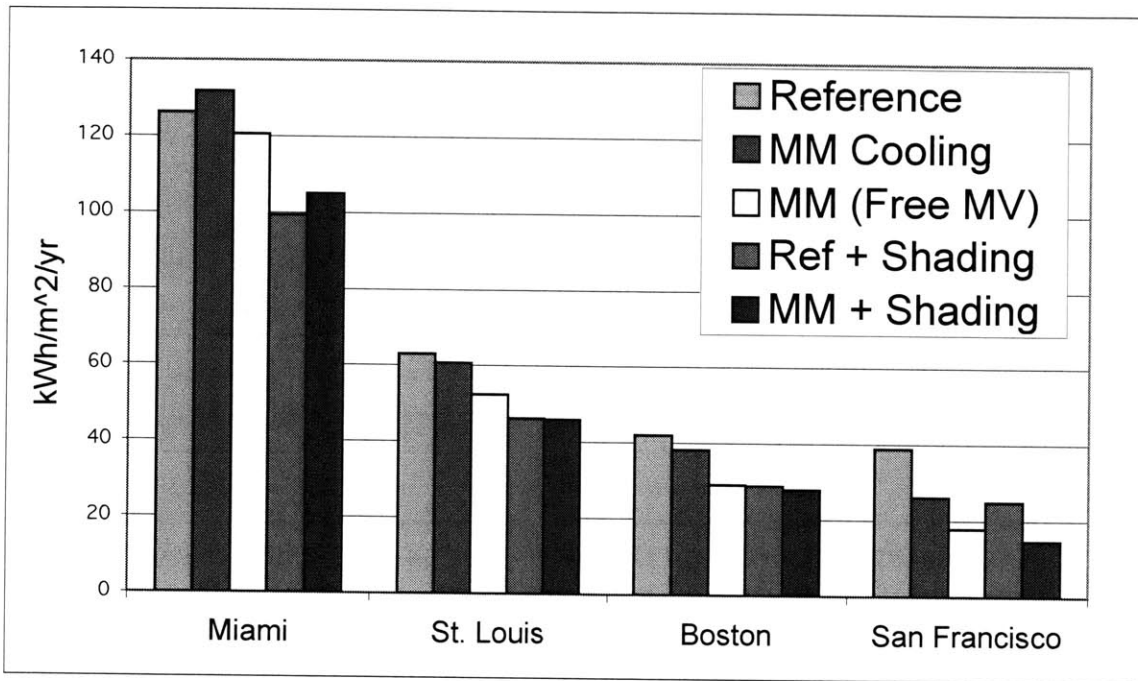


Figure 2-1 Annual cooling energy consumption in four U.S. cities. Plotted for each city are the annual cooling energy consumption of the reference building, the reference building using MM, the reference building using MM (with SFP = 0), the reference building with additional shading, and the reference building with MM and shading.

Several other representations of the potential savings due to MM are provided.

Shown in Table 2-3 is a complete listing of the 40 cities, displaying the cooling energy consumption of the building using AC and MM, as well as the relative savings due to MM and to MM with SFP = 0.

Shown in Table 2-4 is a summary of the benefits of adding MM to the reference building in all 40 cities examined. No shading was added. The range of percentage savings relative to the reference building is shown for each group, as is the group average savings of kWh/m². Cities in each group are listed in order of greatest to least savings.

City	State	Mode	Total Cooling kWh/m ² /yr	Savings	Savings with NV
ALBANY	NY	AC	42	8%	26%
		MM	39		
ALBUQUERQUE	NM	AC	67	7%	22%
		MM	63		
ANCHORAGE	AK	AC	15	47%	65%
		MM	8		
ATLANTA	GA	AC	72	3%	17%
		MM	69		
AUSTIN	TX	AC	99	0%	9%
		MM	99		
BISMARCK	ND	AC	42	10%	28%
		MM	37		
BOSTON	MA	AC	42	9%	31%
		MM	38		
CHARLESTON	SC	AC	82	2%	14%
		MM	80		
CHICAGO	IL	AC	47	5%	23%
		MM	45		
COLORADO_SPRINGS	CO	AC	47	14%	31%
		MM	40		
COLUMBUS	OH	AC	49	4%	25%
		MM	47		
CONCORD	NH	AC	40	12%	28%
		MM	36		
DES_MOINES	IA	AC	51	7%	24%
		MM	48		
DULUTH	MN	AC	27	22%	41%
		MM	21		
GRAND_JUNCTION	CO	AC	61	7%	21%
		MM	57		
GRAND_RAPIDS	MI	AC	41	9%	26%
		MM	37		
HARTFORD	CT	AC	46	9%	25%
		MM	42		
HELENA	MT	AC	38	16%	33%
		MM	32		
HONOLULU	HI	AC	125	-9%	5%
		MM	136		
HOUSTON	TX	AC	97	0%	8%
		MM	97		

City	State	Mode	Total Cooling kWh/m ² /yr	Savings	Savings with NV
INDIANAPOLIS	IN	AC	52	5%	21%
		MM	49		
LAS_VEGAS	NV	AC	102	3%	13%
		MM	99		
LEXINGTON	KY	AC	55	4%	22%
		MM	52		
LITTLE_ROCK	AR	AC	76	3%	15%
		MM	74		
MEMPHIS	TN	AC	78	3%	15%
		MM	76		
MIAMI	FL	AC	126	-4%	5%
		MM	132		
MINNEAPOLIS	MN	AC	44	7%	26%
		MM	41		
NEW_YORK_CITY	NY	AC	50	5%	26%
		MM	48		
OMAHA	NE	AC	55	5%	21%
		MM	53		
PHILADELPHIA	PA	AC	54	5%	22%
		MM	51		
PHOENIX	AZ	AC	120	0%	9%
		MM	120		
PORTLAND	ME	AC	36	16%	34%
		MM	30		
PORTLAND	OR	AC	38	16%	37%
		MM	32		
RALEIGH	NC	AC	67	5%	18%
		MM	64		
RICHMOND	VA	AC	64	5%	18%
		MM	60		
SAN_DIEGO	CA	AC	64	1%	34%
		MM	63		
SAN_FRANCISCO	CA	AC	39	32%	54%
		MM	26		
SEATTLE	WA	AC	31	22%	44%
		MM	24		
ST._LOUIS	MO	AC	63	4%	17%
		MM	61		
TOPEKA	KS	AC	61	5%	19%
		MM	58		

Table 2-3 Annual cooling energy consumption in 40 cities for the reference building using AC and using MM. Also shown are relative savings due to using MM and to using MM with SFP = 0 (“Savings with NV”).

>20% 8 kWh/m²	5% - 10% 3 kWh/m²	1% - 5% 2 kWh/m²
Anchorage, AK San Francisco, CA Duluth, MN Seattle, WA	Boston, MA Hartford, CT Grand Rapids, MI Albany, NY	Columbus, OH Lexington, KY St. Louis, MO Atlanta, GA
10% - 20% 6 kWh/m²	Minneapolis, MN Grand Junction, CO	Las Vegas, NV Memphis, TN
Portland, ME Portland, OR Helena, MT Colorado Springs, CO Concord, NH Bismarck, ND	Des Moines, IA Albuquerque, NM Richmond, VA Chicago, IL Philadelphia, PA Indianapolis, IN New York, NY Topeka, KS Omaha, NE Raleigh, NC	Little Rock, AR Charleston, SC San Diego, CA
		-10% - 0% -3 kWh/m²
		Phoenix, AZ Austin, TX Houston, TX Miami, FL Honolulu, HI

Table 2-4 Summary of annual savings due to the use of mixed-mode cooling in 40 U.S. cities.

The data in Table 2-3 were used in conjunction with Department of Energy survey results to estimate the national savings possible if commercial buildings were cooled with MM rather than AC. The cities shown in the Table 2-3 were grouped according to membership in the U.S. census divisions⁶, and the average annual savings per floor area were computed for each division. Using numbers for the total floor area of air-conditioned commercial buildings in each census division (EIA 1999), an estimate of national savings was computed for the case of MM and MM “with NV”. Percentage savings of cooling energy consumption were 4% and 19%, respectively. According to the Energy Information Administration (EIA) preliminary commercial building end-use consumption estimates for 1999⁷, the national on-site cooling energy consumption was 0.9 Quads, or approximately 2.7 Quads when primary sources are considered. Therefore, the primary energy savings potential due to MM and MM “with NV” in commercial buildings is estimated to be 0.1 and 0.5 Quads out of a national consumption (in 2002) of 97.4 Quads (DOE). These estimates should be treated with caution as they are based on an extrapolation of the performance of a single small office building to all of the commercial buildings in the nation. Furthermore, the census division building area data includes commercial buildings of all types, with the sole requirement that *part* of the building be air conditioned. Further studies are in order, preferably with simulations of multiple commercial building types.

Shown in Figure 2-2 is a geographical representation of the cooling energy savings potential due to MM and to shading. A set of ten representative cities were selected to display.

⁶ http://www.eia.doe.gov/emeu/cbecs/census_maps.html

⁷ http://www.eia.doe.gov/emeu/cbecs/enduse_consumption/intro.html



Figure 2-2 Geographical representation of the annual cooling energy consumption in kWh/m²/yr using AC, MM, AC with additional shading, and MM with additional shading.

As evidenced by Figure 2-1 and Table 2-3, if the building design were to allow for NV with flow rates comparable to those used in MV, then the savings due to MM rise dramatically: from 9% to 31% in Boston, from 4% to 17% in St. Louis, and from 32% to 54% in San Francisco.

The calculations described above reflect the use of both inside-outside temperature difference and humidity levels in the control decision of when to boost air exchange rates. If the outside humidity is not factored into the decision, then more hours of sensible cooling become available, but the latent loads that must be met by AC in the space may increase.

2.6.2 Assessment of Cooling Effectiveness of MV

The success of a cooling strategy necessarily depends on the expectations of the occupants. As some researchers have argued (Brager and de Dear 2000), individuals may adapt to and even prefer warmer and more humid conditions than ASHRAE comfort conditions would indicate (ASHRAE 1997) if they have some control over their environment and have direct access to abundant fresh air in their workspace. In a survey of occupants in an NV office space in the Boston area (the same where experiments for this thesis were performed), occupants' stated comfort levels generally exceeded those predicted by traditional comfort models (Chang 2002).

For these reasons, a space temperature slightly higher than typical was considered comfortable in this assessment. The criteria for comfort were defined to be temperatures at or below 27°C and humidity ratios at or below 0.012. The percentage of occupied hours during the cooling season that met these criteria was calculated for each city.

It was assumed that an office space designed to rely on MV alone for cooling would have increased levels of shading (reflected by a SHGC of 0.2) and increased thermal mass integrated into the space to take advantage of night cooling. The increased thermal mass was modeled here by removing the carpet from the slab floor and adding a 20.3 cm concrete masonry unit layer between the sheetrock and insulation in the walls.

Shown in Table 2-5 is a summary of the percentages of comfortable occupied hours in the office space cooled by MV. Only in San Francisco and Anchorage were at least 90% of the occupied hours during the cooling season comfortable. To provide some context for these results, the same calculations were performed for the office space located in London and Copenhagen (using Meteororm weather data (Meteotest)). In both cases, greater than 90% of the occupied hours were comfortable.

10% - 50%	50% - 70%	70% - 100%
Philadelphia, PA	Boston, MA	Anchorage, AK
Albuquerque, NM	Bismark, ND	San Francisco, CA
Grand Junction, CO	Colorado Springs, CO	San Diego, CA
Raleigh, NC	Portland, ME	Seattle, WA
Charleston, SC	Helena, MT	Duluth, MN
Little Rock, AR	Hartford, CT	Portland, OR
Atlanta, GA	Minneapolis, MN	
Memphis, TN	Chicago, IL	
Richmond, VA	Columbus, OH	
St. Louis, MO	Concord, NH	
Topeka, KS	Grand Rapids, MI	
Austin, TX	Albany, NY	
Las Vegas, NV	Des Moines, IA	
Houston, TX	Indianapolis, IN	
Miami, FL	Lexington, KY	
Phoenix, AZ	New York, NY	
Honolulu, HI	Omaha, NE	

Table 2-5 Percentages of comfortable occupied hours using exclusively MV for cooling.

The building design and its usage strongly influenced the cooling potential of MV. When equipment loads were reduced by 50%, and the SHGC reduced to 0.1, then the office in San Diego and Seattle could be comfortably cooled by MV for greater than 90% of occupied hours.

According to the measure used in this study, MV could be applied successfully in only a very small subset of U.S. cities. From a more positive perspective, however, the results indicated that AC was not required for a significant fraction of the occupied hours during the cooling season.

An additional approach to comfort assessment was inspired by the design targets for the BRE Environmental Office (Bunn 1997). The dry-bulb temperature thresholds exceeded by the inside temperature for 2% and 5% of the occupied hours during the cooling season were calculated. Analysis using these thresholds led to conclusions similar to those above about the effectiveness of MV and NV in most U.S. cities.

2.7 Conclusions

A prototypical office space was modeled in 40 different U.S. cities using three types of cooling strategies: standard air conditioning, mixed-mode cooling, and mechanical ventilation cooling. Total cooling energy consumption of AC and MM were calculated and compared to determine the potential benefits of implementing MM in the U.S. In addition, the effectiveness of cooling via mechanical ventilation alone was assessed.

The principal barrier to the introduction of mixed-mode cooling into the United States is evidently not climatic. It is likely economic, since the first cost of installing multiple cooling systems is often considered too onerous despite the potential life-cycle cost savings. In many U.S. locations, the use of MM rather than AC led to savings between 5 and 50% (and much greater if NV was available). These savings increased when additional building upgrades were

made. However, the majority of locations in the U.S. were found to be ill-suited for cooling exclusively by mechanical or natural ventilation.

This broad assessment of the applicability of MM and MV in the U.S. was feasible due to a number of modeling approximations. The purpose of the work was to give an indication of sites where the cooling techniques may be worthwhile. Further assessment using building and equipment models more carefully tailored to a particular building must be done to determine the ultimate benefits of MM or MV. In addition, a more in-depth analysis of comfort conditions may be warranted.

The ultimate success of the cooling strategies outlined in this chapter relies on the careful implementation of a control strategy to ensure that potential savings are realized. Herein lies the motivation for the work presented in the following chapters.

3 Modeling Natural and Mixed-mode Ventilation in Buildings

3.1 Challenges of Modeling Naturally Ventilated Buildings

3.1.1 Introduction

The purpose of this chapter is to provide the reader with insight into the complexity of modeling the thermal behavior of a naturally ventilated building as well as to propose a method for doing so. As will be shown, the building's thermal response is governed not only by the thermal properties of the building materials and weather conditions, but also by the configuration and size of all apertures throughout the building. As building apertures frequently open or close (presumably in a concerted fashion), the available paths for air to move through the building change accordingly. Naturally, the potential driving forces for airflow for a given set of temperatures are affected by the constellation of the different apertures. Conversely, flow through a building with unchanging apertures may vary in path and direction, as influenced by the temperature distribution within the building, the current outside temperature, and the prevailing wind.

Only the simplest flow scenarios can be addressed with first-principles, physically-based, analytical models. An abundance of analytical models have been proposed to describe natural ventilation driven by thermal buoyancy forces (stack effect), wind-induced forces, or combinations thereof (Linden *et al.* 1990; Allard and Utsumi 1992; Andersen 1995; Santamouris *et al.* 1995; Li 2000; Li *et al.* 2000; Allocca 2001; Leung and Li 2001; Li and Delsante 2001; Chen and Li 2002; Li 2002; Awbi 2003). Airflow in buildings with multiple zones is generally simulated using physically-based network models, (Feustel and Dieris 1992; Herrlin and Allard 1992; Dols and Walton 2002) or with computational fluid dynamics. Even with the simpler network models, the analytical relationship among the physical properties of the various zones is not available.

Since many naturally ventilated buildings do contain multiple zones, such as those with office areas linked to an atrium, it is helpful to study the analytical relationship between the different zones' configuration, temperature and the overall airflow predicted by a simple model. As will become clear in the next section, the governing flow equation for a simple case with purely buoyancy-induced flows is a nonlinear function of the outside temperature, all zone temperatures and building geometry.

Finally, to compound the difficulty of developing physically-based models for even simple building configurations is the difficulty of establishing *a priori* the appropriate parameters to use.

In this chapter, an analytical model for buoyancy-driven airflow in a building with three stacked zones is developed and used to underscore the complexity of modeling natural ventilation. Of particular interest is the nonlinear behavior accompanying flow reversal in a multi-zone building. This abrupt change in flow rates (and associated heat inputs to the zones) is difficult to build into a model unless one knows precisely the conditions that trigger the reversal. Also emphasized in this chapter is the difficulty of establishing and modeling *a priori* all of the different types of

flow regimes possible for all building aperture configurations. Several data-driven (as opposed to physically based, or analytical) modeling approaches are investigated in this chapter, and their performance in handling flow regime change is compared. One method, Principle Hessian Direction Regression Tree, is identified as the most capable of capturing such behavior. The method is described in detail below.

3.1.2 Buoyancy-induced Multi-zone Airflow

As will be described in more detail in the following chapter, the Broadmoor Wildlife Sanctuary building is comprised of three zones, stacked one above the other, and connected by apertures at the junctions. This section is devoted to the development of one set of buoyancy-induced airflow equations for such a space. Even with significant simplifying assumptions, the resulting equations display significant nonlinearity and complexity. A schematic of the building is shown in Figure 3-1.

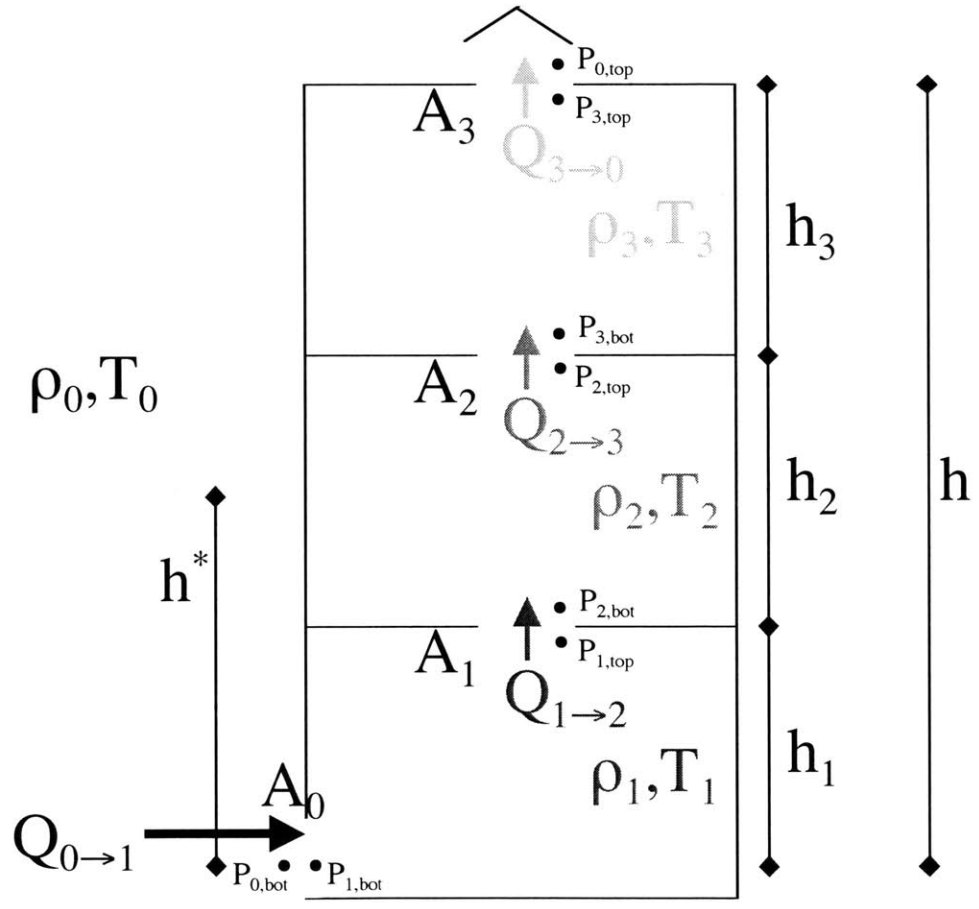


Figure 3-1 Schematic of a three-zone building.

In this figure, it is assumed that the outside temperature, T_0 , is cooler than the inside temperatures. Warmer temperatures are indicated by lighter shading. The temperature in each zone is assumed to be uniform. Wind is ignored. A discharge coefficient, C_d , is associated with each aperture, A_i , in the building: $C_{d,0}$ through $C_{d,3}$. The pressures on both sides of the various

apertures are labeled according to zone; *e.g.*, $P_{1,top}$ is the pressure at the top of zone 1 and $P_{1,bot}$ is the pressure at the bottom of zone 1, in which the temperature is T_1 . Flow from zone i to zone j is denoted $Q_{i \rightarrow j}$. h^* represents the height of the neutral pressure plane, where inside pressure equals outside pressure. The assumed location of h^* relative to h_1 and h_2 is not important. The flow equations derived below are unaffected by this choice.

In the following calculations, it is assumed that a monotonic temperature distribution exists (from zone to zone), and that the system is in steady-state operation.

The orifice equation may be written for each of the four orifices:

$$\begin{aligned}
 Q_{0 \rightarrow 1} &= (C_d A)_0 \sqrt{\frac{2(P_{0,bot} - P_{1,bot})}{\rho_0}} \\
 Q_{1 \rightarrow 2} &= (C_d A)_1 \sqrt{\frac{2(P_{1,top} - P_{2,bot})}{\rho_1}} \\
 Q_{2 \rightarrow 3} &= (C_d A)_2 \sqrt{\frac{2(P_{2,top} - P_{3,bot})}{\rho_2}} \\
 Q_{3 \rightarrow 0} &= (C_d A)_3 \sqrt{\frac{2(P_{3,top} - P_{0,top})}{\rho_3}}
 \end{aligned}
 \tag{3-1}$$

The pressures at the indicated locations may be written:

$$\begin{aligned}
P_{0,bot} &= \rho_0 g h^* \\
P_{0,top} &= -\rho_0 g (h - h^*) \\
P_{1,bot} &= P_{1,top} + \rho_1 g h_1 \\
&= P_{0,bot} - \left(\frac{Q_{0 \rightarrow 1}}{(C_d A)_0} \right)^2 \frac{\rho_0}{2} \\
P_{1,top} &= P_{1,bot} - \rho_1 g h_1 \\
&= P_{2,bot} + \left(\frac{Q_{1 \rightarrow 2}}{(C_d A)_1} \right)^2 \frac{\rho_1}{2} \\
P_{2,bot} &= \rho_2 g (h^* - h_1) \\
P_{2,top} &= -\rho_2 g (h_1 + h_2 - h^*) \\
P_{3,bot} &= P_{2,top} - \left(\frac{Q_{2 \rightarrow 3}}{(C_d A)_2} \right)^2 \frac{\rho_2}{2} \\
P_{3,top} &= P_{3,bot} - \rho_3 g h_3 \\
&= P_{0,top} + \left(\frac{Q_{3 \rightarrow 0}}{(C_d A)_3} \right)^2 \frac{\rho_3}{2}
\end{aligned} \tag{3-2}$$

where g is the acceleration due to gravity.

Substitutions are made into the expressions for $P_{1,top}$ to yield:

$$(Q_{1 \rightarrow 2})^2 = \frac{2[g h^* (\rho_0 - \rho_2) - g h_1 (\rho_1 - \rho_2)]}{\rho_0 \left[\frac{\rho_0}{\rho_1} \frac{1}{(C_d A)_1^2} + \frac{1}{(C_d A)_0^2} \right]} \tag{3-3}$$

Similarly, substitutions are made into the expressions for $P_{3,top}$ to yield:

$$(Q_{1 \rightarrow 2})^2 = \frac{2[g h^* (\rho_2 - \rho_0) - g (\rho_2 h_1 + \rho_2 h_2 + \rho_3 h_3 - \rho_0 h)]}{\rho_0 \left[\frac{\rho_0}{\rho_3} \frac{1}{(C_d A)_3^2} + \frac{\rho_0}{\rho_2} \frac{1}{(C_d A)_2^2} \right]} \tag{3-4}$$

Conservation of mass was used to derive Equations [3-3] and [3-4]:

$$\rho_0 Q_{0 \rightarrow 1} = \rho_1 Q_{1 \rightarrow 2} = \rho_2 Q_{2 \rightarrow 3} = \rho_3 Q_{3 \rightarrow 0} \tag{3-5}$$

Defining D_1 to be the denominator of Equation [3-3] and D_2 the denominator of Equation [3-4],

$$D_1 = \rho_0 \left[\frac{\rho_0}{\rho_1 (C_d A)_1^2} + \frac{1}{(C_d A)_0^2} \right]$$

$$D_2 = \rho_0 \left[\frac{\rho_0}{\rho_3 (C_d A)_3^2} + \frac{\rho_0}{\rho_2 (C_d A)_2^2} \right]$$
[3-6]

h^* is found to be:

$$h^* = \frac{[-g(\rho_2 h_1 + \rho_2 h_2 + \rho_3 h_3 - \rho_0 h)D_1 + g h_1 (\rho_1 - \rho_2)D_2]}{g(\rho_0 - \rho_2)(D_1 + D_2)}$$
[3-7]

Substituting Equation [3-7] into Equation [3-3] yields⁸

$$(Q_{0 \rightarrow 1})^2 = \frac{2g(\rho_0 h - \rho_1 h_1 - \rho_2 h_2 - \rho_3 h_3)}{(\rho_0)^2 \left(\frac{1}{\rho_0 (C_d A)_0^2} + \frac{1}{\rho_1 (C_d A)_1^2} + \frac{1}{\rho_2 (C_d A)_2^2} + \frac{1}{\rho_3 (C_d A)_3^2} \right)}$$
[3-8]

Since $\rho_i \propto 1/T_i$ if T is given in Kelvin and constant pressure is assumed, Equation [3-8] may be rewritten (\dot{m} is the mass flow rate, Q is the volumetric flow rate):

$$\dot{m}_{0 \rightarrow 1} = \rho_0 Q_{0 \rightarrow 1} = \rho_0 \sqrt{\frac{2gT_0^2 \left(\frac{h}{T_0} - \frac{h_1}{T_1} - \frac{h_2}{T_2} - \frac{h_3}{T_3} \right)}{\left(\frac{T_0}{(C_d A)_0^2} + \frac{T_1}{(C_d A)_1^2} + \frac{T_2}{(C_d A)_2^2} + \frac{T_3}{(C_d A)_3^2} \right)}}$$
[3-9]

Equation [3-9] is valid only when the numerator is positive. This is true when air flows into zone 1 from zone 0 (outside). Analogous calculations may be performed to find $Q_{1 \rightarrow 0}$, which is relevant when air enters zone 3 from above and exits from zone 1 to zone 0. These flow conditions occur when inside temperatures are cooler than the outside temperature. In that case,

$$\dot{m}_{1 \rightarrow 0} = \rho_1 Q_{1 \rightarrow 0} = \rho_1 \sqrt{\frac{-2gT_1^2 \left(\frac{h}{T_0} - \frac{h_1}{T_1} - \frac{h_2}{T_2} - \frac{h_3}{T_3} \right)}{\left(\frac{T_1}{(C_d A)_0^2} + \frac{T_2}{(C_d A)_1^2} + \frac{T_3}{(C_d A)_2^2} + \frac{T_0}{(C_d A)_3^2} \right)}}$$
[3-10]

⁸ It was found that Li *et al.* (2000) developed a similar expression, but without explicitly accounting for the density terms in the denominator. The authors established an expression for $(P_{3,top} - P_{0,top})$, then assumed an average ρ in the denominator to calculate flow rates.

and, again, Equation [3-10] is valid only when the numerator is positive.

Finally, it should be noted that the heat input to zone 1 due to airflow is given by:

$$\text{Heat Input} = \left\{ \begin{array}{l} \dot{m}_{0 \rightarrow 1} c_p (T_0 - T_1) : \left(\frac{h}{T_0} - \frac{h_1}{T_1} - \frac{h_2}{T_2} - \frac{h_3}{T_3} \right) \geq 0 \\ \dot{m}_{1 \rightarrow 0} c_p (T_2 - T_1) : \left(\frac{h}{T_0} - \frac{h_1}{T_1} - \frac{h_2}{T_2} - \frac{h_3}{T_3} \right) < 0 \end{array} \right. \quad [3-11]$$

where c_p represents the specific heat of the air. Recall that Equation [3-11] was derived for steady-state flow, with a monotonic temperature distribution in the building.

The pattern of Equations [3-9] and [3-10] may be extended readily to buildings with different numbers of stacked zones. In the simple case of one interior zone, $h_1 = h$, and Equation [3-10] reduces to the familiar form⁹:

$$Q_{0 \rightarrow 1} = \sqrt{\frac{2gh \frac{T_0}{T_1} |T_1 - T_0|}{\left(\frac{T_0}{(C_d A)_0^2} + \frac{T_1}{(C_d A)_1^2} \right)}} \approx \sqrt{\frac{(C_d A)_0^2 (C_d A)_1^2}{(C_d A)_0^2 + (C_d A)_1^2}} \sqrt{\frac{2gh |T_1 - T_0|}{T_1}} \quad [3-12]$$

where the absolute value in the numerator accounts for the two possible flow directions. In either case, Equation [3-12] denotes flow of air into zone 1, be it from above or below.

It is important to emphasize what Equation [3-11] represents. The thermal model of zone 1 has two distinct regimes, determined by the direction of airflow. With relatively cooler outside temperatures, cool air flows into zone 1, then up and out through zone 2. The relevant temperature difference (ΔT) is $(T_0 - T_1)$. In the other case, with relatively hot outside temperatures, the relevant ΔT is $(T_2 - T_1)$. In other words, zone 1 is heated by air from zone 2. The expression determining the direction of airflow is a function of all four temperatures and four height measurements. In essence, it is a mass balance equation: recognizing that $T^{-1} \propto \rho$, the sign of the expression in the inequality is determined by which column of air is heavier—outside or inside.

As a final note, it should be emphasized that we have examined a relatively simple form of airflow in the building. Recall that for the system examined, wind was ignored. Consider also the case where aperture A_1 in Figure 3-1 is closed. In this case, simple, physically-motivated analytical equations can no longer be written to describe the flow of air in zones 2 and 3 (even under no-wind conditions). Flow in zone 1 may be described by single-opening equations, such as in (Allocca 2001). Recognizing that the apertures depicted in Figure 3-1 may be composite apertures, potentially representing multiple apertures in the same zone, perhaps at multiple

⁹ The ratio T_0/T_1 is frequently set to 1.

heights, it becomes clear that model complexity will increase, and that multiple models are required for different patterns of aperture openings.

In the next section, a representative two-zone building will be used to explore the ability of several types of models to capture the type of behavior predicted by Equation [3-11].

3.1.3 Data-driven Models of Buoyancy-induced Flow in a Simplified Two-zone Building

3.1.3.1 Problem Statement

In this section, a two-zone model is further simplified to illustrate a fundamental modeling challenge. The goal is to develop a data-driven model of a system where the magnitude of the flow rate is determined solely as follows (*c.f.* Equation [3-12]):

$$Q \propto \sqrt{\frac{|T_1 - T_0|}{(T_1 + T_0)}} \quad [3-13]$$

The heat input to the space is given by:

$$\text{Heat Input} \propto \begin{cases} Q(T_0 - T_1) : T_0 \leq T_1 \\ Q(T_2 - T_1) : T_0 > T_1 \end{cases} \quad [3-14]$$

The scenario is contrived (note that T_2 plays no role in determining the flow rate), but helps to illustrate an important phenomenon.

3.1.3.2 Models for “Heat Input”

Two types of models were constructed from data generated from Equation [3-14]. In both cases, three sets of data were generated as follows¹⁰. A 3000-point training set and a 3000-point test set were generated, with T_0 uniformly and randomly generated from the set [290, 300] K, T_1 from the set [290, 300] K, and T_2 from the set [280, 310] K. A third dataset of 500 points was generated, with $T_0 = 297$ K, $T_2 = 308$ K, and T_1 evenly spaced on the interval [290, 300] K. The corresponding “Heat Inputs” were calculated for all points. All data points were normalized to the range [0, 1] before models were constructed. All figures and numbers shown in the text represent un-normalized temperatures and “Heat Input”.

3.1.3.2.1 Linear Regression Model

The simplest model to be constructed was a linear regression model (using singular value decomposition adapted from Wunsch (1996)) with the three normalized temperatures as inputs and the normalized “Heat Input” as the model output. In case a), all training data were used to construct a single model to fit the data. In case b), the training data were divided into two portions, with one set comprised of points where $T_0 \leq T_1$ and the other of points where $T_0 > T_1$. The intent was to improve overall model performance by giving one sub-model the task of

¹⁰ See FlowRegimeTest.m

modeling the system when cool outside air enters zone 1, and the other sub-model the task of modeling the system when hot air enters zone 1 from zone 2. Shown below in Figure 3-2 are two plots of the linear regression models tested on the third dataset.

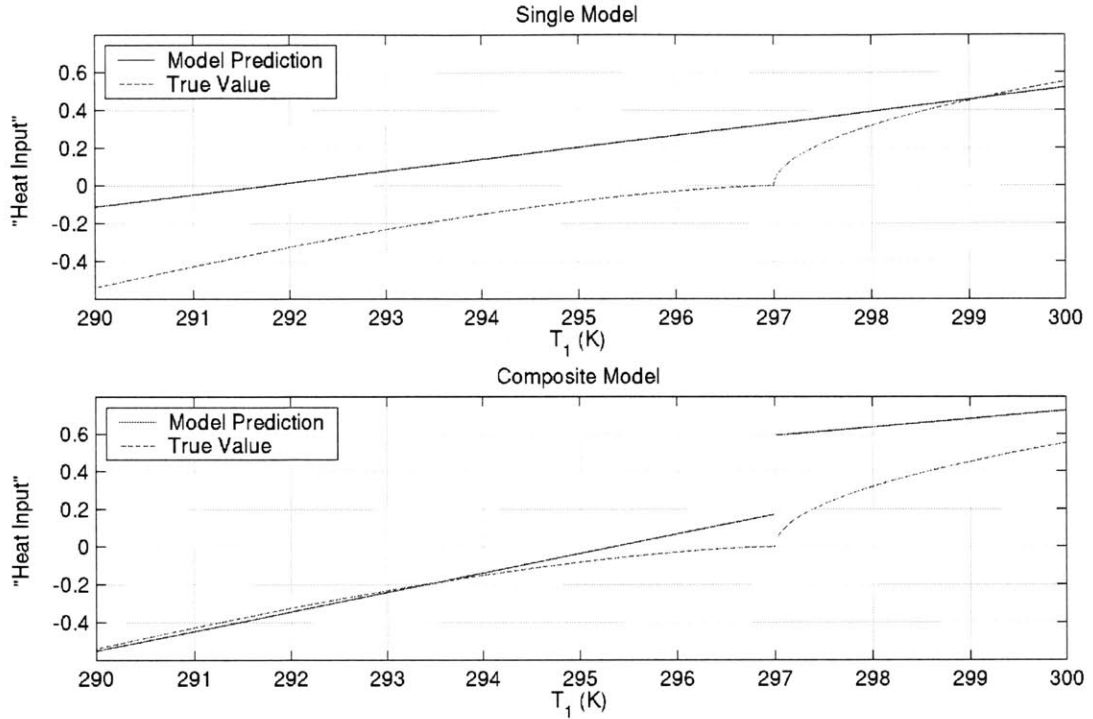


Figure 3-2 “Heat Input” versus T_1 for case a (top) and case b (bottom).

The performance of the two models (cases a and b) on the training and test sets is shown in Table 3-1. Note that the curves shown in Figure 3-2 represent one slice of the 4-dimensional space of the problem. The performance data shown in Table 3-1 were gathered over the entire 4-dimensional problem space. Note that the size of the root-mean-square error (rms error) is approximately halved when two sub-models are used for the two flow regimes (case b) rather than a single model for both regimes (case a).

	Case a: Regression, 1 model	Case b: Regression, 2 models	Case c: KRLS, 1 model	Case d: KRLS, 2 models
Training rms error	0.27	0.15	0.015	0.004
Test rms error	0.27	0.16	0.022	0.013
Training time (s)	0.41	0.06	1411	473

Table 3-1 Performance of four model types on data generated from Equation [3-14]. Cases a and b reflect linear regression models, while cases c and d reflect models made with a form of support vector regression, KRLS (see discussion below).

Examination of the form of Equation [3-14] quickly reveals why a simple linear regression performs as poorly as shown above. If one knew *a priori* the functional form of the generating function, or appropriate basis functions, then one could still approach this problem using linear regression to advantage. Model inputs would no longer be simple temperatures, but rather composite terms, such as

$$\sqrt{\frac{|T_1 - T_0|}{(T_1 + T_0)}}$$

or

$$\sqrt{\frac{|T_1 - T_0|}{(T_1 + T_0)}}(T_0 - T_1) \quad [3-15]$$

Unfortunately, such basis functions are not known, due to the presence of unknown geometry parameters and other factors. Also, the functional form of the basis functions would have to be tailor fit to the particular type of building zone and all the flow regimes it encounters (some influenced by wind). Since such an approach is impractical, a more general approach is considered.

3.1.3.2 Kernel Recursive Least Squares Model

A wide variety of black-box models exist that are capable of learning a nonlinear mapping of a set of inputs to a set of corresponding outputs. Radial basis functions (RBF), neural networks (NN), generalized regression neural networks (GRNN) and fuzzy inference systems (FIS) are all such modeling tools that have been used successfully in various building modeling and control applications, *e.g.*, (Kreider and Haberl 1994; Ahmed *et al.* 1996; Bruant *et al.* 2001; Ferreira and Ruano 2002). Another class of mapping function that has been investigated as part of this research is kernel recursive least squares (KRLS) (Engel *et al.* 2002; Engel *et al.* 2004; Engel *et al.* Unpublished). Their work is an extension of the underlying framework of support vector machines (SVM) (Vapnik 1995), which are used for classification, and the associated regression analog, support vector regression (SVR), *e.g.*, (Gunn 1998; Smola and Schoelkopf 1998). In addition to being a recursive algorithm, KRLS offers a sparse, or compact, model representation relative to SVR. SVR and KRLS offer better generalization properties than tools such as NN's without the difficulty of choosing the appropriate architecture of the NN.

Kernel-based approaches to classification and regression take advantage of the flexibility and computational reduction of generating and indirectly manipulating high-dimensional feature vectors in Hilbert space, H , by performing calculations on vectors in the input space¹¹. By the input space is meant, for example, the measurements taken in an experiment. The feature space then represents all combinations of those input variables that could serve as useful inputs to a model. It is the kernel function that generates these combinations automatically. The simplest common kernel function is the dot product, represented by $K(x',x) = \langle x', x \rangle$. Polynomial kernels, such as $K(x',x) = \langle x', x \rangle^d$, where d represents the degree of the polynomial, are also

¹¹ Much of the content of this discussion has been drawn from (Engel *et al.* 2004) and (Gunn 1998).

common. They are sometimes implemented as $K(x',x) = (1 + \langle x', x \rangle)^d$. The kernel function similar in nature to RBF's is the Gaussian kernel: $K(x',x) = \exp(-\|x' - x\|^2/(2\sigma^2))$.

Given a set of N data points and associated targets, y , the problem can be reduced to a linear regression in Hilbert (or kernel) space:

$$y = \sum_{i=1}^N \alpha_i k(x_i, x) \quad [3-16]$$

One simply must solve for the parameters, α_i . However, with a large dataset, the computational burden of computing the kernel for all N points becomes excessive. The primary contribution of Engel *et al.* (2004) was to develop a method for selecting only a subset of the N data points for use in calculating the kernels in such a way that the minimal information is lost when the other points are discarded. Rather than using Equation [3-16] involving all N points, a comparable model could be formed:

$$y = \sum_{j=1}^{m < N} \alpha_j k(x_j, x) \quad [3-17]$$

where m is the chosen subset of points.

To select the m points in the subset of N points, the authors presented a method of comparing the feature vector of each new point obtained with that of all foregoing points in the “dictionary” to establish whether it is approximately linearly independent of the feature vectors of the “dictionary” points. The tolerance for determining approximate linear independence, v , is set by the user. If the new point’s feature vector is approximately linearly independent, the point is added to the “dictionary”. If not, the point can be discarded (after adjusting some model weights).

The authors have embedded this concept into an online recursive least squares framework, so that as new data points are obtained, they either serve to update the model by adjusting parameters, or to enlarge the dictionary. This approach leads to a very sparse model, with far fewer than N point required to form accurate predictions.

The parameters available to the user are determined by the selection of the kernel. For a polynomial kernel, one has just to select the degree, d . For a Gaussian kernel, selected in this work, one must select both σ , which governs the spread of the Gaussian (as it does for RBF's), and v , the tolerance for admission to the “dictionary”. The smaller the v , the greater the number of “dictionary” members required by the model and the greater the computational load imposed on the algorithm. Similarly, the smaller the σ , the finer the model, and the greater the number of points required to generate the model.

The authors demonstrated the performance of the KRLS algorithm on two well-known time-series tests: the Mackey-Glass time series (Mackey and Glass 1977), and the Santa Fe Laser time series (Weigend and Gershenfeld 1994). In the former case, the KRLS algorithm performed

comparably and better than an SVR package, SVMtorch (Collobert and Bongio 2001), and in the latter, it outperformed the winning solution.

The KRLS algorithm was implemented for this research in MATLAB and used to model the same dataset described above at the start of Section 3.1.3.2, first to generate a single model (case c) and second to generate a composite model (case d) analogous to case b, above. The parameters σ and ν were set to 0.05 and 0.01 respectively. Model performance could have been improved with smaller parameters, but the calculation penalty would have become excessive. The performance of the KRLS models is superior to that of the linear regression models; see Figure 3-3 and Table 3-1. Examination of the table reveals that the accuracy of the KRLS model greatly exceeds that of the linear regression model (with a significant calculation time penalty). Furthermore, it is clear that the use of sub-models (case d) increases accuracy and reduces calculation time relative to case c. The total number of support vectors (or “dictionary” members) is approximately the same for cases c and d.

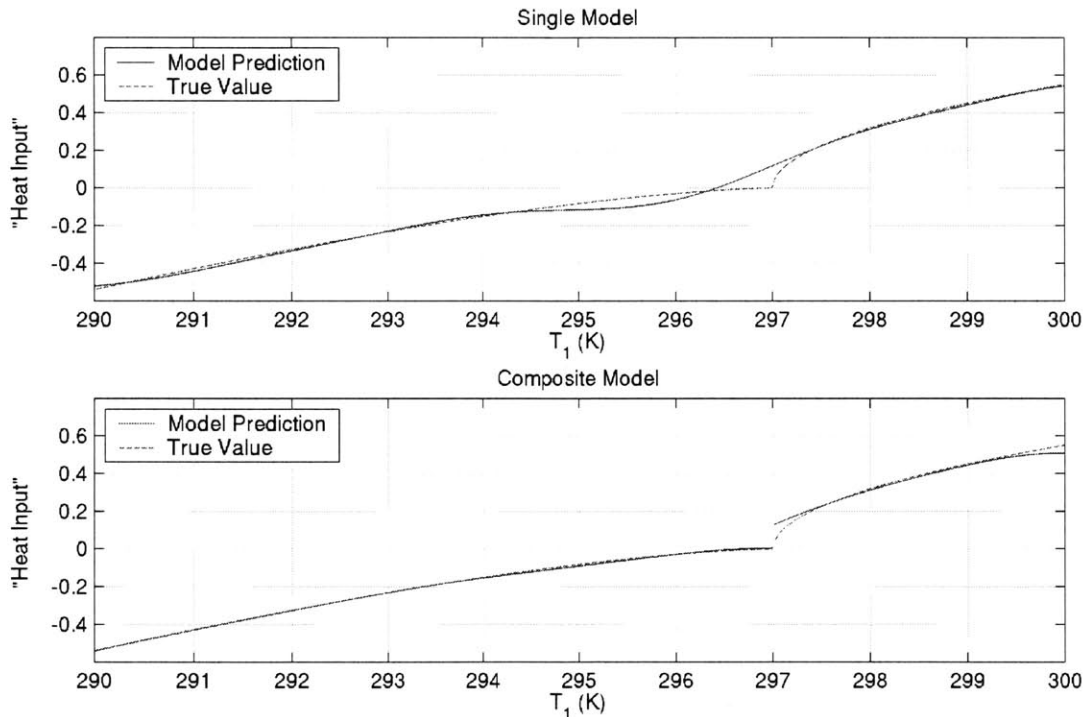


Figure 3-3 “Heat Input” versus T_1 for case c (top) and case d (bottom).

3.1.3.2.3 Model Assessment

As discussed above, the accuracy of the nonlinear KRLS model exceeds the linear regression model by approximately an order of magnitude. Linear regression performance would improve dramatically if appropriate basis functions could be supplied (possible in this simplified case, but not in general). This conclusion indicates that a nonlinear model that maps inputs to outputs is required to obtain high accuracy. Furthermore, there is a clear benefit to dividing the dataset into regimes and fitting a model for each regime—rms errors are reduced by nearly a factor of two by doing so.

The obvious questions to follow are: which nonlinear mapping should be used, and how can the data be segmented into different regimes? The first is a matter of balancing computational load, model accuracy and generalization. The second is more confounding. In this example, it was possible to use a strict comparison between the T_0 and T_1 to segment the data. In a more realistic scenario, the flow regime depends on geometry and temperatures in multiple zones. The use of FIS models was explored, but the antecedents (rules) must be based on individual model inputs, such as:

IF (Input 1 \geq 2 AND Input 2 $<$ 1) THEN use model i
rather than
IF (Input 1 \geq Input 2) THEN use model j

Of course, the inputs could be combined so that Input 1' = (Input 1 – Input 2), rendering the previous rule equivalent to:
IF (Input 1' \geq 0) THEN use model j.

However, such a manipulation requires the user to know exactly which inputs must be combined, and in what strengths (*c.f.* Equations [3-11]). If that were known, then there would be no need to use the FIS architecture—the modeler could directly partition the data, then build models for each segment. A unified approach that answers both questions posed above will be introduced in Section 3.3 below.

3.2 Challenges of Modeling Complex Naturally Ventilated or Mixed-mode Buildings

The issues of demarcating flow regimes are compounded by the fact that when apertures in a more complex building are opened or closed, or when an exhaust fan is turned on to boost natural flows, the airflow pattern changes. Such changes may necessitate entirely different models. For example, the thermal model of zone 1 under the conditions shown in Figure 3-1 may be very different from the model obtained for zone 1 when aperture A_1 is shut (see Figure 3-4). Note that the models for zones 2 and 3 may also change dramatically.

Successful modeling of a complex naturally ventilated building (defined by the presence of multiple operable apertures) or a mixed-mode building requires the ability to handle multiple control modes—or modes of operation—and perhaps two or more flow regimes within each control mode. Recall that the flow regime was defined by the overall pattern of airflow; in Equations [3-9] and [3-10] two regimes were present: one with upward flow and the other with downward flow. The control mode establishes the physical configuration of the space through which the air may flow, perhaps restricting flow patterns (such as was the case in Figure 3-4 when A_1 was shut). Once the physical constraints are determined by the control mode, the current temperature distribution (and possibly wind speed and direction) dictates the flow regime.

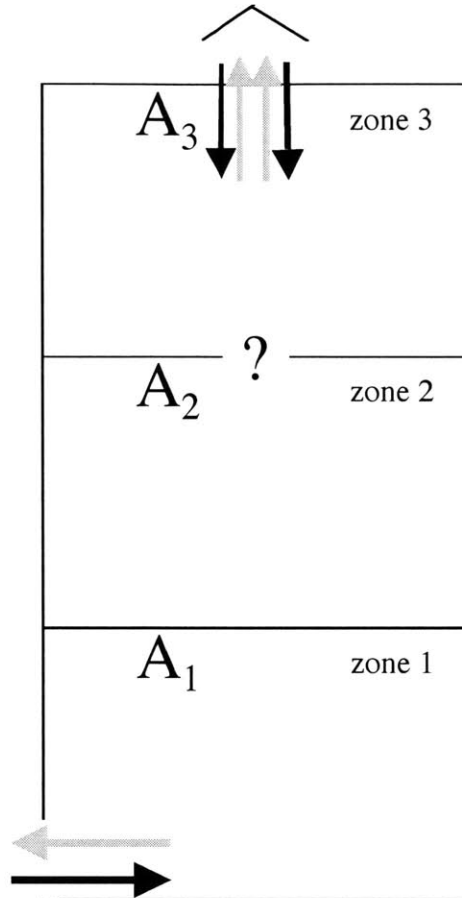


Figure 3-4 Three-zone building with aperture A_1 closed. Warm air is shown in grey, cool in black. The flow pattern through A_2 is not known *a priori*.

Since the control mode is set (and therefore known) by the building operator, it is reasonable to partition the data directly by control mode. However, with N independent apertures, the number of control modes $\propto 2^N$ if the controls are assumed to have binary OFF/ON or OPEN/SHUT settings. For the modeling problem to be tractable, the number of control modes must be restricted, perhaps by grouping a set of apertures together. It may also be that the operation of certain apertures does not significantly alter the flow patterns, so their operation may not require the introduction of a distinct mode. These subtleties will be explored in Chapter 5.

3.3 Proposed Solution

To confront the modeling challenges described in the previous sections, a search of existing modeling approaches was undertaken, with the focus on methods that could handle (and identify) abrupt flow regime changes where the gradient of the modeled function changes discontinuously. A natural solution to examine is one where the variable space is divided into subspaces, each associated with a sub-model for that subspace.

3.3.1 Survey of Methods

Much work has been devoted to finding methods to cluster similar data together. These clusters may be used to simplify modeling (by reducing computational load) or as a first step of

generating a set of local models. These can take the form of crisp clusters, as generated by k-means clustering, *e.g.*, (Moody and Darken 1989), or the simple one-pass method proposed by Specht (1991), or fuzzy clusters, such as c-means clustering, subtractive clustering, ellipsoidal clustering, *e.g.*, (Buckles *et al.* 1994; Chiu 1994; Chiu 1994a; Abe and Lan 1995; Babuska and Verbruggen 1996; Chiu 1996; Linkens and Chen 1998). The primary drawback of all of these methods is that membership in a cluster is based upon proximity in the input space rather than proximity in the output space. Specifically, when a multi-dimensional surface contains an abrupt change in slope (as occurs in the examples of airflow given above), points defined by similar temperatures may, in fact, belong to different flow regimes and therefore should not be clustered together.

Classification and regression trees (CART), presented in (Jang *et al.* 1997), were identified as a possible solution to the problem of inappropriate clustering of proximate data points. As described by the authors, each subspace is defined by axis-orthogonal boundaries, and each sub-model is a constant throughout the region defined by the boundaries. The boundaries are positioned to maximize the accuracy of the model predictions. A higher-fidelity model is generated as the number of subspaces is increased. Ultimately, a piecewise-constant approximation of the function over the entire space is obtained.

Loh presents a more flexible technique, (Loh 2002a; Loh 2002b), wherein the sub-models are linear combinations of all inputs rather than constants. This work is extended by Chauduri *et al.* (1994), who employ a piecewise polynomial fitting for each sub-model.

Jordan (1994) presents an alternative approach using soft boundaries to generate a mixture-of-experts model (*i.e.*, one incorporating and assimilating predictions from local models). As above, though, soft boundaries can lead to poor performance at regime edges.

The principal drawback to the CART algorithms mentioned above is that they require that the “cuts” between subspaces be axis orthogonal. Given a bimodal function described by two different functions on either side of the line $y = x$, the techniques above would require a very high resolution (*i.e.*, a high number of cuts) to capture the change along the $y = x$ line. Recall the scenario described in Equation [3-11] or [3-14] where the flow regime boundaries were not axis-orthogonal. While this problem could be avoided by redefining the axes, the user may not have the necessary information to do so. A more general approach combines the features of the methods mentioned above with an algorithm for determining how to transform the axes automatically to allow the most efficient “cutting” of the variable space, (Lue 1994; Li *et al.* 2000)¹². This method is called Principal Hessian Direction Regression Tree.

¹² The author of this thesis also developed an alternative method of automatically dividing the full space into convex hull subspaces, then performing a linear regression on each subspace. This unpublished work was found to be successful for problems with low input dimension (≤ 2 inputs to the model) and with arbitrary discontinuities, but it was not computationally feasible for systems with the large number of inputs required in the modeling described later.

3.3.2 PHDRT Theory

A brief review of this approach, called Principal Hessian Direction Regression Tree (PHDRT), is provided here. The essence of the technique is to identify the location of the hyperplane that cuts the data space perpendicular the direction of greatest curvature. Each subspace is cut in turn to render the subspaces more and more planar. The procedure terminates when the number of data points within a subspace falls below a pre-defined level or the curvature of the subspace becomes negligible. The points in each subspace are used to form a linear regression model for that subspace. The contribution of the authors cited above was to provide a method that automatically identifies the hyperplane used to cut the data space and to provide a statistical test for the existence of curvature of a subspace. Both methods were modified in the current research to improve modeling robustness and accuracy. The proposed modifications will be presented after the original work has been described.

The curvature of a function $y = f(x)$ is given by the Hessian matrix of the function:

$$H(x) = \frac{\partial^2 f(x)}{\partial x^2} \quad [3-18]$$

where the dimensions of the $H(x)$ are p by p , where p is the dimension of x . $H(x)$ (weighted by the covariance matrix of x) can be decomposed into eigenvectors and eigenvalues. The eigenvectors associated with the largest magnitude eigenvalue define the direction of greatest curvature of $f(x)$. The weighting by the covariance matrix of x is performed to remove dependence on the particular affine transformation of x . Identification of $H(x)$ by step-wise estimation at each point is computationally prohibitive in problems of high dimension. The authors drew on Stein's Lemma, see (Stein 1981; Li 1992) to avoid any direct computation of $H(x)$. Rather than using $H(x)$ to compute its eigenvalues and eigenvectors, they demonstrated they were identical to the eigenvalues and eigenvectors found using the following equation:

$$\Sigma_x^{-1} \Sigma_{rxx} b_i = \lambda_i b_i \quad [3-19]$$

where Σ_x is the covariance matrix of x , b_i the eigenvectors, λ_i the eigenvalues ($i = 1 \dots p$), and

$$\Sigma_{rxx} = \frac{1}{n} \sum_{j=1}^n r_j (x_j - \bar{x})(x_j - \bar{x})' \quad [3-20]$$

and

$$r_j = y_j - \hat{y}_j$$

n is the number of data points available, and \hat{y} is the least-squares estimate of y .

To determine if eigenvector b_i is significant, Theorem 4.2 in (Li 1992) is used to calculate the associated p -value. It is assumed that the λ_i have been sorted in order of decreasing absolute value.

$$p_i = 1 - X_{(p-i)(p-i+1)}^2 \left(\frac{n \sum_{j=i}^p \lambda_j}{2 \sum_{j=1}^n r_j^2} \right) \quad [3-21]$$

(Note that p with a subscript refers to the i^{th} p-value, while p with no subscript refers to the dimension of the input x .)

The p-value reflects the probability that the null hypothesis is true (Portney and Watkins 1993, pg. 349), *i.e.*, that the associated λ_i is zero and that no significant directions exist. For instance, a common threshold for significance (used by Lue (1994), pg. 11) is p-value = 0.05, which indicates that there is less than a 5% chance of the null hypothesis being true.

Once the eigenvector associated with the lowest p-value has been found, the data are projected along this direction to form the transformed variable PHD_1 . What remains is to determine the location along PHD_1 where the dataset should be split, or cut. The method suggested by the authors is outlined in detail in (Li *et al.* 2000). Their algorithm will not be described fully here, as a different method is proposed. In essence, the authors search indirectly for the point yielding a minimal weighted rms error (for the two sides of the cut). This is accomplished by assuming that both sides of the cut will have the same rms error, and that the cut position, when adjusted to meet that goal, will also yield the lowest overall rms error. If the overall rms error falls below a user-specified threshold, the algorithm terminates.

The PHDRT algorithm is summarized here, and is illustrated below.

1. Start with entire dataset.
2. With current points, perform linear regression and calculate residuals. If the rms error falls below a user-defined threshold, proceed to Step 6.
3. Use residuals of all current points to calculate PHD_1 .
4. Check that the p-value of the PHD_1 is ≤ 0.05 . If so, the direction is significant. Proceed to Step 5. If not, no significant curvature exists. Do not subdivide the current points, but rather find the best linear fit. Proceed to Step 6.
5. Locate the cut position (as outlined above, or as proposed below). Generate two subsets of data.
6. Select any remaining subset of data whose number exceeds a minimum set by the user and return to Step 2. If none exists, proceed to Step 7.
7. Terminate modeling.

To evaluate the model's output at a given data point, determine to which subset of points it belongs (as defined by the cuts) and calculate the output using the local regression equation.

Several comments are warranted regarding the technique as outlined so far. Equations [3-19] and [3-21] are strictly valid only for cases when the regressors form a normal distribution. Li *et al.* (2000) argue that the bias introduced when these conditions are modestly violated does not prevent the technique from remaining useful. A strong violation of these conditions does occur

when the regressors are binary, which would be the case if control modes of the building are used as regressors rather than as a tool to divide the dataset before applying a modeling tool such as PHDRT. The application of the PHDRT algorithm to such cases is explored in (Cheng and Li 1995; Filliben and Li 1997).

3.3.3 Proposed Changes to the PHDRT Algorithm

Several changes to the algorithm are proposed here, and are implemented in the calculations in Section 3.3.4.

3.3.3.1 Modified p-value Significance Criterion

It was found that datasets with obvious curvature were associated with eigenvectors and eigenvalues that had p-values > 0.05 . Rather than terminate the algorithm based on the p-value criterion alone, it was found to be more effective to terminate only when the p-value was high (> 0.75) and the full range of the residuals (max-min) exceeded a threshold. The threshold used in the examples below was 0.25. This approach couples a check for statistical significance with the requirement that if large residuals exist, the data must be further subdivided to improve the fit.

3.3.3.2 Automatic Dimension Reduction

Whenever a regression is performed, the user has the option of calculating the $1-\sigma$ confidence intervals for each regression parameter. If the ratio of each parameter's $1-\sigma$ confidence interval to its itself exceeds 1, then those parameters are fixed at zero and all other parameters are recalculated.

An analogous technique may be adopted when the eigenvectors of the weighted $H(x)$ are calculated. The elements of each eigenvector are divided by the element with largest absolute value. Any elements below a user-defined tolerance, *e.g.*, 0.01, are eliminated.

In the calculations in section 3.3.4, the latter method of dimension reduction was employed. Comments will be made regarding the effectiveness of these techniques in Section 3.3.4.4.

3.3.3.3 Alternative Cut Locator¹³

Rather than searching indirectly for the best position to place the cut, it is proposed that a direct (constrained) binary search be performed. The following method was implemented in the calculations illustrated in Section 3.3.4.

1. Locate the positions of the maximum and minimum allowable cuts $\equiv c_{\text{extreme}}$ (dictated by minimum set size).
2. Locate five additional equally-spaced trial cuts within the range defined by step 1. Label the cuts 1 through 5, from minimum to maximum.
3. Perform regression for the data points on both sides of each cut and find the associated overall rms error (weighted by the number of points on each side). Identify the cut with lowest associated rms error $\equiv c^*$. If the label for this cut, i , is

¹³ See MySplitLineSearch2.m

- 1 or 5, proceed to Step 4. Otherwise, the bounding cuts are $i-1$ and $i+1$. Proceed to Step 5.
4. Identify a cut between c^* and the adjacent c_{extreme} such that this cut is surrounded on both sides by cuts associated with higher rms errors. This new cut becomes c^* . (If no such cut exists, c_{extreme} is used as the final cut, c^* . Proceed to Step 7.)
 5. Perform binary search within the bounding cuts. Adjust bounds as necessary, always requiring that c^* be bounded by cuts associated with higher rms errors.
 6. When c^* is changed and the membership of the two subsets of data points does not change, proceed to Step 7. Otherwise, return to Step 5.
 7. Terminate. c^* has been identified.

3.3.3.4 Alternative Error Criterion

Adjustments were made so that the user may choose to use the criterion of maximum absolute error rather than that of rms error.

3.3.3.5 Additional Proposed Improvements

Several other proposed improvements have been noted, but not implemented, namely:

1. Provisions could be made to identify multiple cuts along PHD_1 simultaneously.
2. The entire procedure could be made iterative. It is often found that errors in cut placement (due to noise, for example) generate apparent outliers later in the process. If those outliers could be used in a feedback mechanism to influence earlier cut positions (and eliminate the existence of apparent outliers), the overall accuracy of the algorithm could be improved.
3. It was found that when either point symmetry or axial symmetry existed in the dataset (when the model input dimension > 1), it was impossible to identify significant eigenvectors. Further work should be done to characterize and remedy this problem. Demo 2 in Section 3.3.4.2 is a case in point. The bounds of the input data had to be shifted to break the symmetry. Note that in many interesting physical problems, this symmetry and associated problems will not exist.

3.3.4 PHDRT in Practice

The sections under this heading are included both to demonstrate the PHDRT algorithm and also to compare its performance with that of the KRLS algorithm. For additional examples, including some using experimental data, see (Lue 1994; Li *et al.* 2000).

3.3.4.1 Demo 1¹⁴

The first example was designed to illustrate the steps of the algorithm. The function to be learned by the model is:

¹⁴ Refer to MainPHDRTdemo1.m

$$z = \begin{cases} x: 1 \leq x \\ 1: -1 \leq x \leq 1 \\ 2+x: x \leq -1 \end{cases} \quad [3-22]$$

The function is plotted in the upper portion of Figure 3-5. x is drawn from a uniform random distribution between $[-3, 3]$, with 800 training points and 200 test points. Note that z was not corrupted by any noise. The changes described in Section 3.3.3 were implemented. The abscissas of the plots are labeled PHD_1 since there is only one direction/dimension present in the problem. The lower plot contains the residuals of the single regression model fit to all data points.




Figure 3-5 Top: Function to be modeled and the best linear fit (shown with thin dashed line). Bottom: Residuals based on linear regression using all points.

Figure 3-6 contains the same information as Figure 3-5, but includes also the search for the best cut with which to divide the entire dataset.

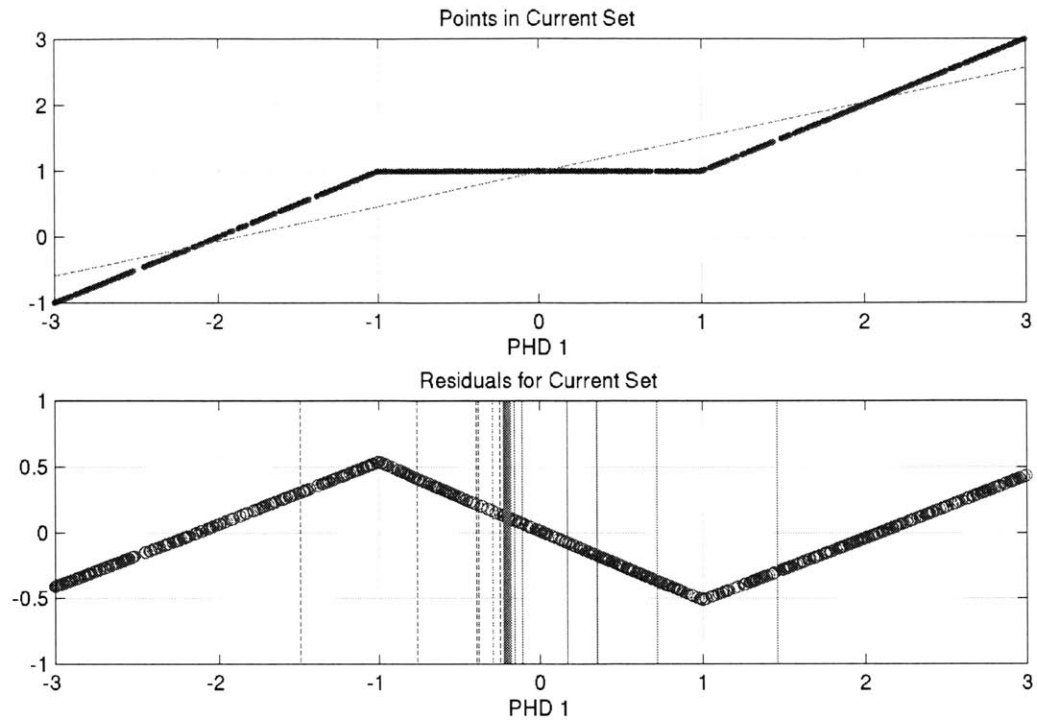


Figure 3-6 Same as Figure 3-5, but showing the search for the best cut to split the data along the PHD_1 direction. Solid vertical lines are right-hand bounds of the search, dashed vertical lines are left-hand bounds of the search, and the thick solid vertical line shows the location of the cut.

The cut identified in Figure 3-6 was located at approximately -0.2 . Figure 3-7 contains the portion of the original data points falling below this cut (upper plot) and the residuals for the linear regression performed on this subset of the data (lower plot). The cut is accurately placed at PHD_1 (or x) = -1 .




Figure 3-7 Top: Points in the current set and the best linear fit (shown with thin dashed line). Bottom: Residuals for the current set ($\text{PHD}_1 \in [-3, -0.2]$), showing the search for the best cut.

The rms errors of the two subsets to the left and right of the cut, $[-3, -1]$ and $[-1, -0.2]$ were within the specified error tolerance of $1e-8$ (actually, they were within machine precision), so no further cuts were made.

Finally, the analogous procedure is depicted in Figure 3-8 for the set of points in the range $[-0.2, 3]$. As before, the rms errors of the two subsets to the left and right of the cut were within error tolerance, so the algorithm terminated.

To summarize, the 800 training points were divided into 4 subsets, each with its own regression equations. The ranges of PHD_1 (or x) were $[-3, -1]$, $[-1, -0.2]$, $[-0.2, 1]$, and $[1, 3]$. The fits were all to within machine precision.




Figure 3-8 Top: Points in current set and the best linear fit (shown with thin dashed line). **Bottom:** Residuals for the current set ($\text{PHD}_1 \in [-0.2, 3]$), showing search for best cut.

The performance of the PHDRT algorithm was compared with that of the KRLS algorithm, using a Gaussian kernel function, with $\sigma \in [0.0125, 0.05]$ and $\nu = 0.05$ ¹⁵. The identical dataset was used, though inputs and outputs were normalized to the range $[-1, 1]$. One of the characteristics of the KRLS algorithm is that σ is generally problem specific and must be identified for each new problem. To do so, σ was changed within the specified range to identify the value that led to the lowest rms error on the test set. The results of the calculations are plotted in Figure 3-9. The optimal value for this dataset was 0.0375. Shown in Figure 3-10 are the errors on the training and test sets using the optimal σ . From this figure, it is apparent that the performance of the KRLS algorithm is worse than that of the PHDRT for this dataset. Table 3-2 contains performance data for the two algorithms. Note that the training time for the PHDRT algorithm is moderately shorter than that for KRLS.

	PHDRT Training	PHDRT Test	KRLS Training	KRLS Test
rms error	0	0	0.007	0.008
max error	0	0	0.071	0.053
training time (s)	3.7		4.2	

Table 3-2 Demo 1: Comparison of PHDRT and KRLS.

¹⁵ Refer to RunDemo1Comparison.m

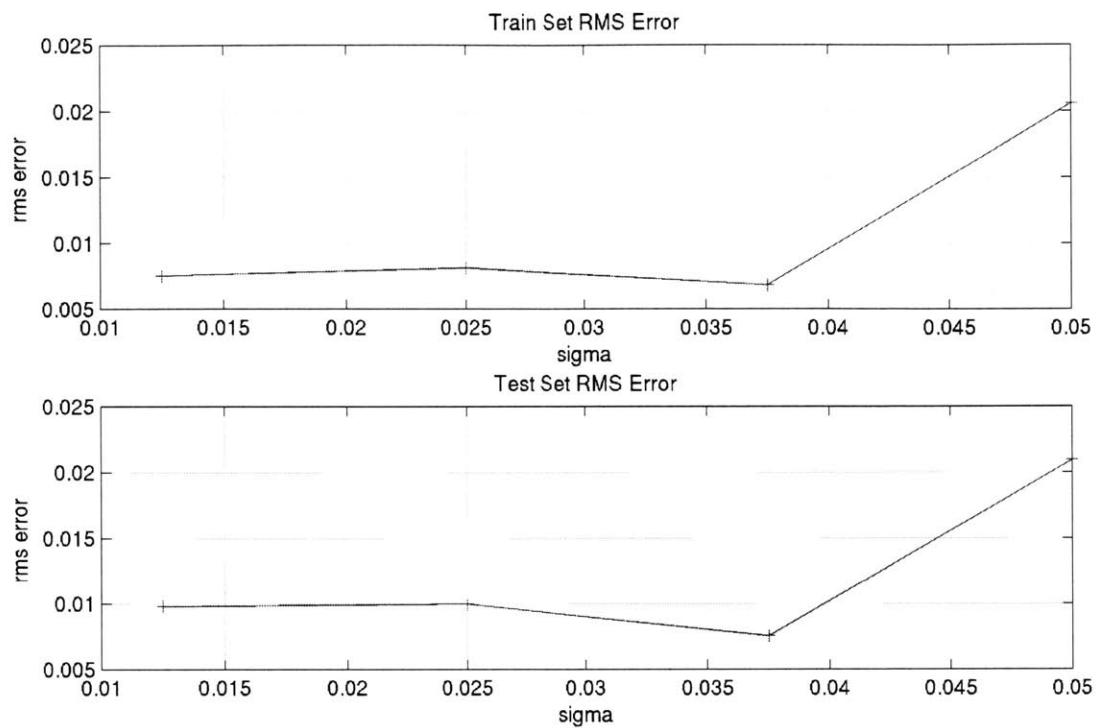


Figure 3-9 Determination of optimal σ for KRLS. The best value is 0.0375.

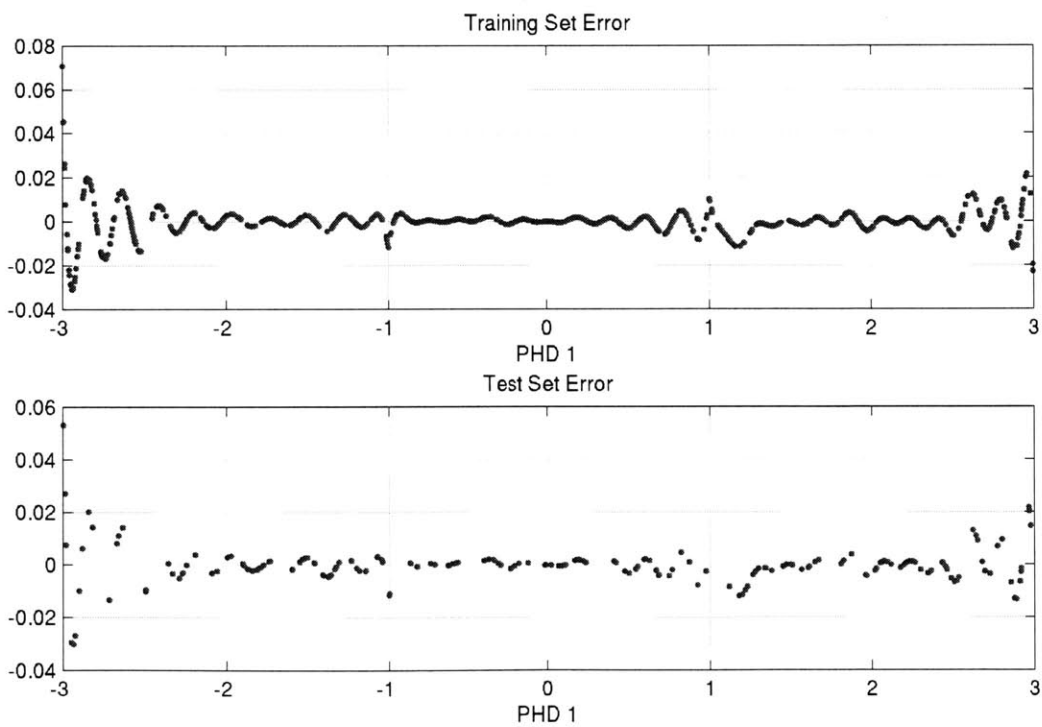


Figure 3-10 Training and test errors for the best KRLS model.

3.3.4.2 Demo 2¹⁶

In this demo, the functional form connecting x to z is identical to that in Section 3.3.4.1, with the exception that four additional confounding dimensions are added as potential model inputs. Each is drawn from the same distribution as x_j . They have no impact on z , but rather make the problem more difficult to model because the algorithm must recognize that the additional inputs are irrelevant. Again, the dataset consisted of 800 training points and 200 test points.

The influence of the additional input dimensions is evident in Figure 3-11, where, in the upper plot, z is plotted versus x_j (true input) and x_2 (one of the noise inputs). The lower plot shows how the presence of the noisy input leads to a less crisp curve after projection of the data points to form PHD₁ (c.f. Figure 3-5). The algorithm proceeds to identify 9 subsets of the data, as opposed to the 4 identified in Demo 1.

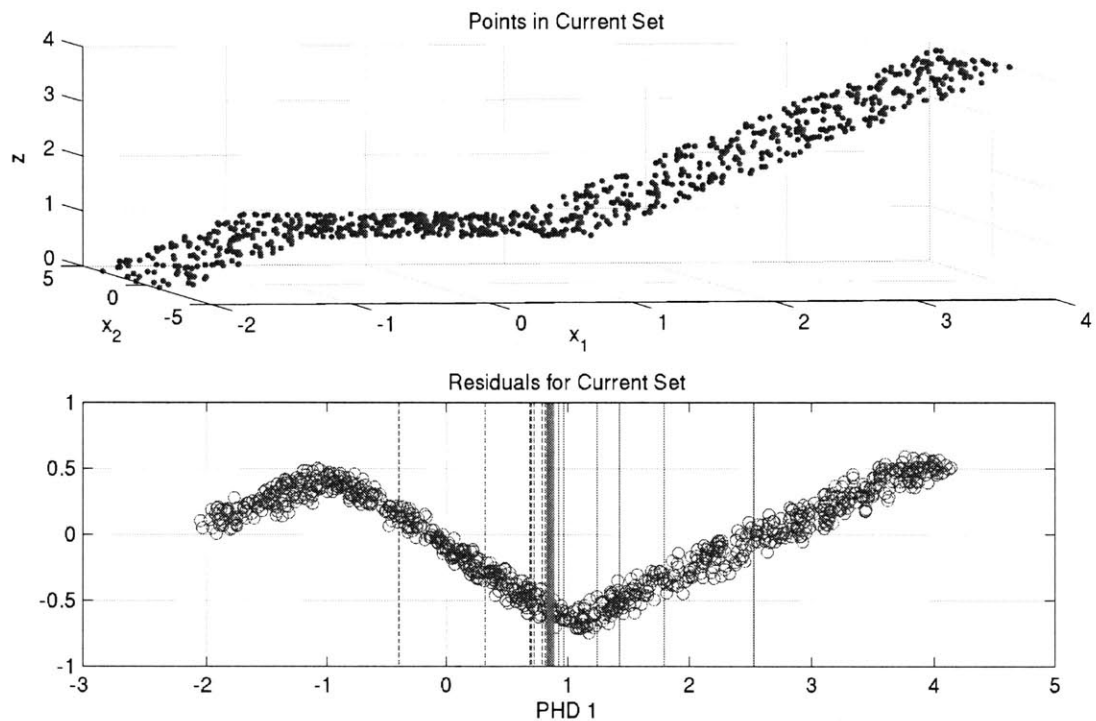


Figure 3-11 Full set of points (above) showing one extra dimension of noise input (x_2). The best cut position is shown with the bold vertical line (below).

The predictions and prediction errors on the training and test sets using the final PHDRT model are shown in Figure 3-12. Note the errors at $x_j = -1$ and 1, where the function's slope changes abruptly. Histograms of the training and test errors are presented in Figure 3-13.

¹⁶ See MainPHDRTdemo2.m




Figure 3-12 PHDRT training set predictions and errors (upper plots) and test set predictions and errors (lower plots).




Figure 3-13 Histograms of training and test errors for the PHDRT model.

As for Demo 1, in Section 3.3.4.1, a KRLS model was generated using the identical dataset (normalized as before)¹⁷. The choice of $\sigma = 0.5$ was made using the same method as before (see Figure 3-14).

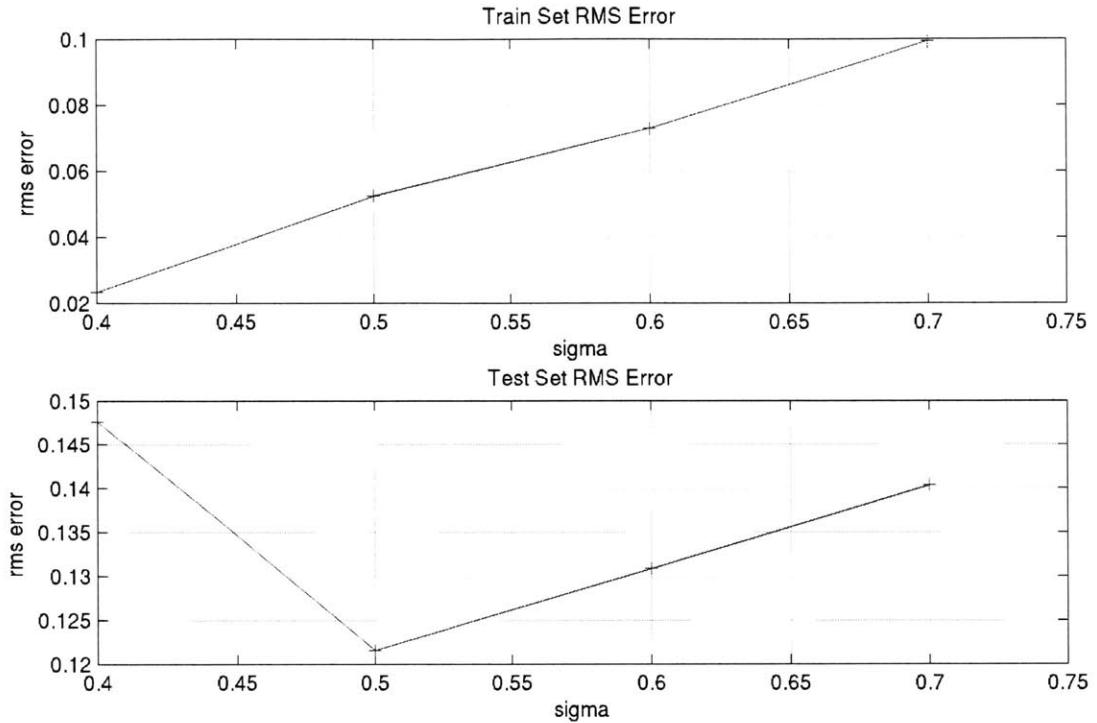


Figure 3-14 Selection of $\sigma = 0.5$ for the KRLS algorithm.

Significantly larger training and test errors for the KRLS model are depicted in Figure 3-15. The larger extent of the prediction errors is demonstrated further in the histograms of training and test error in Figure 3-16. There is a dramatic reduction in accuracy from the PHDRT model errors shown in Figure 3-13.

Overall performance of the two models on the Demo 2 dataset is given in Table 3-3. Note the large disparity in training times as well as in the errors, where the KRLS errors are roughly an order of magnitude greater than those for PHDRT.

	PHDRT Train	PHDRT Test	KRLS Train	KRLS Test
rms error	0.004	0.007	0.053	0.123
max error	0.051	0.064	.203	0.514
training time (s)	11		79	

Table 3-3 Comparison of PHDRT and KRLS on second demo problem.

¹⁷ See RunDemo2Comparison.m

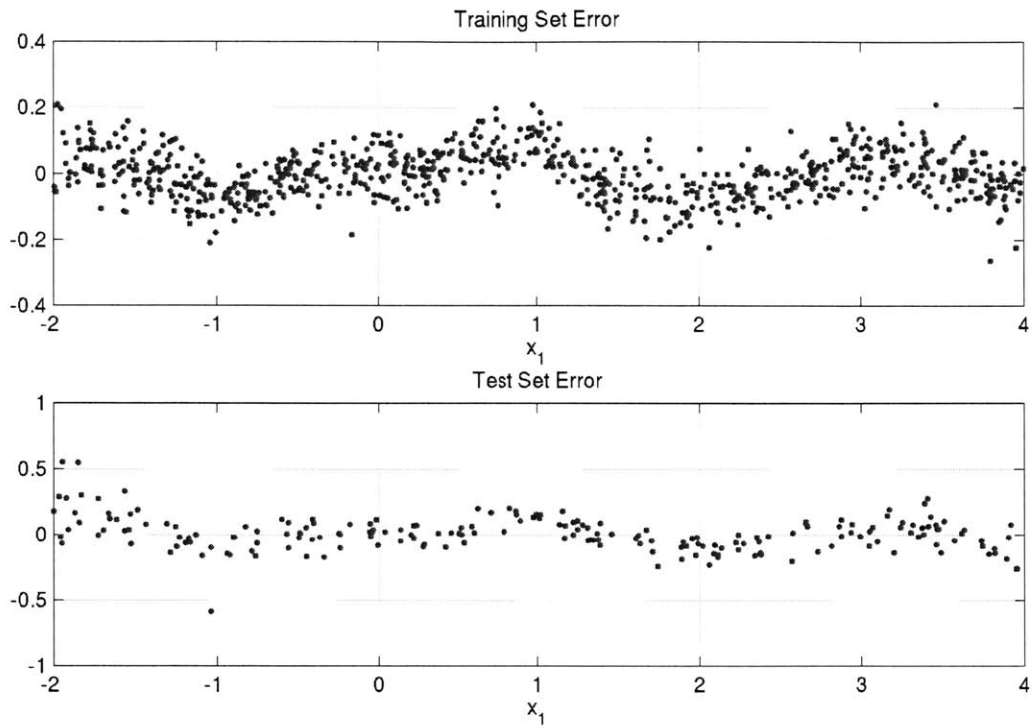


Figure 3-15 Training and test errors for the best KRLS model.




Figure 3-16 Histograms of best KRLS errors (*c.f.* Figure 3-13).

To investigate the impact of measurement noise on the modeling approaches, noise from a normal distribution with $\sigma^2 = 0.1$ was added to z before training the PHDRT and KRLS models. $\sigma = 0.6$ was found to be optimal for the KRLS model. Model performance was assessed in two ways. In the first, the agreement between the model predictions and the noise-contaminated signal was used for calculating residuals and rms errors. These data are shown in the first row of numbers in Table 3-4. In the second, the ability of the models to reconstruct the noise-free function was examined by comparing model predictions with the true function values evaluated at training and test points. These data are shown in the final row of Table 3-4. It is seen that the KRLS algorithm predicts the true function's behavior with rms errors approximately three times greater than those associated with the PHDRT algorithm.

	PHDRT Training	PHDRT Test	KRLS Training	KRLS Test
rms error (noisy)	0.097	0.103	0.114	0.167
rms error (noise free)	0.033	0.048	0.095	0.145

Table 3-4 Modeling performance in the presence of measurement noise.

3.3.4.3 Demo 3: Modeling of Two-zone Data¹⁸

In this final demo, the identical dataset used in Section 3.1.3.2 (linking temperatures T_0 , T_1 , and T_2 to the “Heat Input” to zone 1) was used to train a PHDRT model. The error tolerance was set to 0.001, otherwise settings were as described in Demos 1 and 2. The algorithm divided the data into 102 subsets, each with its own model. Training time was 89 seconds, rms errors were 0.002 and 0.003 for training and test sets, respectively. Maximum error for training and test sets were 0.027 and 0.032 respectively.¹⁹ These results compare favorably with those in Table 3-1. For reference, the best KRLS model (two models, one for each portion of the dataset—that would have to be identified by the user) yielded rms errors for the training and test sets of 0.004 and 0.013 after 473 seconds. When the PHDRT model is used, the test rms error is improved by a factor of about 4 and the training time is reduced by a factor of about 5.

Histograms of the training and test errors for the PHDRT model are shown in Figure 3-17.

¹⁸ See MainPHDRTdemo3.m

¹⁹ If the error tolerance is increased to 0.01, 17 subsets of the data are formed in 22 seconds, yielding rms training and test errors of 0.007 and 0.0072.




Figure 3-17 Histograms for Demo 3.

3.3.4.4 Further Comments on the PHDRT Algorithm

A brief study of the benefits of some of the proposed algorithm improvements (those in Section 3.3.3.2) is presented here. The impact of the dimension reduction algorithms proposed above was examined for the datasets of Demo 2 and Demo 3. Two dimension reduction approaches were proposed: the regression dimension reduction and the eigenvector dimension reduction. Experiments using all combinations of these reduction techniques were performed and are tabulated below.

Regression Reduction	Eigenvector Reduction	Demo 2 Test Set rms error	Demo 3 Test Set rms error
No	No	0.011	0.007
No	Yes	0.010	0.007
Yes	No	0.001	0.008
Yes	Yes	0.001	0.008

Table 3-5 The impact of the proposed dimension reduction techniques on the datasets used in Demo 2 and Demo 3.

The performance improvement associated with using the regression dimension reduction technique in the Demo 2 dataset (column 3) does not appear in the Demo 3 dataset (column 4). The strong improvement in the case of the Demo 2 set is related to the nature of the dataset. In that case, there is no interaction among the various input variables, or dimensions. The model output is a function of one input alone, so any tool to reduce the impact of the others is

beneficial. In Demo 3, all input dimensions contribute to the model, so the benefit of the proposed techniques in this case is negligible. The techniques should be applied on a case-by-case basis.

An assessment of the impact of the other techniques proposed in Section 3.3.3 is not presented here.

3.4 Summary and Proposed Approach

In this chapter, we have seen that the considerable complexity encountered in naturally ventilated or mixed-mode buildings is not likely to be well captured by a single analytical model. The building may operate in several operational or control modes, each of which is associated with multiple airflow patterns or regimes. It was found that ensembles of local linear models, such as were generated by the PHDRT algorithm, were best suited for capturing the thermal impact of the airflow patterns generated by simple analytical models. While such models excel under the conditions of continuous inputs, they are not appropriate for treating binary control mode inputs.

Based on the information presented in this chapter, a two-tiered approach to modeling a naturally ventilated building:

1. The user selects the dominant (or most prevalent) control modes of the building and sorts the complete dataset into subsets, one per mode. Particularly in the case where multiple openings exist, simplifications must be made to reduce the overall number of modes.
2. With each subset identified in Step 1, use the PHDRT algorithm to create a model of the system as it operates in each specified mode.

This approach will be investigated in the later chapters using experimental data collected from two mixed-mode buildings.

4 Broadmoor: Experimental Description

4.1 Introduction

The experimental portion of this thesis was performed at the Saltonstall Nature Center at the Massachusetts Audubon Society's Broadmoor Wildlife Sanctuary²⁰. In this chapter, a description of the center is provided, followed by a detailed description of the experimental setup. The data collected as described in this chapter are used in the next chapter to model the thermal behavior of the building.

4.2 Building Description

The Broadmoor Wildlife Sanctuary (to be referred to as Broadmoor) is situated in South Natick, Massachusetts, approximately 20 miles southwest of Boston. A 1911 horse and chicken barn on the site of the sanctuary was retrofitted in 1983 to create office and assembly space for the staff, volunteers, visiting public and summer campers. The building retrofit was led by sanctuary director Elissa Landre, Bill Giezentanner (former planner for Massachusetts Audubon Society) and architect Gerard Ives, all of whom have been extremely generous throughout the project—sharing the building, their insights and time.

The building was first monitored by Chang (2002) during the summer of 2001. His work included thermal comfort assessments and occupant surveys as well as development of a physically-based, second-order model of the first floor. The monitoring capabilities at the time were somewhat limited, as solar radiation data, electrical heat gains, and aperture schedule had to be estimated. Despite the shortcomings of the data, qualitative agreement between model predictions and measurements was achieved.

4.2.1 Physical Description

The retrofitted building was designed to be nearly entirely passive, with heating assistance only from a basement woodstove. Triple-glazed windows and copious insulation (R-30 in walls, R-45 in ceiling) were used. Distribution fans are used in the winter to blow warm air from a sunspace and attached water wall into the occupied areas. Another fan is used to move heated sunspace air to a rockbed below the basement slab. In the summer season, no mechanical systems are used for cooling (with the exception of personal fans).

The rectangular building is 40' x 70', with the long dimension facing South²¹. There are two occupied levels (the basement and the first floor) and two unoccupied levels (unfinished attics). The ceiling heights of the basement and first floor are 7'7" and 9'9". The first attic (Attic 1) ceiling height is 9', and the distance from the second attic (Attic 2) floor to the center of the cupola apertures is 17'. The two attics are connected by a stairway with an opening size of approximately 70" x 36" and by numerous gaps in the flooring. A sunspace including water tubes for heat storage was built along much of the South side to harvest solar energy in the wintertime. Photographs of the building exterior are shown below in Section 4.2.1.1. The

²⁰ http://www.massaudubon.org/Nature_Connection/Sanctuaries/Broadmoor/index.php.

²¹ The direction is actually 164°, or SSE, not adjusting for magnetic declination.

building floor plan (first floor) and section are shown in Section 4.2.1.2. Photographs of the building interior are shown in Section 4.2.1.3.

4.2.1.1 Photographs of Building Exterior

Shown in this section are photographs of the building exterior, moving around it in a counter-clockwise direction. Numeric labels on the photographs correspond to controlled apertures. A complete list of the controls (apertures and fan) appears below in Table 4-1 on page 78.




Figure 4-1 Left: Exterior view of main entrance of building on the West side. Right: West and South sides, showing the two-story Sunspace. The three upper windows on the West face are in Attic 2. The smaller ones are operable. Note the cupola in the photo at right.



Figure 4-2 Left: View of building from Southwest. Note aperture #2 on the wall perpendicular to the front door. The white arrow shows the louvered door kept open all summer to vent hot air from the Sunspace. Right: View of the louvered door to the right of the main entrance.




Figure 4-3 View of the Sunspace from the Southeast. The white arrow shows a louvered door kept open all summer to vent hot air from the Sunspace.



Figure 4-4 Rear (East) of building, showing operable Attic 2 windows (#10), Conference and Assembly room doors (#3), and the East louvered door (#1). The unlabeled windows (two larger Attic 1 windows and one small Attic 2 window) are not operable.



Figure 4-5 North face of building. Open windows at right are to vent the restrooms. These apertures were not carefully monitored since the restrooms are somewhat isolated from the rest of the building. The white asterisk (*) indicates the Attic 1 window above which the outside temperature thermocouple (thermocouple 10) was placed. It was placed under the eaves, protected from direct sunlight.

4.2.1.2 Architectural Plans: Floor Plan and Building Section

The floor plan for the first floor and a building section are shown in the following two figures. Numeric labels on the drawings correspond to controlled apertures (with #) and thermocouple locations (with arrow). A complete list of the controls (apertures and fan) appears below in Table 4-1 on page 78. A complete list of the thermocouple array appears below in Table 4-2 on page 82.

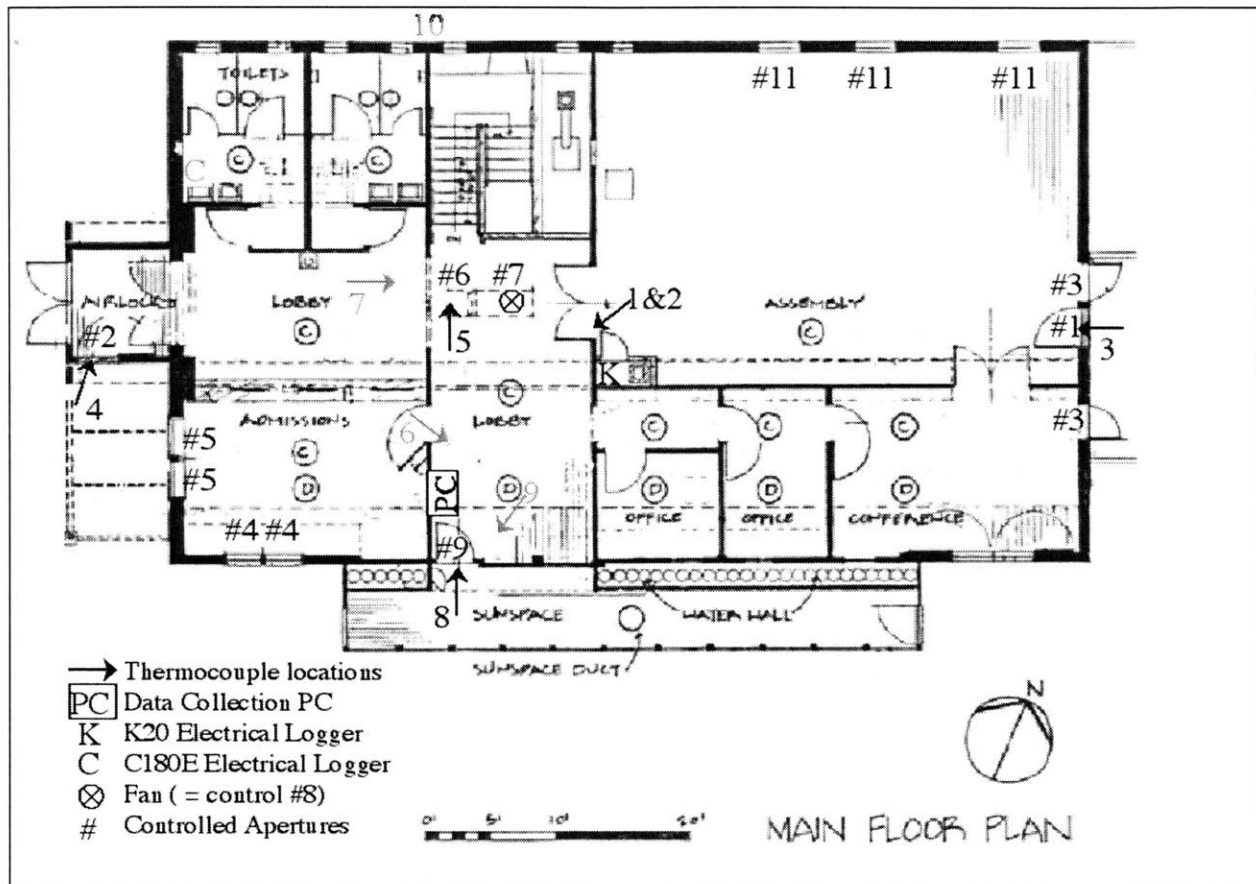


Figure 4-6 First floor plan. Also shown in this figure are the locations of thermocouples used to measure temperatures in different locations of the building. Grey arrows and numbers (rather than black) are used to indicate the position of thermocouples projected onto the first floor plan. Specifically, 6 is in Attic 1, 7 is in Attic 2, 9 is in the Basement, and 10 is outside, above an Attic 1 window. Black arrows and numbers are used for thermocouples that are located on the first floor. Note that the positions of the electrical data loggers are also provided. C (in grey) is located in the Basement. (Floor plan courtesy of Gerard Ives.)

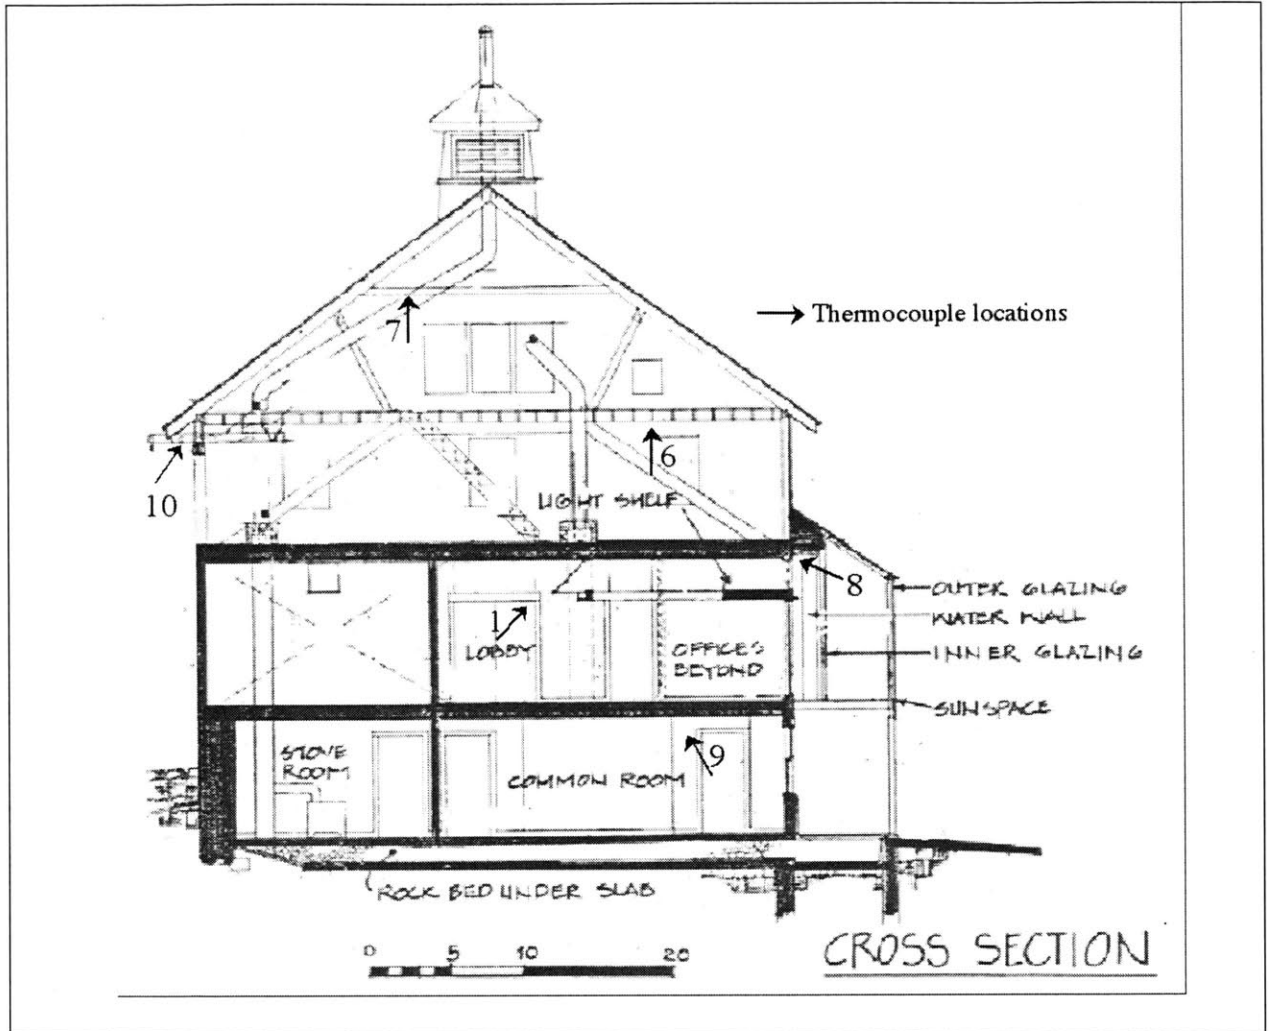


Figure 4-7 Building cross section, showing locations of selected thermocouples. (Section courtesy of Gerard Ives.)

4.2.1.3 Photographs of Building Interior

Photographs of the interior of the building are provided in this section. Numeric labels on the photographs correspond to controlled apertures. A complete list of the controls (apertures and fan) appears below in Table 4-1 on page 78.




Figure 4-8 Left: View from main entrance into the Lobby and Assembly Room beyond. Right: Immediately to the viewer's right is the West Louvered Door (#2). Note thermocouple (4) near the top of the door.




Figure 4-9 Two views of the Assembly Room, with labeled apertures. Left: North wall. Right: East wall.




Figure 4-10 Left: Conference Room and Right: Office. Through the window, one water tube of the water wall is visible.



Figure 4-11 Left: Lobby area, facing Sunspace (Sunspace door is #9). PC used for data collection can be seen. Right: Admissions Office.



Figure 4-12 Upper level of Sunspace. Slats in the floor (1/2") permit some exchange of air with the lower level of the Sunspace. Left: View towards East. In the Fall and Winter, a fan moves hot air from the top of the Sunspace through the duct in the center of the photo to the rockbed below the Basement. The water wall is on the left side of the photo. The closed door at the East end of the Sunspace was open throughout most of the monitoring period to allow escape of hot air. Right: Detail of water wall, with slider open to show water-filled tubes.

9

Figure 4-13 Common Room in Basement. Note Basement Door (#12) and Basement thermocouple location (9).



4.2.2 Design Intent

The design intent for the cooling season was that the building be cooled by natural ventilation, with cool air entering apertures #1 and #2 (louvered, screened doors), replacing warm air rising up through the Attics and out the cupola. Cross ventilation would also contribute to cooling the building, with the West Louvered Door (#2) serving as an inlet to the (generally westerly) winds and the East (#1) as an outlet.

This design intent was somewhat thwarted by the presence of only a very small opening (1.9 ft²) that allowed warm air to pass from the occupied first floor up into the Attics. Despite this, Chang (2002), page 169, found that stack ventilation was the dominant mode of airflow in the building. Due to the building's orientation on an East-West axis, and cross-ventilation openings restricted primarily to those ends of the building, the flow resistance through the building was too great to allow effective cross ventilation. Furthermore, he found that even when West winds blew, air actually exited the West Louvered Door. The origin of that behavior may be linked to the fact that only the West face of the cupola was open—the other three were boarded up. This situation created a “wind scoop” whereby westerly winds were directed into the Attic, blowing hot (foul-smelling²²) Attic air into the occupied area.

4.2.3 Building Usage and Operation

Prior to the monitoring and experimental period begun with this research, the building's operation was primarily passive and seasonal. The louvered doors and small Attic opening generally remained open for the entire summer season. Other apertures, particularly the West office windows (#5), conference room and Assembly doors (#3) and, less frequently, the Assembly room windows (#11) were used for temporary cooling.

In 2003, the hours of operation for the building were 9-5 Tuesday-Friday (several staff members and I were often present on Mondays), 10-5 weekends and holidays (staffed by a single volunteer). The first floor was generally occupied by at most 3-4 staff members or volunteers, and the Basement by 1 staff member. Due to the nature of the mission of the Sanctuary, staff members were frequently not working inside. Guests visited the building for admissions, restrooms, and special functions. With the exception of special events, there were usually fewer than 6 guests in the building during the time of the experiments.

For six weeks in July and August (July 7 – August 15, 2003), summer camps for children were held Monday through Friday. During that time period, hours of operation on weekdays were 8-4. The occupant load on the building was very sporadic, as the children were usually busy outside. At times, two or three dozen children were in the building.

On average, the electrical usage during the summer months was approximately 700W, split roughly evenly between first-floor and Basement zones.

²² Fowl-smelling may be more appropriate, as the barn previously housed chickens.

4.2.4 Experimental Setup

Described in this section are a set of building modifications that were needed to both improve building performance and to allow a broader range of experimental options. Following this initial material (Sections 4.2.4.1 - 4.2.4.3) is a detailed description of the experimental setup, including all measurement equipment and software used to carry out the research.

4.2.4.1 Modification of Openings to the Attic

Of primary interest when considering use of this building for experimental purposes was the possibility of having larger, adjustable openings between the first floor and Attic 1. The pre-existing opening is both fixed and small (1.9 ft²). It is located centrally in the building, in the ceiling above the lobby (refer to Figure 4-6).



Figure 4-14 Pre-existing aperture adjacent to pull-down stairs. The large opening at left was not available for airflow since the Attic stairs were closed except for infrequent trips to the Attic for maintenance or storage. Shown at right (white arrow) is the small, fixed aperture that was previously kept unobstructed throughout the summer season. At the top of the photo is an insulated exterior door that swings down to seal all the openings in the winter season.

Broadmoor's director agreed to the removal of the pull-down stairs shown in Figure 4-14.

The opening was covered from below with sliding panels permitting straightforward manual adjustment of aperture size. A whole-house fan, described below, was installed on the East end of the opening. During fan operation, it was necessary to prevent air from flowing from the Attic back down through the large opening into the first-floor space. To accomplish this, two independent sliding panels were installed, one to cover the fan (#7) and one to cover the rest of the aperture (#6).



Figure 4-15 Retrofitted apertures to the Attic, with sliding covers. View from below with both sliders open and closed. Left: Note the compartment made to isolate the fan from the rest of the opening (to prevent backflow of Attic air). The large slider covers the pre-existing aperture (at left) and a portion of the opening created by the removal of the stairs.

Figure 4-16 View of the slider configuration used during fan operation.



The dimensions of the apertures are given in the sketch below, in Figure 4-17.




Figure 4-17 Dimensions of apertures to Attic 1. Not to scale. Note that the diameter of the fan opening is greater than the size of the rectangular opening covered by the Fan Slider. Shaded regions represent obstructions to airflow.

4.2.4.2 Fan Installation and Properties

As mentioned in the previous section, a whole-house fan was installed to allow experiments using mechanical ventilation. A 24" CertainTeed Ventilation fan (model WH242MLX) was purchased. As delivered, the fan specifications were 600 rpm/3200 cfm at low speed; 775 rpm/4500 cfm at high speed. High-speed power consumption was rated at 240W.

Variable speed control of the fan was desired, so a General Electric 1/3 hp ECM programmable motor (GE Model # 5SME39DSO451, S/N 012173), donated by Nailor Industries Inc. for these experiments, was installed in place of the pre-existing direct-drive motor. Also included from Nailor was their EPIC™ volume controller card. As controlled, the torque of the motor varies linearly with the input voltage (0-5000mV) supplied to the card. The rated maximum rpm was 1050.

The relationship between the fan speed and the control signal was measured using a Solomat 327™ optical tachometer attached to a Solomat MPM 500e. Measurements were made with the fan installed at Broadmoor, with both louvered doors open. See Figure 4-18.

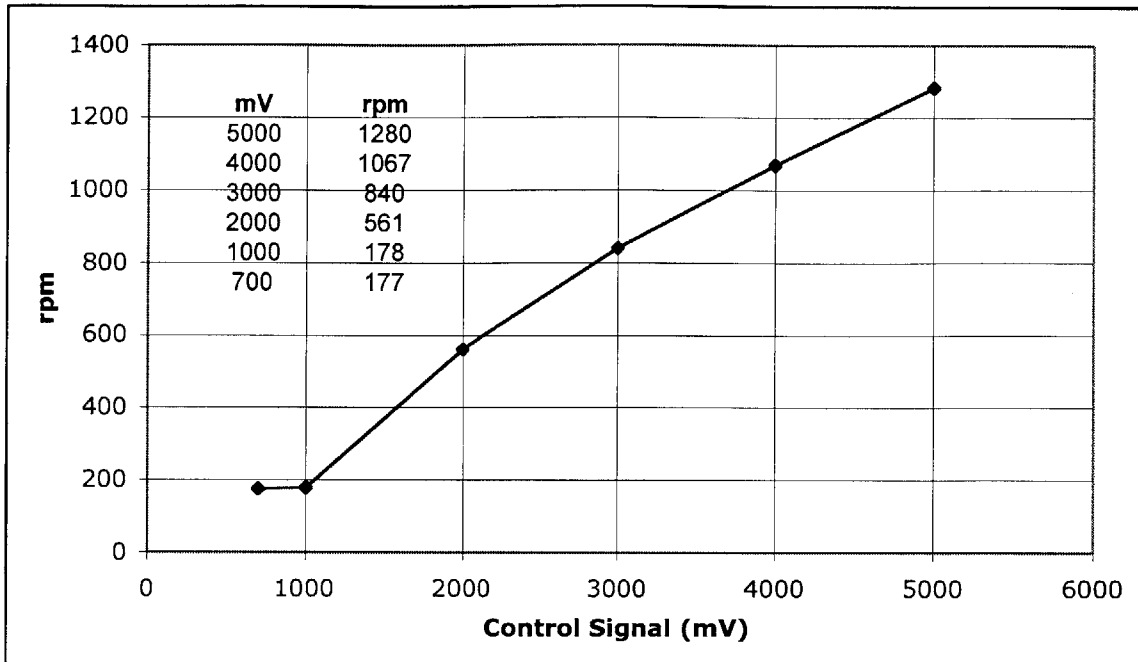


Figure 4-18 Fan speed versus control signal (voltage supplied to controller card). Note the behavior below 1000mV.

No airflow measurements were made when the fan was on. The specifications for the original fan indicate that the flow at 600 rpm is 3200 cfm (corresponding to a 2140mV control signal in Figure 4-18), and that the flow at 775 rpm is 4500 cfm (corresponding to a 2768mV control signal in Figure 4-18). By application of the fan law stating that flow is proportional to fan speed (McQuiston and Parker 1994), the maximum flow rate for the fan (at 5000mV control signal) is found to be between 6830 and 7430 cfm. A flow rate of 7100 cfm corresponds to 15.6 air changes per hour for the first-floor zone.

Using data from the same experiment used to generate Figure 4-18²³, the relationship between the control signal and the fan power was measured using the electrical power logging equipment described below in Section 4.2.4.8. The results are shown in Figure 4-19.

²³ September 24, 2003, 6:53-7:17 a.m. Voltage was maintained at each level for two minutes. To avoid transient effects, the power measurements for the second minute only were used.

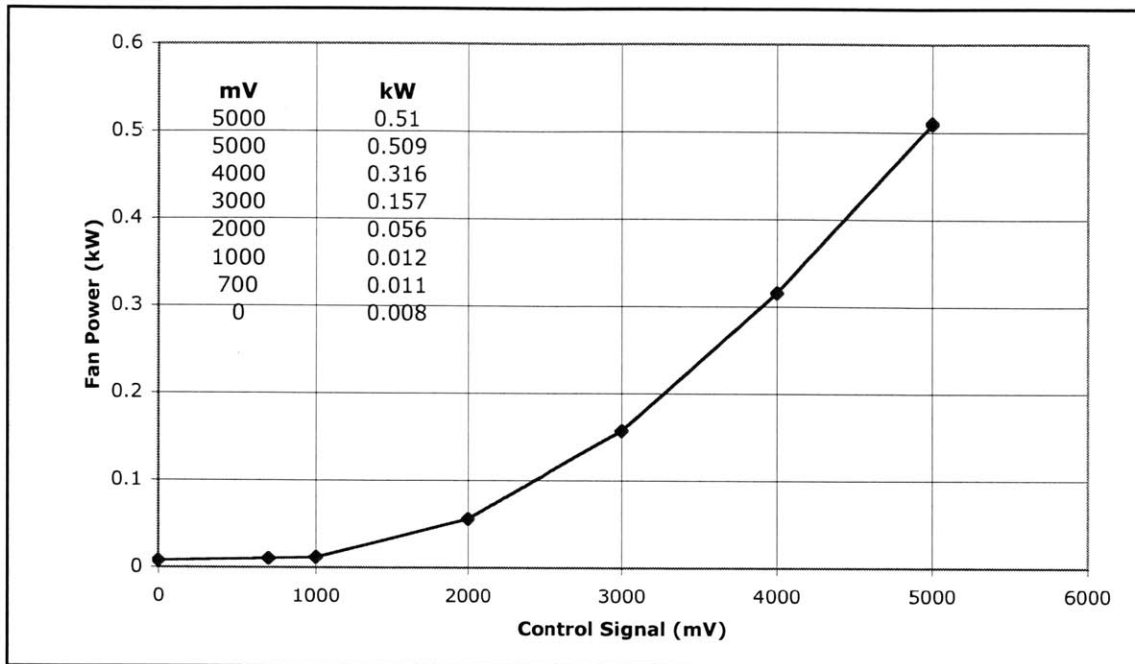


Figure 4-19 Fan power versus control signal (voltage supplied to the controller card). Note the behavior below 1000mV.

Note that by linear interpolation, the fan power consumption of the ECM-driven fan is 134W at the rated full speed of the original fan²⁴.

The fan was installed as shown in the following figures.



Figure 4-20 Installation of ECM-driven fan. Left: In the foreground, note the black motor controller on an improvised heat sink. This element was not in the airflow path and became very hot.

²⁴ Here, the association from Figure 4-18 was used that 2768mV corresponds to 775 rpm, the rated full speed of the original fan. The implication of this calculation is that the ECM-driven fan uses only 56% of the power of the original fan.




Figure 4-21 Controller card (left arrow) and 24V transformer (right arrow). The control signal to the card was transmitted through the white Cat.5e cable shown in the foreground.

4.2.4.3 Cupola Improvements

To eliminate the problematic “West-wind scoop” and to create a greater aperture for exhausting hot Attic air, the paneling covering the North, East and South faces of the cupola was removed. The size of each of the resulting four openings was 4’ x 3’. The cupola itself was approximately 6’ square. The distance from the center of the cupola openings to the Attic 2 floor was approximately 17’. Half-inch screening was installed on the interior of the louvers on the four sides of the cupola. The louvers were spaced approximately four inches apart. Several photographs below provide interior and exterior views of the cupola.



Figure 4-22 Left: View looking up into cupola after paneling was removed: three interior faces. Right: Unobstructed interior face of one side of the cupola, showing the 1/2” screen.




Figure 4-23 Left: South and East faces of cupola. Right: West and South faces.

4.2.4.4 Controlled Elements

To influence the airflow and temperature within the building, twelve different building elements were controlled throughout the experiments. With the exception of the fan, all were apertures. The fan was controlled via computer, while the apertures were adjusted manually. All controlled elements are listed below in Table 4-1, which provides the location, opening size, range of settings, and references to figures where the element is pictured.

ID #	Controlled Item	Location	Opening Size	Range	Shown in Figures:
#1	E Louver Door	E end of Assembly Room	33" x 72 1/2" (1)	0/1	Figure 4-4 Figure 4-6 Figure 4-9
#2	W Louver Door	S side of Airlock, adjacent to Front Doors	33" x 72 1/2" (1)	0/1	Figure 4-1 Figure 4-2 Figure 4-6 Figure 4-8
#3	Conference Room and Assembly Room Doors	E end of Assembly Room and Conference Room	Conference: 34 1/4" x 79 1/4" (with insect screen) Assembly Room: 35" x 79 1/2" (no screen)	0/1/2	Figure 4-4 Figure 4-6 Figure 4-9 Figure 4-10
#4	S Office Windows	2 double-hung windows on S side of Admissions Office	29 1/2" x 23 3/4" (with insect screen, 16x16 squares per in ²)	0	Figure 4-2 Figure 4-6 Figure 4-11
#5	W Office Windows	2 double-hung windows on W side of Admissions Office	29 1/2" x 23 3/4" (with insect screen, 16x16 squares per in ²)	0-1	Figure 4-1 Figure 4-2 Figure 4-6 Figure 4-11
#6	Attic Slider	Wooden Slider covering W side of opening to Attic 1	5.9 ft ² (see Figure 4-17)	0/1	Figure 4-6 Figure 4-15 Figure 4-16
#7	Fan Slider	Wooden Slider covering the Fan on the E side of the opening to Attic 1	23 1/2" x 25 1/2"	0/1(2)	Figure 4-6 Figure 4-15 Figure 4-16
#8	Fan	Fan located on E side of the opening to Attic 1	24 3/4" diameter	0-5000mV	Figure 4-6 Figure 4-16 Figure 4-20
#9	Sunspace Door	Adjacent to PC in Figure 4-6	35" x 79"	0/1/2 (3)	Figure 4-6 Figure 4-11
#10	Attic 2 Windows	2 Windows on both E and W ends of Attic 2	4 x (22" x 29")	0/1	Figure 4-1 Figure 4-2 Figure 4-4 Figure 4-5
#11	Assembly Room Windows	3 double-hung Windows on N side of Assembly Room	3 x (33 1/2" x 27 3/4") (with insect screen, 16x16 squares per in ²)	0-1	Figure 4-5 Figure 4-6 Figure 4-9
#12	Basement Door	Next to SW end of Sunspace	35" x 79 1/4"	0/1	Figure 4-2 Figure 4-13

Table 4-1 List of control settings. All apertures connect inside zones to the outside, with the exception of #6 - #8, which are located between the first floor and Attic 1. In the column entitled "Range", the notation using "/" indicates that the aperture was adjusted discretely. For example, aperture #1 was either open or shut. Apertures #3 were either both closed, both open, or just one was open. Numbers shown in parentheses refer to notes directly below, in the text.

- (1) Louvers: 7 3/4" long, 3/4" thick, separated by 2 3/4". Horizontal depth of louvers (exterior-most point to screen): ~ 6". There were 15 openings between the louvers. Screen over openings was 18x18 squares per square inch.
- (2) On only one occasion was this opened fractionally.
- (3) Generally, the Sunspace door was shut. On September 10, 2003, the door was opened to warm the building. The louvered door on the upper level of the Sunspace was shut. The label "1" reflects that situation. To increase the warming, all the sliders connecting the Sunspace to the water wall were opened. The label "2" reflects that situation.

Several rules were observed in setting the controls:

1. Whenever the fan (#8) was on, aperture #7 was always open and aperture #6 always shut.
2. The fan (#8) was used only at night, or during the early morning hours (before staff arrived).
3. Apertures #1 and #2 were open whenever any aperture to the Attic (#6/#7) was open.
4. Apertures #1 and #2 were generally opened/closed at the same time (specifically, 1 minute apart).
5. Apertures #10 remained open at all times until September 9, 2003, when they were all shut.
6. Apertures #4 were not used throughout the experimental period.
7. No apertures leading to the exterior were left open when the building was unoccupied, with the exception of #1 and #2.

The monitoring and adjusting of all the apertures was a manual process, occurring on a minute time step. With the assistance of the staff and volunteers, recordkeeping for most apertures was quite accurate. As it was located on another level, the Basement door (#12) was difficult to monitor, and records for it were less complete. As there were grilles in the floor between the Basement and first floor, opening aperture #12 may have affected airflow through the first floor. Furthermore, doors closing off the stairs to the Basement level were generally closed, but not at all times. Additional flow through the stairway may have been induced when the stairway doors were open at the same time as #12.

The restroom windows were not systematically controlled and were not monitored. The restrooms were somewhat—though not entirely—isolated from the rest of the building.

In order to generate data rich enough for constructing models of the building (but also to provide reasonable comfort to the occupants), the major apertures (louvered doors (#1 and #2), Attic openings (#6 and #7)) were generally adjusted several times daily. The minor apertures, such as office windows (#5), conference and Assembly doors (#3) and the Basement door (#11), were adjusted by staff, as desired. Of particular interest was the benefit of aggressive night cooling, *i.e.*, shutting all building apertures when outside temperatures exceed interior temperatures, and opening them as much as possible at night, allowing night air to cool the building in preparation for the following day.

4.2.4.5 Temperature Measurements

Temperatures in ten locations were measured throughout the experimental period. Special Limits of Error copper-constantan thermocouple wire (Omega TT-T-24-SLE), was used for all measurements²⁵.

The thermocouple placement was planned to avoid direct radiation from the sun and incandescent lights. It was found that the thermocouples are influenced by the temperatures of adjacent surfaces (Arens 2000). Arens found that coating the thermocouple with silver paint reduced this effect. In the research cited, exposure to a surface of temperature approximately 12°C greater than dry-bulb temperature increased the uncoated thermocouple's temperature reading by approximately 0.8°C, while the temperature of the coated thermocouple was increased by approximately 0.4°C. The 38 cm² surface was held at a distance of 36 cm from the thermocouples. Note that temperature difference between the dry-bulb and surface temperatures is much larger than expected in the occupied zones at Broadmoor. Chang (2002) performed surface temperature measurements in the Basement and Assembly zones and found that the temperature difference was less than 2°C, often considerably less. The temperature differences in the unoccupied zones are expected to be higher (in the Attics and Sunspace) due to the presence of surfaces heated directly by the sun.

Especially when large temperature differences are anticipated, care should be taken to paint or otherwise shield the thermocouples when the modeling and control framework presented in this thesis is implemented permanently in an occupied building. Doing so will minimize radiation-induced errors that could cause true zone temperatures to violate desired temperature set points.

This issue was probed further in an attempt to quantify the effect of radiation on the apparent air temperature measured by the thermocouple. Assuming steady-state conditions (and ignoring heat transfer through the thermocouple), the following heat balance may be written:

$$h_{conv}A_{sensor}(T_{air} - T_{sensor}) + \epsilon_{sensor}A_{sensor}\sigma(T_{surfaces}^4 - T_{sensor}^4) = 0 \quad [4-1]$$

where h_{conv} is the convection coefficient, A_{sensor} is the surface area of the thermocouple, ϵ_{sensor} is the emissivity of the thermocouple, and σ is the Stefan-Boltzmann constant. It is assumed that all surroundings have the same surface temperature. For this example, the air temperature was 298K, and the surface temperature was 296K (consistent with the ΔT observed by Chang (2002)) and 293K (equivalent to a 5K temperature difference). Convection coefficients for the thermocouple of diameter $\sim 1/32$ " were estimated using equations 9.34 and 7.45b from Incropera and DeWitt (1996) for natural convection and forced convection at 0.2 m/s, respectively. The resulting convection coefficients were 26 and 65 W/m²K. For a range of emissivities, the resulting sensor temperatures are plotted in Figure 4-24 for the cases where surface temperatures are 296 and 293K and for h_{conv} equal to 26 and 65 W/m²K. An estimated lower bound for the

²⁵ The manufacturer's stated accuracy for this material was $\pm 0.5^\circ\text{C}$, with the particular batch found to have a deviation of 0.4°F at 200.3°F . No measure of precision (variance, or repeatability) was available. From the measurements taken at Broadmoor, the variance was found to be very low and negligible for these experiments.

sensor emissivity is 0.24, for NRL leafing aluminum paint²⁶. An estimated upper bound for the soldered thermocouple tip is 0.8. When the surface temperatures are 296K, if the tip is painted, the apparent air temperature is 279.89 and 279.96K for two convection coefficients, respectively. When the tip is not painted, the apparent air temperature is as low as 297.69 and 297.87K. The effect of painting the tip is to reduce the influence of radiation by a factor of approximately three. This same reduction is observed for the case with the lower surface temperatures. According to these calculations, the effect of radiation should be considered even if small surface-air temperature differences are anticipated.




Figure 4-24 Sensor temperature versus sensor emissivity for two surface temperatures and two convection coefficients.

²⁶ http://www.electro-optical.com/bb_rad/emissivity/matlemisivty.htm

The locations to be monitored were selected by their perceived importance in governing the building's thermal dynamics, weighed by the channel number limitations of the data logger. Shown in Table 4-2, below, are the names and locations of the ten thermocouples, as well as the figures in which they appear. Following the table are additional figures depicting the placement of the various thermocouples.

ID #	Name	Location	Shown in Figures:
1	Assembly	Suspended 1 3/4" inches below Assembly room door frame, 3/4" from the South side of frame, 78 3/4" above floor.	Figure 4-6 Figure 4-25
2	Mass	Directly above #1, 2 1/8" above Assembly room door trim, 86 3/4" above floor, sealed inside sheetrock.	Figure 4-6 Figure 4-25
3	East Louver Door	Woven in screen of door, 12" below top of opening, 67 1/2" above floor.	Figure 4-6 Figure 4-25
4	West Louver Door	Woven in screen of door, 12" below top of opening, 67 1/2" above floor.	Figure 4-6 Figure 4-8 Figure 4-25
5	Attic Opening	Extended E 2 1/2" from center of W edge of aperture #6, 3" below trim into which wire staple is driven (Figure 4-26, right), 128" above floor.	Figure 4-6 Figure 4-26
6	Attic 1	Suspended from Attic 2 floor joists, 89" above Attic 1 floor, 115" from S wall, 22' from W wall.	Figure 4-6 Figure 4-27
7	Attic 2	Suspended 1 1/2" below Attic 2 collar tie, 90" above floor, 21'10" from W wall, 14'3" from N wall.	Figure 4-6 Figure 4-27 Figure 4-28
8	Sunspace	Suspended below dropped ceiling outside door to Sunspace (#9), 82" above floor, 23" S of plane of door, 4 3/4" from partition on E side.	Figure 4-6 Figure 4-29
9	Basement	In N fork of "Y" of SE support column, 85" above ground.	Figure 4-6 Figure 4-13 Figure 4-30
10	Outside	Under eaves and above Attic 1 window, 11" above inner edge of window trim, 4 1/2" from N wall, directly above inner corner of trim, 21' above ground.	Figure 4-5 Figure 4-6 Figure 4-31

Table 4-2 List of thermocouple locations.





Figure 4-25 Left: Assembly room air thermocouple (1) and Mass thermocouple (2). Right: West Louver Door thermocouple (4) (placement identical on East Louver Door (3)).





Figure 4-26 Left: Attic Opening thermocouple (5), from below. Right: From above.




Figure 4-27 Left: Attic 1 thermocouple (6). Right: Attic 2 thermocouple (7).




Figure 4-28 Attic 2 thermocouple (7).




Figure 4-29 Sunspace thermocouple (8).





Figure 4-30 Basement thermocouple (9).




Figure 4-31 Outside thermocouple (10). Left: Window pictured is in Attic 1.

4.2.4.6 Weather Station Measurements

A weather station was installed on a pole in an open field, approximately 100' WSW from the SW corner of the building. To measure wind direction and speed, a wind vane and anemometer (Davis Instruments, model 7911) were installed. To measure solar radiation, five LI-COR Inc. pyranometers were installed, facing the four cardinal directions²⁷ as well as directly upwards. Since the building is oriented in (almost) the same way, no geometrical calculations were required to determine incident solar radiation on the building faces. Photographs of the weather station are provided in Figure 4-32, below.




Figure 4-32 Left: Wind vane, anemometer and pyranometers. Right: Detail of pyranometer setup. The pyranometers faced North, East, South, West and directly upwards. The spikes on upper surface were to discourage birds from roosting on the pole.

4.2.4.6.1 Technical Details: Wind Vane and Anemometer

The wind vane and anemometer were mounted 80" above the ground in an open field.

²⁷ Not adjusted for magnetic declination.

Four wires were connected to the wind station:

1. Yellow: 2500mV supplied to unit
2. Red: Ground
3. Black: Pulsed output (from anemometer— 1 pulse per full rotation)
 - a. 26.667 revolutions per minute \Leftrightarrow wind speed of 1 mile per hour, therefore:
 - b. x revolutions per minute $\Leftrightarrow x \left(\frac{1 \text{ mph}}{26.667 \text{ rpm}} \right) \left(\frac{0.447 \text{ m/s}}{1 \text{ mph}} \right) = 0.0168x \text{ m/s}$
4. Green: Potentiometer voltage from wind vane
 - a. Measurement in mV is multiplied by (360°/2500mV) to yield direction:
 - b. 90° \Leftrightarrow East
 - c. 180° \Leftrightarrow South
 - d. 270° \Leftrightarrow West
 - e. 360° \Leftrightarrow North
 - f. Directions were converted to sin(direction) and cos(direction) to allow time-averaging of direction without introducing wrap-around errors near 0°/360°.

The accuracy of the wind vane was listed as $\pm 7^\circ$, while that of the anemometer was listed as $\pm 5\%$.

4.2.4.6.2 Technical Details: Pyranometers

The pyranometers were mounted 8' above the ground away from any local shading. Specifics of the five pyranometers used are listed in the following table:

Serial Number	5829	5822	5801	4626/5626	4951
Orientation	North	East	South	West	Vertical
Resistor Used (Ω)	199.4	199.4	199.3	199.2	199.2
Nominal Calibration (nA/(W/m ²))	65	56.5	-47.3	72.4	-68.6
Final Calibration (nA/(W/m ²))	61.55	53.30	-45.17	52.65	-68.6
Multiplication Factor	81.48	94.09	-111.09	95.36	-73.18

Table 4-3 Pyranometer information

The pyranometers listed in Table 4-3 were calibrated in the following way:

1. 497 sets of measurements were taken with all pyranometers facing directly upwards. The output of the pyranometers is current, so the voltage across resistors of approximately 200 Ω was measured in mV. To convert the measurement in mV to W/m², the following calculation was performed:

$$(\text{Measurement in mV}) \frac{\left(\frac{1\text{V}}{1000\text{mV}} \right) \left(\frac{1}{\text{Resistor}(\Omega)} \right)}{\left(\frac{\text{Calibration number} \times 10^{-9}\text{A}}{\text{W/m}^2} \right)} = \text{Associated Radiation in W/m}^2 \quad [4-2]$$

2. Nominal manufacturer's calibration numbers were used to calculate incident solar radiation in W/m^2 . The pyranometer with the greatest radiation measurement was identified. This was #4951. It was assumed that the output current of the pyranometers would only degrade after the manufacturer's calibration, and consequently that #4951 was closest to its original output.
3. Using the measured data, the calibration numbers of the other four pyranometers were adjusted by least squares to minimize any differences between the radiation measurements of each pyranometer and that of #4951. The final calibration numbers are provided in the penultimate row of Table 4-3. The multiplication factors used to convert the millivolt measurements to W/m^2 (from Equation [4-2]) are provided in the final row.

The pyranometers were affixed to a rectangular metal box with silicone, with the orientation shown in the second row of Table 4-3. Leads were soldered to the Cat.5e cable leading to the data logger.

Pyranometer noise led to non-zero measurements of solar radiation during the hours between sunset and sunrise. The following figure illustrates the situation.




Figure 4-33 North-facing pyranometer data showing non-zero measurements during nighttime hours. The lower plot is logarithmic on the ordinate to accentuate the problem. (SolarFilter.m)

To avoid spurious influence on the model, this noise was removed in the following way:

1. All measurements $< 20 \text{ W/m}^2$ were set to zero.
2. A second pass was made over the data to zero any points with radiation levels $< 60 \text{ W/m}^2$, provided that at least one measurement in the previous three minutes and one measurement in the following three minutes was zero. This step eliminates relatively large, isolated spikes in the data.

Another way to accomplish the same goal would be to calculate local sunset and sunrise times and set to zero any measurements between those times. Such a method would be compatible with real-time operation of the model.

The filtered measurements are shown in the following figure for comparison.




Figure 4-34 Filtered North-facing pyranometer measurements.

4.2.4.7 Campbell Scientific 21X Data Logger

This instrument performed multiple roles during the experiments. It served as a data collection device (*e.g.*, temperature and solar radiation measurements), a DC voltage generator (*e.g.*, control signal to fan controller card, supply to wind vane), and a pulse counter for the anemometer. All actions were performed on a 60-second time step. Measurements were made every second, and were averaged to generate the value of the measurement for the minute preceding the time stamp (in the case of the anemometer, pulses were summed, not averaged).

With the exception of the temperature measurements, all sensors were connected to the 21X using Cat.5e twisted-pair cables. Outgoing signals (DC voltages) were also conducted on this cable. Included in Appendix 2 is a catalogue of all wiring connections to the 21X. A photograph of the logger is provided below.




Figure 4-35 Campbell Scientific 21X Data Logger.

The logger was connected via serial cable to the Dell PC running the experiments. The logger was programmed using PC208W software (Campbell Scientific, Inc). The code loaded onto the logger to run all experiments is provided in Appendix A2.2.

4.2.4.8 Electrical Measurements

In order to provide a measure of the heat gains to the building from electricity consumption, the power delivered to the relevant circuits was monitored and recorded using two electrical power loggers. Average power use on each channel was calculated on a 60-second time step. Provided below are the details of the monitoring equipment and measurements.

A Synergistic Control Systems, Inc. C180E data logger (S/N 1630) with 16 channels was used to monitor circuits in the two Basement electrical panels. A Highland Technology, Inc./Enernet Corporation K20 data logger (S/N 10046) with 8 channels was used to monitor circuits in the Assembly room closet. MagneLab current transducers (*e.g.*, SCT-0750-000 and SCT-1250-000) were used. The loggers are shown in Figure 4-36, below. The building's circuits and details of the logger settings are provided in Table 4-4, below.

The parameter files for the two loggers are included in Appendix 1.





Figure 4-36 Left: Electrical panels in the Basement (Clivus Room). Grey panel at left is the secondary breaker panel, the one at right is the primary breaker panel. The box labeled C is the C180E data logger. Right: Upstairs breaker panel in Assembly Room closet. The box labeled K is the K20 data logger.

Breaker	Label	Phase	CT Number	Max Amps	Channel	Breaker	Label	Phase	CT Number	Max Amps	Channel
Clivus Room Primary Box											
	Primary A (lead to Box)	A	51	100	C0						
	Primary B (lead to Box)	B	54	100	C1						
1	Panel Outlet	B	52	20	C2	2	Parking lighting outside	B	53	20	C3
3	Sub Main (to Secondary Box)	A	58	100	C5	4	Networks outlet	A	56	20	C4
5		B	59	100	C6						
Clivus Room Secondary Box ("Common Room Breaker")											
1	Upstairs main panel	B	48	100	C8	11	Closet light upstairs	B	57	20	C11
2		A	49	100	C7	12	Exit Lamp				
3	Lights Clivus tank room, wood stove					13	Lights and Plug storage	B	46	20	C12
4	Receptacle Common Room and Hot Water					14	Lights storage	A	45	20	C13
5	Fans in pit					15	Temp lights bsmt				
6	Lights Common Rm					16	Receptacle bathroom	A	47	20	C14
7	Fans clivus tank rm	B	60	20	C9	17	Receptacle below panel telephone				
8	Hot Water (Instantaneous)	B/A	55	40	C10	18	Receptacle below panel telephone + Attic	A	50	20	C15
9						19	Range sub panel				
						20					
Upstairs Box (Assembly Room Closet)											
	Upstairs A (lead to Box)	A	42	100	K0						
	Upstairs B (lead to Box)	B	41	100	K1						
1	Recept. assembly					2	Lights and Recept Lobby South and Sunspace				
3	Floor recept assembly					4	Recept. Admission				
5	Recept. Assembly					6	Strip Lights				
7	Recept + lights Conf Rm					8	Spot Lights				
9	Recept E and W office					10	Airlock and Admission Lights				
11	Bath and Stair lights					12	Strip Lights				
13	Track Lights					14	Timer to outside lights	A	44	20	K2
15						16	Fire alarm				
17	Transformer feed--emergency lights					18	Sunspace Fans				
End-use Measurement											
	Attic Fan	B	43	10	K3						
		(mistake)									

Table 4-4 Circuit breaker and electrical logger details. Channels labeled C0-C15 were C180E channels, those labeled K0-K3 were K20 channels. (C10 max amperage was 50A rather than the 40A recorded during experiments. Corrections were made accordingly.)

Note that the phase of the Attic fan measurements (K3) was incorrect. If the channel for the Attic (C15—which was calculated with proper phase) had been used instead, the measured power consumption would have been approximately 5% lower than that shown in Figure 4-19.

To calculate the impact of electrical usage on the two occupied building zones, measurements on different channels were combined according to the location of the circuits in the building. The load in the Basement was calculated as follows:

$$Q_{\text{Basement}} = \text{Primary A (C0)} + \text{Primary B (C1)} - \text{Parking Lighting Outside (C3)} - \text{Upstairs Main A (C7)} - \text{Upstairs Main B (C8)} - \text{Fan Clivus (C9)} - \text{Hot Water (C10)} - \text{Closet Light Upstairs (C11)} - \text{Receptacle Bathroom (C14)} - \text{Attic (C15)}$$

The load in the first floor was calculated as follows:

$$Q_1 = \text{Upstairs Main A (C7)} + \text{Upstairs Main B (C8)} + \text{Closet Light Upstairs (C11)} + \text{Receptacle Bathroom (C14)} - \text{Timer to Outside Lights (K2)}$$

The hot water heater on the first floor (C10) was an instantaneous heater (~9kW) whose typical use led to very brief, large spikes of power usage. To prevent the first-floor signal from being overwhelmed by the noise introduced by the hot water heater, its contribution to the first-floor load was ignored. This may be rationalized by the fact that at least a portion of the hot water exits the building before it warms the building.

The load in Attic 1 from the fan was ignored.

It was necessary to filter the Q_{Basement} measurements because the circuit containing a small storage water heater was not individually monitored. The operation of the heater to offset standby jacket losses led to large power spikes, occurring roughly periodically. These spikes appear in the composite electrical power term: Q_{Basement} . Had the Basement circuits been monitored otherwise, this problem could have been avoided and no filtering would have been necessary. The filtering method proposed and implemented is acausal, so would not be appropriate for real-time operation of the models.

The rationale for eliminating the influence of the water heater from Q_{Basement} is sound since the actual thermal gains to the space arising from the water heater are the roughly constant standby losses. Such a constant term may be readily absorbed by the model.

The situation is depicted in the following figure, where the large spikes of up to 2kW are due to the water heater operation. Note that no phase change is introduced by the filter (middle plot), and that the underlying shape of the actual signal is preserved. Some smaller spikes, not necessarily due to the water heater, were also removed.




Figure 4-37 Basement electrical loads. Top: Actual. Middle: Filtered. Bottom: (Actual – Filtered). (QfilterDemo.m)

The algorithm developed and used to remove the spikes in Q_{Basement} is summarized below²⁸:

1. A 12-minute window is stepped over the entire data set with a step size of one minute. (12 minutes were selected due to the typical duration of the water heater reheat time.)
2. At each step, the 12 points are examined to determine the location of the largest drop in power. Define the time index of the point on the low side of the drop t_{drop} .
3. Compare the power of the first point in the window $P(1)$ with $P(t_{\text{drop}})$. Calculate $P_{\text{max}} = \max(P(1), P(t_{\text{drop}}))$.
4. If any power measurements between point 1 and t_{drop} exceed P_{max} , then set them equal to P_{max} .

Return to step 1 until the beginning of the window reaches the last point of the data set.

²⁸ See SpikeFilter1.m for the MATLAB implementation of the algorithm.

4.2.4.9 Weather Forecasts

Local weather forecasts were automatically downloaded from a NOAA website on an hourly basis²⁹. The following information from the forecasts was relevant to this work:

- Air Temperature (°F)
- Wind Direction (compass direction)
- Wind Speed (mph)
- Cloud Cover (“OV”, “B2”, “B1”, “SC”, “CL”—representing cloud cover ranging from overcast to clear).

Provided in Table 4-5 is a sample of the predictions for Southeast Middlesex county (appropriate for South Natick).

Eastern Essex-Eastern Norfolk-Eastern Plymouth-Southeast Middlesex-Suffolk-																						
347 AM EDT Wed May 14 2003																						
	Wed 05/14/03						\ Thu 05/15/03						\ Fri 05/16/03									
EDT	03	06	09	12	15	18	21	00	03	06	09	12	15	18	21	00	03	06	09	12	15	18
Pop 12hr					40			30					20			20					20	
Qpf 12hr				.10	-.24		.01	-.10					0			0					0	
Max qpf				.10	-.24		.01	-.10					0			0						
Snow 12hr				00	-00		00	-00					00	-00								
Mx/MN					60			49					56			45					52	
Temp	55	57	59	59	56	54	52	49	52	55	55	56	52	48	47	45						
Dewpt	44	47	48	49	50	48	47	47	48	48	47	47	46	44	43	41						
Rh	67	70	67	70	81	81	83	93	87	78	75	72	80	86	86	86						
Wind dir	W	NW	N	E	E	E	NE	NE	NE	E	E	E	E	NE	NE	NE						
Wind spd	5	10	5	10	10	5	2	2	5	15	15	15	10	10	10	10	20	20	20	20	20	20
Clouds	B2	B2	B2	B2	B2	B2	B2	B2	B1	B1	B1	B1	B1	B1	B1	B1	B1	B1	B1	B2	B2	B2
Rain shwrs	C	C	C	C	C	C	C	C														
	\Sat 05/17/03\Sun 05/18/03\Mon 05/19/03\Tue 05/20/03																					
EDT	00	06	12	18	00	06	12	18	00	06	12	18	00	06	12	18	00	06	12	18	00	06
Pop 12hr	30	30			20	10	10	10					10	10								
Mx/MN	46	55			46	63	48	65					49	66								
Clouds	B2	B1			SC	SC	SC	SC					SC	SC								
Rain	C	C																				

Table 4-5 Sample of NOAA weather predictions that were downloaded every hour.

These data were provided for every third hour of the day, extending 48 or 72 hours into the future. Data for the next 24-hour period were interpolated to hourly values, which were assumed constant over that hour. The cloud cover prediction was used in conjunction with a clear-sky solar radiation model, *e.g.*, (McQuiston and Parker 1994; ASHRAE 1997), that calculated incident solar radiation on the faces of the solar radiation sensors (pyranometers). The cloud-cover codes listed above were converted to clear-sky radiation multipliers with the guidance of NOAA staff: 0.1, 0.3, 0.5, 0.8, and 1.0.

²⁹ <http://www.erh.noaa.gov/er/box/fcsts/BOSRDFBOX.html> (format and site address have since changed to: <http://www.erh.noaa.gov/box/fcsts/BOSAFMBOX.html>.)

Since these predictions were ultimately not used in the optimal control portion of the research, rigorous comparisons between forecasts and measurements were not performed.

4.2.4.10 PC Data Acquisition and Control

At the center of the data acquisition and control architecture was a Dell PC running MATLAB (Release 13)³⁰ on a Microsoft Windows 2000 operating system. Due to fatal flaws in MATLAB’s serial port communication routines, additional software was required for serial port communication with the Campbell data logger. Microsoft QuickBASIC v4.5 was selected for its simplicity and functionality. An invaluable reference in programming QuickBASIC was (Nameroff 1989).

Shown below in Figure 4-38 is the overall architecture of the data acquisition and control. Each element will be discussed, with frequent references to appendices for supporting code.

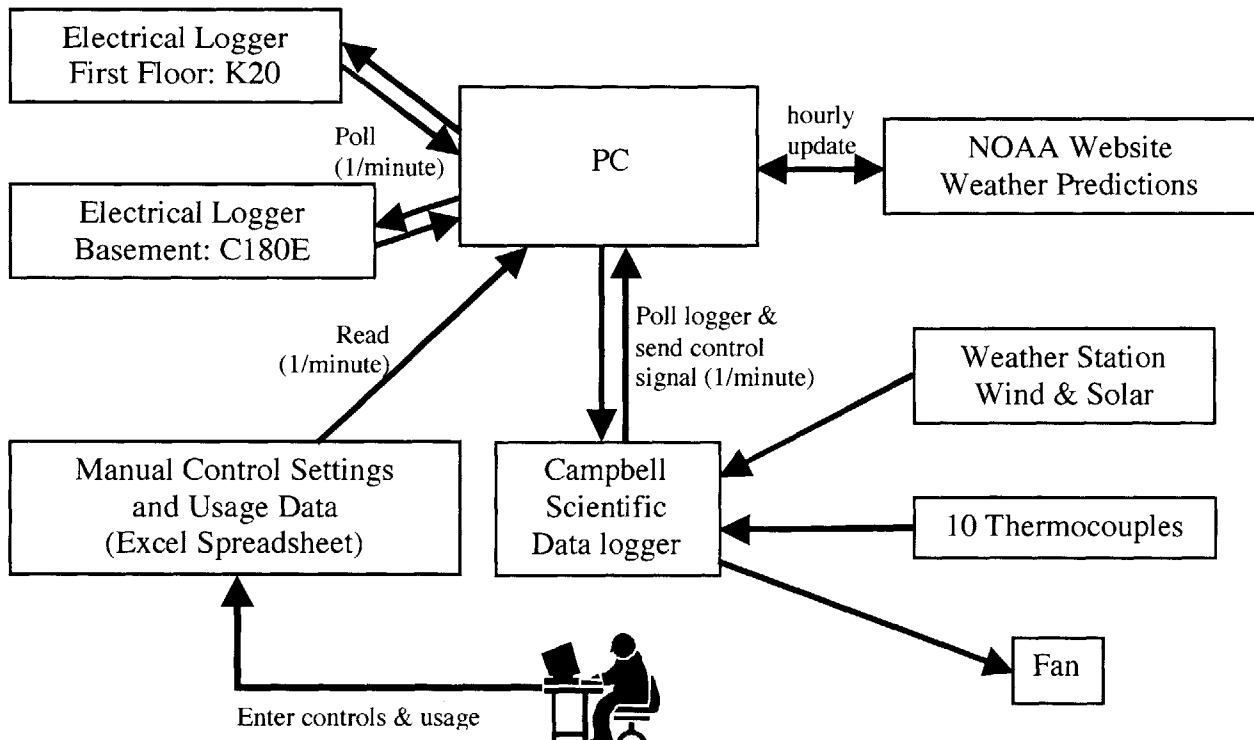


Figure 4-38 Data acquisition and control architecture.

At the core of the architecture was software continuously running in MATLAB. Use was made of timers operating on an hourly and 60-second time step.

4.2.4.10.1 Startup Processes

Among the initial tasks performed were the downloading of all existing data from the three loggers. As mentioned before, all communication with the Campbell logger was performed

³⁰ The Mathworks, Natick, MA.

using QuickBASIC. The essence of the approach for retrieving data from the 21X was to call a QuickBASIC program (from MATLAB) that read the 21X data and then wrote it to a file. The code for reading the data from the 21X is FIRSTDLD.BAS, a text version of which is included in A3.1. The data file generated by FIRSTDLD.BAS was then read into the MATLAB workspace.

The process of downloading the data from the electrical data loggers was facilitated by software from Fishbaugher and Associates, Oakland, OR (kdump.exe v1.3.7). This program was called during the initialization period using C180.BAS and K20.BAS, for the two loggers. Text versions of these codes are provided in A3.3 and A3.5. The parameter files needed by kdump.exe for the C180 and K20 (SC1630.DAT and S10046.DAT, respectively) are provided in A3.4 and A3.6. C180.BAS and K20.BAS were coded to call kdump.exe on minute intervals after completing the initial download³¹.

4.2.4.10.2 Processes Occurring Every Hour

Every hour, the NOAA website was checked automatically for new predictions, and the predictions were stored on the computer. The success of these checks (and other operations) was confirmed by automatic email to the author (and sometimes others).

4.2.4.10.3 Processes Occurring Every Minute

Every minute, the code UPDATE.BAS (text provided in A3.2), was called to retrieve data from the 21X and write it to disk. In addition to downloading data, the code performed a clock check and delivered the desired control signal from MATLAB to the 21X (for transmission to the fan controller card). The file with data from the 21X was then read by MATLAB.

As mentioned above in the initialization section, kdump.exe was called every minute to download data from the two electrical loggers and write it to disk³². These files were read by MATLAB every minute and loaded into the workspace for storage.

In addition to the automatic data collection by the loggers, a manual entry procedure was required to monitor the existing control settings (apertures) as well as the desired settings (fan). Data from a Microsoft Excel spreadsheet was loaded into the MATLAB workspace every minute. Each day during the experimental period had a dedicated spreadsheet with entries for every minute of the day. Table 4-6 provides an excerpt of such a spreadsheet.

³¹ The reason these routines were programmed in QuickBASIC rather than in MATLAB was to allow the processes to run in parallel every 60 seconds, rather than sequentially. The latter approach would have been unnecessarily time-consuming.

³² Due to a bug in kdump.exe associated with the data wrap-around in the loggers' storage areas, MATLAB had to periodically check to determine if kdump.exe was functioning properly, and if not, kill the process and restart it.

minute starting at:	Egrill:Wgrill:Conf/AssDoor:BsmtDoor:Att1Win:Att2Win WinS,WinW,WinAss,AtticSlider,FanSlider	Fan	Setpoint	Override (fan/setpt)	daytype	confirm
0:00	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1
0:01	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1
0:02	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1
0:03	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1
0:04	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1
0:05	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1

Table 4-6 Excerpt of one day's manual monitoring and control spreadsheet.

A detailed explanation of the spreadsheet and its codes is provided in Appendix 4.

4.2.4.11 Graphical User Interface

Monitoring and error-checking the vast amount of data gathered by the procedures described in Section 4.2.4.10 was facilitated by a graphical user interface (GUI) developed for this purpose. The GUI permitted the user to view any of the data collected on a user-specified time period and in the desired units (for temperatures). For instance, any of the electrical measurements (amps, voltage, power on any monitored circuit), all NOAA forecasts, and any of the thermocouple measurements could be viewed. Four windows, each with up to six simultaneous plots were available. New data were incorporated into the graphs every minute. This tool was invaluable. A screenshot is given below in Figure 4-39.

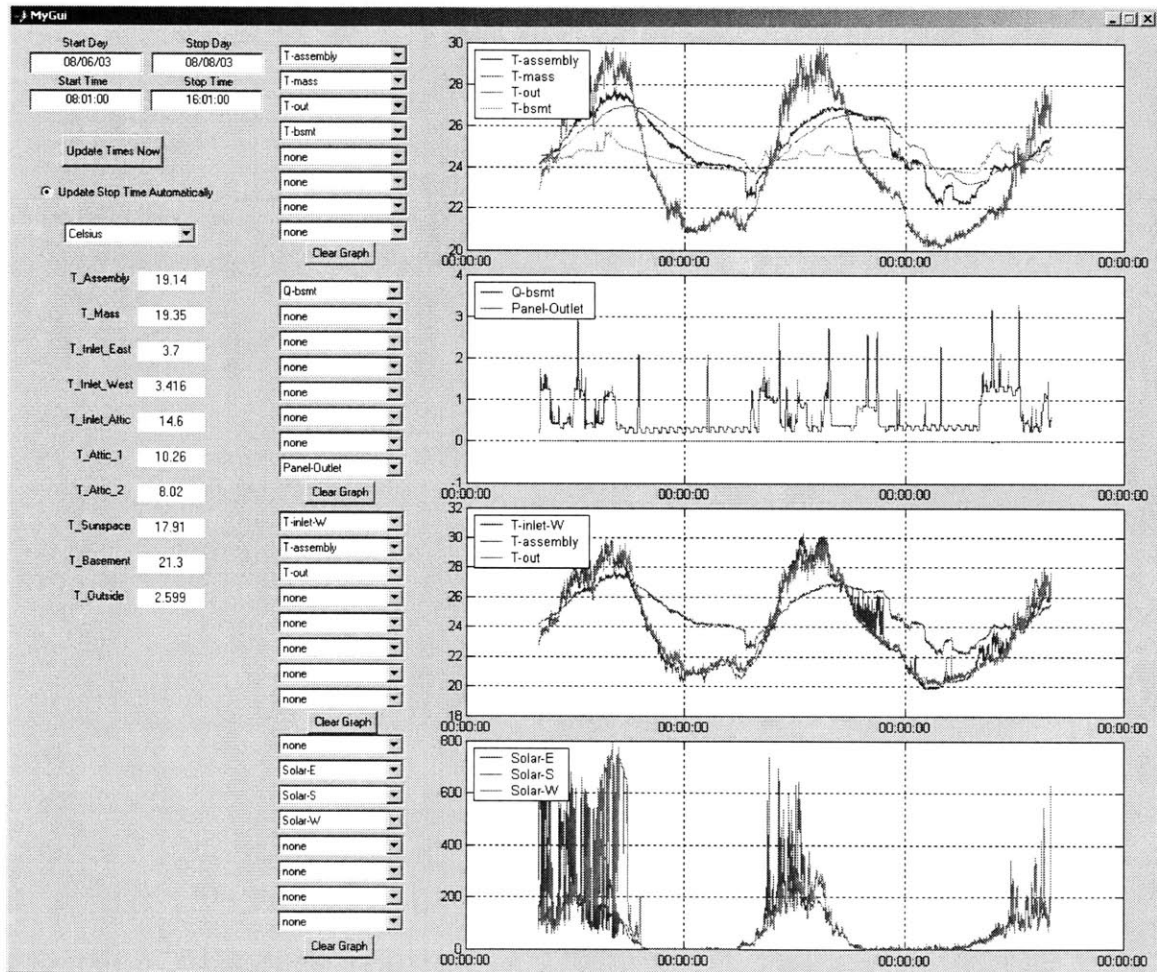


Figure 4-39 Graphical user interface showing data on August 6-8, 2003. The uppermost plot shows several temperature channels. The second plot shows the Basement electrical loads (the saw-tooth pattern is the refrigerator cycling). The third plot shows the temperature at the West louvered door (#2) oscillating within the range defined by $T_{Assembly}$ and $T_{outside}$. The fourth plot depicts the solar radiation incident on the East-, South-, and West-facing pyranometers. (The GUI was color-coded.)

4.2.4.12 Summary of Data Collected

Data were collected over the time period from August 6 to September 30, 2003, though there were several periods when data collection was interrupted or incomplete. The NOAA data collection was least reliable due to multiple changes in the website format, though since real-time implementation of predictive control was not possible, the lack of these data was not reason to discard other data. The five complete data sets (\pm NOAA data) were:

Batch	Start Time	End Time
1	Aug. 06, 06:01	Aug. 08, 16:01
2	Aug. 10, 19:57	Aug. 22, 18:01
3	Aug. 29, 17:00	Sep. 04, 09:00
4	Sep. 04, 12:00	Sep. 26, 17:01
5	Sep. 27, 03:15	Sep. 30, 13:00

Table 4-7 Complete batches of data (some NOAA weather predictions missing).

Data collection was terminated on September 30 when the fans used to move warm Sunspace air into the building were first turned on. This mode of operation was beyond the scope of the project.

In preparation for the heating season, the fan in the Sunspace that moves hot air into the rockbed was turned on at 10:22 on September 12. This fan operated on a thermostat and ran only when the Sunspace temperature was warm enough. Evidence of the operation of this fan may be found in Q_{bsmt} ; the fan's own circuit was not individually monitored. For the purposes of these experiments, the dominant effect of running the fan was to cool the Sunspace.

One small anomaly in the data may be pointed out on September 14. A large spike in the Basement temperature was associated with repairs requiring that the doors be opened to the outside (#12) and possibly between the Basement and the Sunspace.

A large set of measurements was recorded every minute (1440 points per day). Table 4-8 below details all of the measurements that were made. Some of these were used (viewed) only in the GUI, while others, marked with "x", were identified as potential model inputs or as potential components of model inputs.

ID	Recorded Measurement	Potential Modeling Input	ID	Recorded Measurement	Potential Modeling Input	ID	Recorded Measurement	Potential Modeling Input
1	T-assembly	x	45	Office-wind-S-act (#4)	x	87	Primary-A	
2	T-mass	x	46	Office-wind-W-act (#5)	x	88	Primary-B	
3	T-inlet-E	x	47	Attic-slider-act (#6)	x	89	Panel-Outlet	
4	T-inlet-W	x	48	1/0 Fan-slider-act (#7)	x	90	Parking-Outside	
5	T-inlet-attic	x	49	Fan speed-act (#8)	x	91	Networks	
6	T-attic1	x	50	2/1/0 Sunspace-door-act (#9)	x	92	Secondary-A	
7	T-attic2	x	51	1/0 Attic2-wind-act (#10)	x	93	Secondary-B	
8	T-sunspace	x	52	Assembly-wind-act (#11)	x	94	Upstairs-Main-A	
9	T-bsmt	x	53	1/0 Bsmt-Door-act (#12)	x	95	Upstairs-Main-B	
10	T-out	x	54	Blank2-act		96	Fan-Clivus	
11	T-out-NWS		55	Blank3-act		97	HotWater	
12	Solar-N	x	56	Predicted Qinternal bsmt	NA	98	Upstairs-Light	
13	Solar-E	x	57	Predicted Qinternal first	NA	99	Lts/Plug-Storage	
14	Solar-S	x	58	Setpoint		100	Lts-Storage	
15	Solar-W	x	59	1/0 E-grill-des		101	Bath-Recept	
16	Solar-horiz	x	60	1/0 W-grill-des		102	Attic	
17	Solar-Dir-N-NWS		61	2/1/0 Conf-door-des		103	Volts-A-C180	C180
18	Solar-Dir-E-NWS		62	Office-wind-S-des		104	Volts-B-C180	
19	Solar-Dir-S-NWS		63	Office-wind-W-des		105	Amps-Primary-A	
20	Solar-Dir-W-NWS		64	Attic-slider-des		106	Amps-Primary-B	
21	Solar-Dir-horiz-NWS		65	1/0 Fan-slider-des		107	Amps-Panel-Outlet	
22	Solar-Dif-N-NWS		66	Fan speed-des		108	Amps-Parking-Outside	
23	Solar-Dif-E-NWS		67	2/1/0 Sunspace-door-des		109	Amps-Networks	
24	Solar-Dif-S-NWS		68	1/0 Attic2-wind-des		110	Amps-Secondary-A	
25	Solar-Dif-W-NWS		69	Assembly-wind-des		111	Amps-Secondary-B	
26	Solar-Dif-horiz-NWS		70	1/0 Bsmt-Door-des		112	Amps-Upstairs-Main-A	
27	Cloud Cover-NWS		71	Blank2-des		113	Amps-Upstairs-Main-B	
28	sinWind-Dir	x	72	Blank3-des		114	Amps-Fan-Clivus	
29	cosWind-Dir	x	73	Campbell datenum		115	Amps-HotWater	
30	Wind-Speed	x	74	Campbell day num		116	Amps-Upstairs-Light	
31	Wind-Dir-NWS		75	Campbell hour		117	Amps-Lts/Plug-Storage	
32	Wind-Speed-NWS		76	Campbell secs		118	Amps-Lts-Storage	
33	Q-bsmt	x	77	C180 datenum		119	Amps-Bath-Recept	
34	Q-first	x	78	C180 day num	Timings	120	Amps-Attic	
35	Q-fan	x	79	C180 hour		121	Upstairs-A (K20)	
36	1/0 Office day weekday	x	80	C180 secs		122	Upstairs-B (K20)	
37	1/0 Camp/Office day weekday	x	81	K20 datenum		123	Timer-OutsideLts (K20)	
38	1/0 Holiday/Weekend	x	82	K20 day num		124	Attic-Fan (K20)	
39	1/0 Assembly Meeting	x	83	K20 hour		125	Volts-A-K20	K20
40	time of day	x	84	K20 secs	126	Volts-B-K20		
41	day of year		85	NWS day num	127	Amps-Upstairs-A (K20)		
42	1/0 E-grill-act (#1)	x	86	NWS hour	128	Amps-Upstairs-B (K20)		
43	1/0 W-grill-act (#2)	x			129	Amps-Timer-OutsideLts (K20)		
44	2/1/0 Conf-door-act (#3)	x			130	Amps-Attic-Fan (K20)		

Table 4-8 List of all measurements recorded on a 60-second interval. The “x” marks in every third column represent measurements that were treated as potential inputs (or components thereof) of the models in the following chapter.

(The entries ending with “-des” were intended to be the desired control settings as determined automatically by the control strategy optimization. Since this was not implemented in real-time, only the actual settings were relevant—those ending with “-act”.)

4.2.5 Comparison with Prior Measurements

A comparison of the relationship between inside and outside temperatures in 2001³³ versus that in 2003 was performed. Shown in Figure 4-40 are plots of Assembly room temperature versus outside temperature for the full data sets of 2001 and 2003 (excluding the period when the Sunspace door was open).




Figure 4-40 Comparison of Assembly temperature versus Outside temperature in 2001 (above) and 2003 (below). (Data in upper plot are identical to those in Figure 198 of (Chang 2002).) (Comparison2001_2003.m)

A linear fit of the 2001 data yields $T_{\text{Assembly}} = 0.435T_{\text{outside}} + 14.795$, while for 2003, the best linear fit is $T_{\text{Assembly}} = 0.448T_{\text{outside}} + 13.262$. These fits imply that, on average in 2001, when the outside temperature exceeded 26.1°C, the Assembly temperature was cooler than the outside temperature. In contrast, in 2003, the Assembly temperature was on average cooler than the outside temperature whenever the outside temperature exceeded 24°C. In other words, it was relatively cooler in the Assembly room for a given outside temperature in 2003 than in 2001.

Several changes in the experimental setup may contribute to the observed difference. First, the temperature sensors were placed in different locations. In 2001, the outside temperature sensor was in a sheltered area about 6' above the ground on the East side of the building. The Assembly room temperature sensor was at approximately the same height as in 2003, but was located about 12' to the East. Second, the physical alterations to the building (unobstructed cupola and enlarged opening to the Attic) may have contributed to the improved comfort within the building as well. Finally, the method of controlling the building's apertures during the

³³ Data courtesy of J.-C. R. Chang.

experiments in 2003 played a role in the improved average temperatures. Recall that, in 2001, no changes were made to the primary openings (louvered doors and opening to Attic) during the experimental period.

To further illustrate the differences between the data from the two years, the data were binned by outside temperature: [24, 26), [26, 28), ... [32, 34)°C. For each outside temperature bin, the average and standard deviation were calculated of the difference ($T_{out} - T_{Assembly}$). The results are presented in Figure 4-41. From this figure, the greater difference between outside and Assembly temperatures observed in 2003 is evident. The standard deviations in 2003 are also smaller, possibly to due improved control of the building.




Figure 4-41 ($T_{out} - T_{Assembly}$) versus outside temperature, using 2°C bins. Full data sets were used to generate plots. Data points are offset from the centers of the bins (located at odd-numbered temperatures) for clarity. (Comparison2001_2003.m)

Finally, when the 2001 data from the months of August and September exclusively were used (the only months for which data were available in 2003), the results of these comparisons were more striking. In August and September 2001, the outside temperature above which Assembly temperatures were cooler than outside temperatures was 26.5°C. Plots analogous to those in Figure 4-41 are shown below in Figure 4-42.




Figure 4-42 ($T_{out} - T_{Assembly}$) versus outside temperature, using 2°C bins. Data from August and September only were used to generate plots. Data points are offset from the centers of the bins (located at odd-numbered temperatures) for clarity. (Comparison2001_2003.m)

4.3 Summary

In this chapter, the Saltonstall Nature Center at the Massachusetts Audubon Society's Broadmoor Wildlife Sanctuary was introduced. Building modifications were described and complete details of the experimental instrumentation and execution were provided.

In the next chapter, the 2003 Broadmoor data sets are used to generate models that predict the thermal behavior of the building.

5 Broadmoor: Thermal Modeling

In this chapter, the PHDRT algorithm described in Chapter 3 will be applied to the data gathered at Broadmoor, with the intent of creating high-fidelity models for ultimate use in evaluating control strategies in the building. Unexpectedly, this machinery is not essential to develop accurate thermal models for Broadmoor. A straightforward alternative approach, which incorporates linear regression with simple nonlinear terms in the information matrix, is developed and shown to be highly effective.

Following the literature review, the first sections of the chapter lay the groundwork, establishing the vocabulary and framework for the remainder of the chapter. The remaining sections are devoted to the modeling of the system as well as to understanding some of the model properties.

5.1 Literature Review

The research described in subsections 5.1.1 and 5.1.2 concerns the development of models for particular buildings whose indoor climate was modified by natural ventilation or by mixed-mode cooling. The ultimate purpose of constructing the models was to control the modeled spaces or to evaluate various control strategies.

Two types of models are prevalent in the literature: physically-based models and data-driven models. These are also known as forward and inverse models, respectively. The principal benefit of the former model type is that these models may be used to guide design. The shortcoming of the physically-based models is that their structure and the relevant physical parameters are often unknown, or poorly known. To achieve high performance (*e.g.*, for use in optimal control), physical models often require extensive calibration with measured data. On the other hand, data-driven models, by their very nature, do not exist before the building is built. They cannot be used for design purposes as they are developed for a particular building, perhaps under particular conditions. However, data-driven models may achieve higher levels of accuracy without requiring of the user painstaking identification, estimation and entry of relevant building parameters.

In the third subsection, 5.1.3, the scope is expanded to the data-driven modeling of a broader set of buildings, including those whose indoor climate is modified by traditional mechanical systems. Finally, in subsection 5.1.4, a summary of the existing literature is provided.

5.1.1 Physically-based Models

A model based on single-zone analytical airflow calculations from CIBSE (1997) and a three-time-constant thermal model (Tindale 1993) was developed for the purpose of evaluating night-cooling strategies using natural ventilation (Kolokotroni 2001). The design-stage tool was applied to a portion of an educational building that was naturally ventilated using stack, single-sided and cross ventilation. Inside temperature predictions were compared with measured data and found to display modest agreement. Maximum and minimum temperatures were over- and under-predicted by approximately 1K and 2.5K.

The Grong School in Norway was modeled (Jeong and Haghghat 2002) using ESP-r (ESRU 1997). For the period modeled, the building was in the heating season, with several modes of

heating available: radiators and preheated ventilation air. The building was modeled with 21 zones. As above, agreement between predicted and measured temperatures was modest. In this instance, the measured exhaust air temperatures were compared with predictions. Discrepancies over a single day were as great as 3-4K. One of the complications encountered was that there was no way to incorporate into the simulation tool the CO₂-concentration-driven damper control governing room airflow rates. This shortcoming was thought to contribute to the poor predictive performance of the model. Furthermore, well-mixed conditions were assumed (by ESP-r) despite the fact that the building used displacement ventilation.

An analytical airflow model was coupled with TRNSYS (Klein 1990) to predict the thermal behavior of a single room in Athens cooled by natural ventilation through a single operable window (Dounis *et al.* 1996). The model was developed to investigate the use of fuzzy control strategies to regulate indoor air quality. To calibrate the model to experiments, a correction factor was added to the calculation of flow rates. The authors reported an average prediction error for the room air temperature of $0.29 \pm 0.38^{\circ}\text{C}$ for a simulation over a single day³⁴.

While no comparison with experimental data was provided, an equation-based thermal model of a single room with an operable window with an adjustable blind was presented by Dounis *et al.* (1994). The model was constructed to evaluate the impact of visual comfort control on thermal comfort within the room.

Another group (Bruant *et al.* 2001) used TRNSYS to model the temperature of individual mixed-mode test cells in Lyon. The model was used to design a fuzzy controller for thermal and air quality control. This room was cooled by a combination of fan-driven ventilation air (at three speeds) and a mechanical air conditioner. Over three test periods, average prediction errors varied from $1.59 \pm 1.38^{\circ}\text{C}$ to $1.95 \pm 1.36^{\circ}\text{C}$. The authors suggested that the large bias in the predictions may have been due to the difference between the air inlet temperature and its (remotely) measured value.

Skrjanc *et al.* (2001) compared the performance of an analytical model and a Takagi-Sugeno fuzzy model (see, *e.g.*, Jang *et al.* 1997) in predicting temperatures within a small test chamber. The chamber consisted of a fixed window with an operable shade, which was moved during experiments. Error for the analytical model over several one-day periods was reported to be in the range of 5-20%. Error figures were not reported for the two types of fuzzy models developed, though in simulations of a single day, errors of up to 10°C (approximately 50%) were shown. (This example did not involve natural ventilation or mixed-mode cooling, but it appeared relevant to the discussion.)

While not developed for the purpose of temperature control, Chang (2002) developed a second-order model for the Broadmoor Wildlife Sanctuary. Its performance is compared with that of the models developed in this thesis in Section 5.1.3.

³⁴ Note that this is an average error rather than an rms error.

5.1.2 Data-driven Models

Several authors have applied neural networks to the localized modeling of airflow and temperatures within rooms. In one study, a test chamber with four fixed openings was equipped with temperature and velocity sensors in multiple locations (Kalogirou *et al.* 2001). Models were developed to map the following inputs: outside temperature, humidity, pressure, wind speed and direction and the air velocity and direction through the openings to the following outputs: temperatures and velocities at the various locations in the room. No dynamics were incorporated into the steady-state model. Temperature predictions were within 2.3°C of measured values. Others have used neural networks to model the effect of room and aperture geometry on interior air velocities (Kindangen 1996; Krauss *et al.* 1997). In these cases, simulations with computational fluid dynamics rather than experiments provided data for model training.

Through experiments designed to inform PI controller design for temperature control in a single naturally ventilated room, it was found that flow through the upper and lower openings of the test room's one window was independent of local wind speed, wind direction and room temperature (van Paassen and Lute 1993). Flow was linearly related to the size of the (tilted) window openings, or in other words, proportional to the total window area multiplied by the sine of the opening angle. Such a result would not have been predicted from analytical methods. The control strategy was used in conjunction with a thermal model to assess the heat removal potential of the system.

An ARMAX model was constructed using data from measurements in a naturally ventilated test cell in Delft (Lute and Paassen 1995). The single-room zone was equipped with a moveable awning for shade, an automatic window with two openings and a heater. All three components influenced the temperature of the cell, for which a linear predictive control strategy was developed. Nonlinear trigonometric equations were used to calculate the solar radiation striking the window based on solar and awning positions. The ventilation contribution to the heat balance at the interior temperature node was proportional to the window opening area multiplied by the inside-outside temperature difference. The authors hypothesize that the significant model prediction errors observed at times of large window openings were related to the linear approximation made of the ventilation cooling effect and the fact that the window was minimally open for the bulk of the training data. The model was used primarily to make one-step-ahead predictions, but in one instance it was shown to make accurate predictions for four time steps (15 minutes each) into the future. Further details of the modeling and control strategies are available in (Lute 1992).

Considerable effort has been devoted to the closely related problem of model building for greenhouse climate control.

Neural networks were shown to provide an accurate mapping of wind speed, wind direction, greenhouse opening area (10 levels) and fan speed (32 levels) to the ventilation rate of the greenhouse (Linker *et al.* 1999). In this instance, CO₂ was used as a tracer gas for flow rate measurements. A simple first-order thermal model of the greenhouse incorporated the flow rates as predicted by the neural network. The flow model was used in a robust control framework to control both temperature and CO₂ concentration in the greenhouse.

Ferreira published several articles regarding temperature prediction models using radial basis function neural networks (RBFNNs) in greenhouses. In (Ferreira and Ruano 2002), RBFNNs were used to map inside temperature, outside temperature, inside humidity and outside solar radiation to the inside temperature. A variety of time lags on the different inputs were investigated in an attempt to establish the most effective set of inputs to use. The best model was found to employ radiation inputs at times k and $k-21$, humidity inputs at times k and $k-3$, outside temperature inputs at times k and $k-2$ and inside temperature inputs at times $k-1$ and $k-2$. The time step was 5 minutes. Over a 3257-point test set, the pure simulation rms error (defined below in Section 5.7.1.1) was 1.48°C . In (Ferreira *et al.* 2002), a variety of on-line and off-line training methods were discussed for the same greenhouse problem. It was found that a Levenberg-Marquardt training algorithm outperformed all others investigated for on-line or off-line training in terms of error performance, model size and generalization capability. Finally, in (Ferreira *et al.* 2003), multi-objective genetic algorithms were employed to assist in the selection of inputs (and the best number of lag terms to use) as well as the RBFNN architecture. For the same greenhouse—single zone with no operable components—model prediction rms error was improved to 1.30°C for the full test set. In almost all models formed, the wind speed and direction were not identified by the genetic algorithm as useful inputs. One additional input used in this last study was the solar radiation measured at an inside location.

Seginer has also devoted much work to the thermal modeling of greenhouses using feedforward neural networks. Seginer *et al.* (1994) formed a neural-network model (with no dynamics) mapping outside solar radiation, outside dry-bulb temperature, outside wet-bulb temperature, wind speed, heater heat flux, aperture opening angle, water misting time fraction, leaf area, Julian day and hour of day to the inside dry-bulb and wet-bulb temperatures, the inside solar radiation and the soil temperature of an Avignon greenhouse. Another model was formed for a Silsoe greenhouse (inputs: outside solar radiation, outside long-wave sky radiation, outside dry-bulb temperature, outside dew-point temperature, wind speed and direction, and aperture opening angle; outputs: inside solar radiation, inside dry-bulb and dew-point temperatures, crop canopy temperature and soil temperature). By examining weights in the input layer of the neural nets, the authors argued that certain inputs were less useful than others and could be discarded. In an additional study of ventilation rates, the authors found that wind direction was not a relevant input; nor was the inside-outside temperature difference. This finding was suggested by the input weights of these model inputs and confirmed by the training of neural networks without those inputs. The best training errors reported for the inside dry-bulb temperature predictions at Avignon and Silsoe were 0.93 and 1.16°C . No errors on a test data set were reported.

Seginer (1997) presented a general discussion of the application of neural network modeling to greenhouses for the purpose of climate control. He discussed breaking a data set into odd and even points for training and validation to prevent overtraining and to capture the system's behavior under all observed conditions. The odd and even sets are not statistically independent, however. The alternative approach—breaking the data set into two distinct groups (time 1 to time 2, then time 2 to time 3)—may be problematic for nonlinear models such as neural networks since they do not extrapolate well to conditions not observed in the training data. He also suggested using a form of neural network called “bottleneck neural networks” (Kramer 1991; Oja 1991) to reduce model order and, consequently, the number of parameters to train.

5.1.3 Data-driven Modeling of Conventional Buildings and Building Systems

Numerous publications have been devoted to the application of neural networks to the modeling of buildings and building systems for the purpose of load and temperature estimation as well as process variable estimation. Several reviews of the topic include (Kreider *et al.* 1992; Kalogirou *et al.* 2001; Krarti 2003).

Ahmed *et al.* (1996; 1998a; 1998b; 1998c) showed that Generalized Regression Neural Networks (GRNN) were capable of identifying the (static) nonlinear mapping between damper or valve positions and flow rates for a given damper or valve authority. This information was used to improve control of a laboratory VAV system. Li and Braun (2002) has also used GRNN in conjunction with polynomial models to generate models for rooftop air conditioner fault detection.

Building energy consumption prediction has been the focus of many authors investigating neural networks. The “Great Energy Predictor Shootout” (Kreider and Haberl 1994) spawned much research in this area, most implementing some form of neural networks to predict electricity, chiller water energy and hot water energy consumption of two commercial buildings (Feuston and Thurtell 1994; Iijima *et al.* 1994; Kawashima 1994; MacKay 1994; Ohlsson *et al.* 1994; Stevenson 1994; Dodier and Henze 2004).

Curtiss has also demonstrated the use of neural networks to predict the power consumption of a central plant (chiller compressor, condenser fans, supply air fan and chilled water pumps) given set points and uncontrolled variables such as outdoor temperature (Curtiss 1992; Curtiss *et al.* 1993b; Curtiss *et al.* 1994). Chiller power consumption and ice storage tank charge/discharge rates were also modeled for a central plant with ice storage capabilities (Massie *et al.* 1998). These models were utilized by a neural network supervisory controller to determine process set points that led to minimal energy consumption.

Curtiss also presented evidence that neural networks could be valuable tools for modeling individual processes, such as the behavior of a hot water coil (Curtiss 1992; Curtiss *et al.* 1993a; Curtiss *et al.* 1993b). These models were integral to a predictive local control strategy. A method of updating the models over time was suggested whereby data on which the model performed poorly was set aside for future training (Curtiss 1996).

Another model developed for local control purposes was that of an air-cooled chiller condenser (Henze and Hindman 2002). The model architecture used was an RBFNN.

The only instance (of which the author is aware) where neural networks were used to predict room temperatures in occupied buildings was presented in (Teeter and Chow 1998), where a functional link neural network was trained to predict room temperatures as generated by a simulation. The essence of the functional link neural network is that connections between nodes are eliminated if the associated weights are smaller than a certain tolerance. Additionally, the inputs to the network are those associated with a second-order polynomial kernel (see, *e.g.*, (Gunn 1998)) rather than on the inputs themselves. The authors suggested that it was beneficial

to train the model to predict the temperature *change* occurring over the next time step rather than the temperature at the end of the next time step³⁵.

An example of the use of fuzzy modeling for HVAC-related problems is provided by Sousa *et al.* (1997b), where the exiting air temperature from a fan-coil unit is modeled for use in an internal model control (IMC) scheme.

While their use has not been published in the field of HVAC or building modeling, support vector regression (SVR) and the related kernel recursive least squares (KRLS) may prove to be useful tools in dynamic and static modeling. The usefulness of the tools in time-series prediction has been demonstrated in (Mueller *et al.* 1997; Mueller *et al.* 1999; Engel *et al.* 2004).

Linear models have also been widely used to make predictions of building thermal response.

Braun developed a “robust inverse model” to model both commercial building loads and temperatures (Braun *et al.* 2001; Chaturvedi and Braun 2001; Lee and Braun 2004). The boundaries of the zone were modeled with 2-C/3-R models. The parameters of all building components were assigned a range of potential values. Within the range so defined, a global direct search (Aird and Rice 1977) selected the building parameters to minimize the difference between measured loads and those predicted by the transfer function model. A local, nonlinear regression was used to optimize estimated parameters further. The authors reported that the predicted loads of a simulated building (with well-known inputs) were within 2% of the actual loads. A two-week training period was used, followed by a four-week test period. For a real building, the loads were predicted to within about 10%. Temperatures were predicted to within 1°F (approximately 8% of full range). Braun *et al.* (1990; 2001) also presented polynomial models to model power consumption of the air handling unit and the total cooling plant. The models were used to investigate optimal control of thermal storage.

Coley and Penman (1992) proposed a method of modeling a building using a 2-C/3-R model, and demonstrated the performance of recursive least squares for identifying the parameters during building operation. Errors on training data only were provided: 1°C rms error. The information required to develop the model consisted of inside and outside temperatures and all heat gains acting on the inside temperature.

Madsen and Holst (1995) proposed a 2-C/2-R model for the simple building in their study, whose temperature was influenced by solar radiation, electrical heaters and the outside temperature. They argued that a continuous-time-domain modeling approach was appropriate and provided a method for parameter estimation using maximum likelihood estimation. Error statistics were not provided for the pure simulation performed.

Dexter and Haves (1989) presented a first-order ARMAX model for use in self-tuning control of HVAC systems. In order to ensure that the parameters identified on-line did not acquire unreasonable values, a series of “jacketing” rules were implemented. Model identification

³⁵ This recommendation may be based, in part, on the authors’ experience training a model on noise-free data generated by simulation.

performance degraded over time, leading to a reduction of control quality. The authors suggested that this shortcoming was the consequence of continuously updating the parameters.

Chen (1997; 2001) presented a fourth-order time-series model for a test chamber heated by a radiant floor panel and by solar radiation. The model was generated and updated continuously using recursive least squares coupled with supervisory rules for robust estimation. The rms prediction error over a 24-hour simulation was 0.27°C. It was not stated whether the simulation was over training or validation data. The model was developed for optimal predictive control of space temperature.

Armstrong (2004) developed a linear time-series model for a one-zone space employing thermodynamic constraints to ensure physically valid results.

5.1.4 Literature Review Summary

A wide variety of techniques have been demonstrated in modeling the thermal behavior of buildings. Physically-based models encountered difficulties due to the limitations of the simulation tools' capabilities, unknown physical parameters and the need to calibrate the model based on observed differences with predictions and measurements.

With the exception of the work by Lute and van Paassen, data-driven models for naturally ventilated and mixed-mode buildings were based on neural networks or fuzzy inference systems. The systems examined were all single-zone rooms or buildings. The flexibility of the nonlinear models was used to incorporate the thermal behavior of the space when it was operated in different modes. The ARMAX model used by Lute and van Paassen did not explicitly account for the different operational modes of the one-room test chamber. The issue of model input (or feature) selection was addressed directly by Ferreira and, to a lesser degree, by Seginer.

Data-driven modeling of conventional buildings was also restricted to single-zone buildings. In the work reviewed, the issue of model input selection was not addressed and the issue of multiple operating modes (associated with different airflow patterns) was irrelevant. Operational mode switching did exist (for example, in a room cooled by a VAV-system), but the only impact on the zone was a potentially variable amount of heat delivered to the room.

In the remainder of this chapter, a data-driven multi-zone model will be developed for a building with multiple active operating modes. Broadly speaking, the building is much more complex than those examined in the literature, particularly those for which a model's predictions have been compared with measured data. As will be shown, the accuracy of the predictions for this building is as good or better than any reported above.

5.2 Definition of Modeling Terminology

In this thesis, the term “feature” will be borrowed from the machine learning vocabulary and will be used to describe an input to a model. The full set of inputs will be referred to as the “feature vector”. When a model is trained, the model parameters are adjusted so that the errors—the difference between the model outputs and the “targets” provided—are minimized. See Figure 5-1.




Figure 5-1 Illustration of model training, with Features 1-10 serving as inputs to the model, and Target(t) providing guidance in setting the model parameters. Shown in the figure is a single training pair: Feature Vector(t-1) and Target(t). Multiple pairs (for all relevant times, t) are used throughout model training.

A feature may be a raw measurement, such as a temperature, or may be a composite term such as one related to heat input to a space when the fan is running: $(\text{fan speed}) * (T_{\text{outside}} - T_{\text{Assembly}})$. Wherever practical, a full listing of the features used to generate a figure or table will be included for reference.

After training is complete, the model is given new features (typically different from those used to train the model), and adjustment of model parameter ceases. The output of the model may or may not be compared with targets. See Figure 5-2.

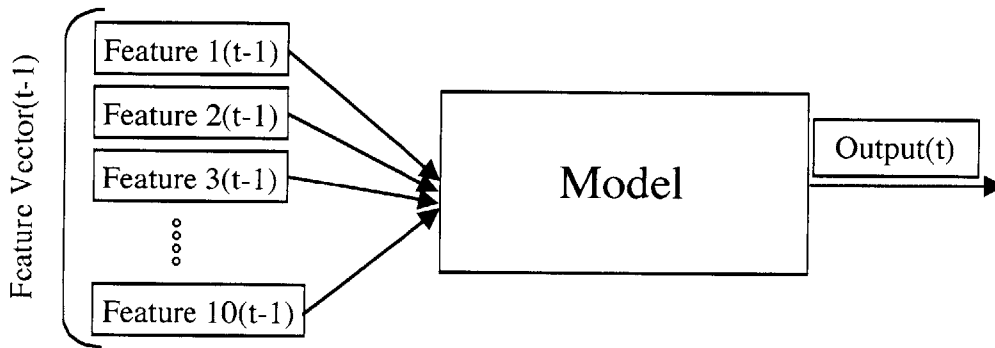


Figure 5-2 Illustration of model use after training is complete. No parameter adjustment is performed.

In this chapter, the targets for the model outputs are the six “zone” temperatures:

- 1) Assembly
- 2) Mass
- 3) Attic 1
- 4) Attic 2
- 5) Sunspace
- 6) Basement

The Assembly temperature is representative of the occupied space temperature to be controlled—the other zones are needed to varying degrees (as will be shown) to improve predictions of Assembly temperatures.

The role of the model is to map the Feature Vector at time $t-1$ to the associated outputs at time t , those outputs being the temperature of a given zone at time t . In other words, given the information available at time $t-1$, what will be the state of the system at time t ? Recall that the step size used in the previous chapter was one minute.

In equation form, the structure of a linear model implementation is given below.

$$\begin{aligned}
 T_1(t) &= a_{11}T_1(t-1) + a_{12}T_1(t-2) + \dots \\
 &\quad + a_{21}T_2(t-1) + a_{22}T_2(t-2) + \dots \\
 &\quad \dots \\
 &\quad + a_{n1}T_n(t-1) + a_{n2}T_n(t-2) + \dots \\
 &\quad + b_{11}I_{sun,horiz}(t-1) + b_{12}I_{sun,horiz}(t-2) + \dots \\
 &\quad + b_{21}I_{sun,south}(t-1) + b_{22}I_{sun,south}(t-2) + \dots \quad [5-1] \\
 &\quad \dots \\
 &\quad + c_{11}Fan(t-1)(T_{out}(t-1) - T_1(t-1)) + \dots \\
 T_2(t) &= \textit{similar} \\
 &\quad \dots \\
 T_n(t) &= \textit{similar}
 \end{aligned}$$

In Equation [5-1], T_1 refers to Assembly temperature, T_2 to Mass temperature, and T_n to Basement temperature. Some exogenous inputs ($I_{sun,horiz}$ and $I_{sun,south}$) are given as examples of the many exogenous system inputs. The influence of controlled elements is shown (e.g., the fan term). Note the time terms and the coefficient notation. The prediction of T_1 at time t is a function of T_1 at previous times. If one lag term is important for accurate predictions, then the term $T_1(t-1)$ would be included in the model and the associated coefficient a_{11} would be determined through model training. If five lag terms are important, then $T_1(t-1)$, $T_1(t-2)$, $T_1(t-3)$, $T_1(t-4)$ and $T_1(t-5)$ would all be included, along with the associated coefficients: a_{11} , a_{12} , a_{13} , a_{14} and a_{15} .

Note that the prediction of $T_i(t)$ may also require terms such as $T_i(t-k)$, where $k \in \{1, 2, \dots, p\}$, with p referring to the maximum number of lag terms. To avoid the complexity involved in solving all six zone temperatures simultaneously, it was decided that $T_j(t)$ would have no influence on $T_i(t)$ ³⁶. In other words, the index k may never be zero for temperature terms. It may

³⁶ Of course this is not strictly true as the temperature measurements obtained were averages over the previous minute. The interactions among zone temperatures during a particular minute are ignored.

be fruitful to allow the index k be zero for the exogenous inputs³⁷. For simplicity, this will be avoided.

The controlled inputs had a constant value over each time step, so for these inputs, $k \in \{1, 2, \dots, p\}$ were the only reasonable options.

5.3 Normalization

To avoid scaling problems in model construction, all features and outputs were scaled to the range of approximately [0, 1]. For each feature, the observed range of values was established, e.g., from $feature_{min}$ to $feature_{max}$. The scaled feature was determined:

$$\text{scaled feature} = \frac{(feature - feature_{min})}{(feature_{max} - feature_{min})} \quad [5-2]$$

Other forms of scaling could be used, such as

$$\text{scaled feature} = \frac{feature}{\max(|feature_{max}|, |feature_{min}|)} \quad [5-3]$$

$$\text{scaled feature} = \frac{2(feature - feature_{min})}{(feature_{max} - feature_{min})} - 1 \quad [5-4]$$

Equation [5-4], which normalizes data to the range [-1, 1], is used by MATLAB's `premnmx.m` routine.

Whenever outputs or output errors are reported in this document, the scaling is removed first.

Rather than implementing different scaling factors for each temperature measurement, all temperatures were scaled using the same range. A summary of the scaling ranges associated with the elementary (not composite) features is provided below. The ranges were generally selected to be slightly greater than the full range actually observed.

³⁷ One could make the argument that when using the model in a predictive (rather than retrospective) fashion, all exogenous variables must be predicted for future times. In that case, it is no more difficult to predict $I_{sun, horiz}$ at t than at $t+5$, for example. Therefore, $I_{sun, horiz}(t)$ could be used to help predict $T_i(t)$.

ID	Feature	Scaling Range	ID	Feature	Scaling Range
1	T_assembly	[275, 330]	16	Q_bsmt	[0, 3]
2	T_mass	[275, 330]	17	Q_first	[0, 7]
3	T_attic1	[275, 330]	18	E Louver Door	[0, 1]
4	T_attic2	[275, 330]	19	W Louver Door	[0, 1]
5	T_sunspace	[275, 330]	20	Conf/Ass. Door	[0, 2]
6	T_basement	[275, 330]	21	Office Windows S	[0, 1]
7	T_out	[275, 330]	22	Office Windows W	[0, 1]
8	Solar-N	[0, 300]	23	Attic Slider	[0, 1]
9	Solar-E	[0, 800]	24	Fan Slider	[0, 1]
10	Solar-S	[0, 950]	25	Fan Speed	[0, 5000]
11	Solar-W	[0, 1100]	26	Sunspace Door	[0, 2]
12	Solar-Horiz	[0, 1100]	27	Attic2 Windows	[0, 1]
13	sin(Wind-Dir)	[-1, 1]	28	Assembly Windows	[0, 1]
14	cos(Wind-Dir)	[-1, 1]	29	Basement Door	[0, 1]
15	WindSpeed	[0, 6]			

Table 5-1 Ranges used to scale the elementary features according to Equation [5-2].

5.4 Development of Potential Features

In addition to the elementary features listed in Table 5-1, several other types of features were considered to be potentially useful. These include terms to account for the operation of the fan in a more direct way.

To motivate this type of feature, assume that the type of model used is a linear model, and that Fan Speed is one of the features used by the model. In this case, the impact of fan operation on the Assembly temperature at time $t+1$ can at best be the average impact—the average effect of running the fan at a particular speed under all temperature conditions. It is more powerful to provide a feature to the model approximately proportional to the heat flow due to mass transfer through the Assembly room: $(\text{fan speed}) \cdot (T_{\text{out}} - T_{\text{Assembly}})$. This crude approximation could be improved by using the flow rate associated with a given fan speed (see Figure 4-19).

When the fan was operated, all the air exiting the Assembly room entered Attic 1, and all air exiting Attic 1 entered Attic 2. Consequently, the following three composite features were introduced:

For calculating T_{Assembly} : $(\text{fan speed}) \cdot (T_{\text{out}} - T_{\text{Assembly}})$

For calculating $T_{\text{Attic 1}}$: $(\text{fan speed}) \cdot (T_{\text{Assembly}} - T_{\text{Attic 1}})$

For calculating $T_{\text{Attic 2}}$: $(\text{fan speed}) \cdot (T_{\text{Attic 2}} - T_{\text{Attic 1}})$

Since the aperture sizes of the Sunspace were fixed throughout the experimental period³⁸, it may be useful to include a term proportional to the heat input due to buoyancy-induced flow:

³⁸ With the small exception of the few hours on September 10, 2003 when the Sunspace door was opened to warm up the building.

$$(T_{Sunspace} - T_{out}) \sqrt{\frac{T_{Sunspace} - T_{out}}{T_{Sunspace} + T_{out}}} \quad [5-5]$$

Such terms are not appropriate for use in zones that form part of a multi-level airflow, such as when air enters the louvered doors, rises up into the attics, then out the cupola. They may be appropriate for use in limited situations such as when calculating $T_{Assembly}$ when both the Fan Slider and Attic Slider are closed. For simplicity reasons, inclusion of such terms will be avoided.

Though not yet investigated, it may be useful to include the building usage term associated with meetings in the Assembly room.

5.5 Definition of Control Modes

As described in Chapter 3, the building's operation can be segmented into various control (or operational) modes. Recall that several flow regimes may exist within each mode, depending on the temperature profile and ambient conditions. In this section, the control modes at Broadmoor are defined. The benefit of these definitions will be explored in Section 5.7.3.3.

Table 5-2 lists the major modes of operation of the building. Figure 5-3 provides a graphical representation of some of the same information.

Mode Number	Fan On	Attic or Fan Slider Open	E Louvered/ W Louvered/ Conference/ Assembly Door Open	Attic 2 Windows Open	Number of Points in Mode	Relevant Zone
1	Y	Y	Y		931	1,2,3,4
2	N	Y			44847	1,2,3,4
3	N	N			14600	3,4
4	N	N	N		7936	1,2
5	N	N	Y		6664	1,2
6	N	Y	Y		44829	1,2,3,4
7	N	Y	N		18	1,2,3,4
8	N	Y		Y	26789	1,2,3,4
9	N	Y		N	18058	1,2,3,4
10	N	N		Y	8083	3,4
11	N	N		N	6517	3,4
12					60765	5,6

Table 5-2 List of control modes. In columns 3 and 4, at least one of the listed apertures must be open to obtain a "Y". To obtain a "N", no listed aperture may be open. The penultimate column contains the number of data points in that mode, while the final column lists the building zones whose thermal behavior may be linked to a particular mode. A blank entry implies that the aperture could be either open or shut. In all cases, it was required that the Sunspace door be shut.

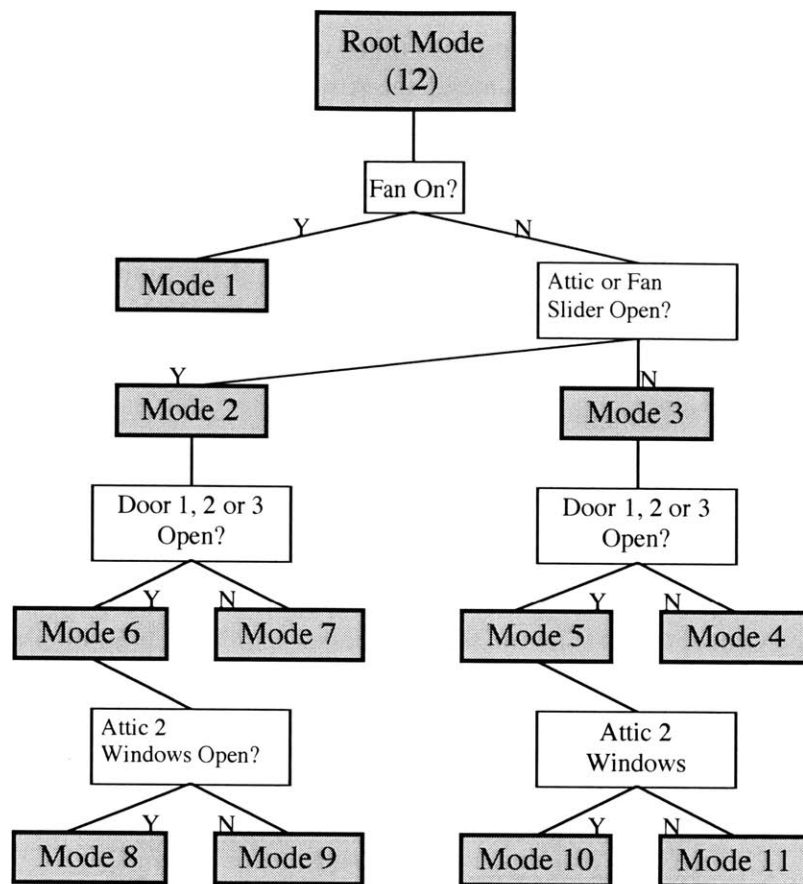


Figure 5-3 Tree representation of building modes.

Mode 1 corresponds to the set of data points when the fan was running. The fan's operation directly affected zones 1 through 4. Whether or not the Attic 2 windows were open during fan operation likely had a very small impact on the thermal behavior of the Attics (and even less on zones 1 and 2) and was ignored.

Mode 2 was not ultimately used for modeling. It represents the general case when some aperture is open to the Attics. The arrow leading from Mode 2 indicates how this set of points was further subdivided into Modes 6 and 7. Mode 6 corresponds to the scenario where air may enter at least some of the openings on the first floor, pass up through an opening to Attic 1, then up through Attic 2 and out the cupola. Mode 7 consists of the (negligibly few) points where there is an opening to the Attics, but there is no major inlet on the first floor.

The existence of Mode 7 is ignored and it is assumed that whenever an aperture to Attic 1 is open, at least one inlet on the first floor is open. In effect, Mode 6 is equivalent to Mode 2. Mode 6 is further subdivided to account for the cases where the Attic 2 windows were open (Mode 8) and shut (Mode 9).

Returning to Mode 3, it is seen that this mode corresponds to the situation where no flow can occur between the first floor and Attic 1. This mode is subdivided into Mode 4, corresponding to the building being entirely shut, and Mode 5, corresponding to the possibility of cross ventilation through the first floor.

The thermal behavior of the attic zones at those times when they are isolated from the first-floor zones may be further influenced by the state of the Attic 2 windows. Two additional modes are introduced to handle this situation: Mode 10 for Attic 2 windows open, Mode 11 for Attic 2 windows shut. Note that Modes 10 and 11 are relevant only for zones 3 and 4.

Finally, Mode 12, which contains all data points, is included for use with zones 5 and 6, which are relatively unaffected by the operational mode of the rest of the building. The Sunspace (zone 5) is completely isolated—no airflow exchange with the building occurs. The use of a single mode for this zone appears well justified. For the Basement zone, however, it may be beneficial to implement two modes: one with the Basement door open and the other with it shut. It was found that a high-fidelity Basement zone model was not required to make accurate predictions of the Assembly temperature, so this refinement was not made. Furthermore, the recordkeeping for the Basement door was the least complete of any aperture, rendering impossible an accurate subdivision of the data into two modes.

To summarize the main point of this subsection, each building zone may experience a set of control modes. The concept is as follows: the data points within each mode are used to construct a model for that particular *zone* as it evolves in that particular *mode*. Depending upon which mode the building is operating in, the appropriate models are selected and are used to make predictions for those operating conditions. Table 5-3, below, is included to summarize which modes are associated with the six building zones.

Zone ID	Zone Name	Available Modes
1	Assembly	1,4,5,8,9
2	Mass	1,4,5,8,9
3	Attic 1	1,8,9,10,11
4	Attic 2	1,8,9,10,11
5	Sunspace	12
6	Basement	12

Table 5-3 Modes available for the six building zones.

Despite the large number of controls (apertures and fan) available at Broadmoor, it was not necessary to generate 2^{11} different modes, accounting for each possible combination of the control settings. A simple assessment of the building’s operating characteristics was adequate to reduce the node number dramatically to fewer than six modes per zone.

It would be desirable to automate this procedure. A practical method for doing so would be to employ a type of regression tree, with exclusive focus on the control inputs. The proposed procedure is outlined below:

- 1) Select a zone with which to begin mode selection.
- 2) Identify all reasonable feature types, including 5 lag terms of each, for example.
- 3) The entire data set is called the “root” node. Let the root node also be the current “parent” node.
- 4) Create a control set, populated with each of the N controls available.

- 5) Using an appropriate model structure³⁹, form a model for the data associated with the current parent node, and calculate the model's error (perhaps the error on a portion of the data that was reserved for testing).
- 6) For each of the controls in the control set, divide the data into subsets (child nodes), one with that control on (or aperture open) and one with that control off (or aperture closed).
- 7) Using an appropriate model structure, form a model for each child node and identify the model's error (perhaps error on a portion of the subset of data that was reserved for testing).
- 8) Form the composite error for each pair of child nodes, weighting the node errors by their size (number of data points therein).
- 9) Select that pair that reduces the error the most, and permanently split the data accordingly. The control that split the parent is eliminated from the control set on that branch of the tree. The two child nodes so identified become new parent nodes.
- 10) For each of the parent nodes, return to step 5. Repeat the process, splitting each parent node until the splitting produces no improvement in the selected error measure.
- 11) At the end of the algorithm, all child nodes correspond to optimal control modes.

Note that the algorithm as enumerated relies strictly on Boolean “AND” decisions to select the components of modes. Since it was useful above (Table 5-2) to employ “OR” decisions when defining modes, this proposed algorithm should be modified to do so. A modified approach will be developed and presented below in Section 5.10.

5.6 Application of PHDRT

Since the operation of the Sunspace was the simplest of all the zones, it was selected as a good candidate for applying the PHDRT algorithm. Recall that a single mode encompassed all data points associated with this zone.

5.6.1 Sunspace

The data set employed for this test included Batch 1 and 2, as well as data from August 25 (00:31) to September 10 (05:47). Note that a portion of this set precedes Batch 3. The problem with these data that excluded them from general use did not affect the Sunspace⁴⁰. The data set concludes in the middle of Batch 4 to avoid the period when the Sunspace door was opened. A total of 43800 feature vector/target pairs were available.

Using the features listed in Table 5-4 as model inputs and Sunspace temperatures as targets, the PHDRT algorithm was run. Before making the first cut of the entire data set, of the 75 possible principal Hessian directions (PHD's), 50 were identified as significant (with p-values < 0.05). The plots of residuals versus each of these 50 PHD's are shown in Figure 5-4.

³⁹ To be discussed below.

⁴⁰ Uncertainty in aperture settings (Attic/Fan sliders and E/W Louvered Doors).

ID	Feature	Number of Lag Terms	ID	Feature	Number of Lag Terms
1	T_assembly	5	16	Q_bsmt	
2	T_mass	5	17	Q_first	
3	T_attic1	5	18	E Louver Door	
4	T_attic2	5	19	W Louver Door	
5	T_sunspace	5	20	Conf/Ass. Door	
6	T_basement	5	21	Office Windows S	
7	T_out	5	22	Office Windows W	
8	Solar-N	5	23	Attic Slider	
9	Solar-E	5	24	Fan Slider	
10	Solar-S	5	25	Fan Speed	
11	Solar-W	5	26	Sunspace Door	
12	Solar-Horiz	5	27	Attic2 Windows	
13	sin(Wind-Dir)	5	28	Assembly Windows	
14	cos(Wind-Dir)	5	29	Basement Door	
15	WindSpeed	5			

Table 5-4 Features used in Figure 5-4.

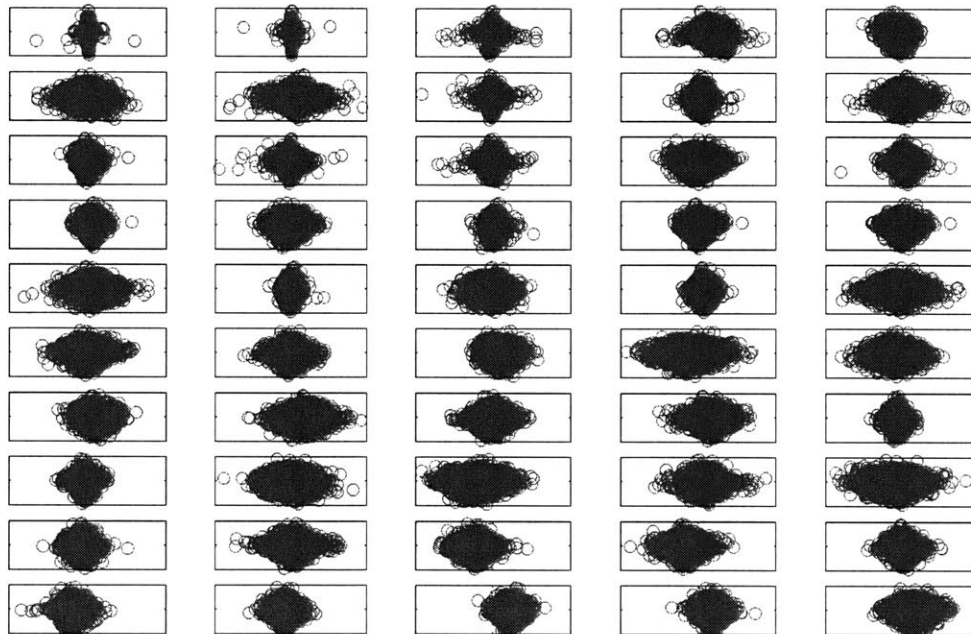


Figure 5-4 Sunspace data set. Prior to making the first cut, residuals versus PHD, for the 50 significant PHD's. (RunPHDRTTestSSpace.m)

Recall that the PHDRT algorithm determines a cut point along a PHD to minimize the composite error (of the two sides of the cut). However, in each of the 50 plots shown, there is no apparent curvature apparent and no obvious location to place a cut. This situation is in sharp contrast to that found with Demo 3 in Chapter 3, where clear curvature and/or abrupt changes in slope were

visible when residuals were plotted versus the PHD's. It was found to be impossible to locate cuts that significantly decreased the composite error. Furthermore, it is not yet understood why so many PHD were deemed statistically significant (with p-values < 0.05).

To probe whether multiple cuts in the data were required before curvature became apparent, the following calculations were performed. The requirement that the composite error after a cut be smaller than 96% of the error before the cut was relaxed to: the composite error must not be larger than the original error. With this relaxed requirement, the algorithm ran to completion. No qualitative changes in the residual plots were observed (they resembled those in Figure 5-4, with no visible curvature). Cutting was generally terminated for two reasons: 1) no significant PHD's were found, and 2) the minimum subset size was reached (100 points in this case).

The data points (feature vector/target pairs) were divided into two sets, one for training and the other for testing. The odd points were used for training and the even for testing⁴¹. The resulting PHDRT model displayed no improvement over a simple regression using all the data (with no cuts). Simple regression yielded 0.1261°C and 0.1259°C for the rms training and test error, respectively, while the PHDRT algorithm yielded 0.1210°C and 0.1265°C. The rms error will be defined in Section 5.7.1.1. The PHDRT algorithm required approximately 27 minutes.

It is important to draw conclusions from this experiment with caution. If the PHDRT algorithm were to find no directions with significant curvature, then one could conclude that the data set was generated by a linear system. That was not the case. Rather, the algorithm did find that there were directions with significant curvature, but that it was not possible to use that information to advantage by subdividing the variable space to make each portion more linear. Either the significance measure failed⁴² and the system was in fact linear, or the system was nonlinear and the subdivision process failed. Another possible alternative is that the system was nonlinear, but the nonlinearities were not strong enough to be observed above the noise. The practical implication is that a simple linear model performs just as well as the piecewise linear PHDRT model.

To investigate whether the strong correlation among the features was a factor, the same test was repeated, except only a single lag term was used for three features: T_{Sunspace} , T_{out} and $I_{\text{sun, South}}$. No other features were used. There still exists a correlation among the remaining features, such as between Sunspace and outside temperatures, but it is weaker than that between T_{Sunspace} at times t-1 and t-2, for example. Table 5-5 shows the features used. Figure 5-5 shows the qualitative correlation between the features mentioned.

As before, the residual versus PHD plots lacked clear curvature and form. The simple regression yielded 0.136 and 0.134 for the rms training and test sets, respectively, while the PHDRT algorithm yielded 0.132 and 0.132. Again, it was found that no substantial benefit was obtained from using PHDRT on this data set.

⁴¹ The rationale for this sort of training set/test set division will be discussed below.

⁴² Perhaps due to broken assumptions—the features were highly correlated. For instance, the Sunspace temperature at time t-1 was highly correlated with that at t-2. See Figure 5-5, below.

ID	Feature	Number of Lag Terms	ID	Feature	Number of Lag Terms
1	T_assembly		16	Q_bsmt	
2	T_mass		17	Q_first	
3	T_attic1		18	E Louver Door	
4	T_attic2		19	W Louver Door	
5	T_sunspace	1	20	Conf/Ass. Door	
6	T_basement		21	Office Windows S	
7	T_out	1	22	Office Windows W	
8	Solar-N		23	Attic Slider	
9	Solar-E		24	Fan Slider	
10	Solar-S	1	25	Fan Speed	
11	Solar-W		26	Sunspace Door	
12	Solar-Horiz		27	Attic2 Windows	
13	sin(Wind-Dir)		28	Assembly Windows	
14	cos(Wind-Dir)		29	Basement Door	
15	WindSpeed				

Table 5-5 Features used to investigate role of feature-feature correlation on PHDRT performance.




Figure 5-5 Above: Sunspace temperature at time t-2 versus at t-1. Below: Outside temperature at time t-1 versus Sunspace temperature at time t-1. (CorrelationPlot.m)

5.7 Linear Modeling Approach

As observed in the previous section, no benefit was found when the PHDRT algorithm was used instead of a simple linear regression. Such a result was not expected; indeed, the development in

Chapter 3 strongly indicated that a nonlinear or piece-wise linear model would be required to make accurate predictions of the thermal behavior of a naturally ventilated building such as Broadmoor. For the time being, we will put aside the question of why a linear model could possibly be appropriate and focus instead on the issue of whether a linear model can provide the performance required for control purposes.

5.7.1 Definitions for Model Assessment

Subsections 5.7.1.1 through 5.7.1.4 provide a set of definitions and preliminary information required for the following sections. Error measures, types of simulation, data-set partitioning and regression are all discussed.

5.7.1.1 Error Measures

Through training, a model “learns” (or adjusts its parameters) to map features to the desired targets in some optimal sense. With real data, the mapping is invariably imperfect due to the presence of noise or errors in the model structure. The difference between the model predictions and the targets will be called the prediction error = (estimated target – target). A compact summary of a model’s prediction errors over an entire data set with N points is the root-mean-squared (rms) error:

$$\text{rms error} = \sqrt{\frac{\sum_{i=1}^N (\text{prediction error}_i)^2}{N}} \quad [5-6]$$

An additional measure illustrating model performance is the maximum absolute error:

$$\text{max error} = \max_{i \in \{1,2,\dots,N\}} |\text{prediction error}_i| \quad [5-7]$$

5.7.1.2 Simulation Types

The models developed in the remainder of the chapter will be evaluated in several ways. These are:

- One-step-ahead simulation. With one-step-ahead simulation, the model makes the same types of predictions it made during training. Specifically, the features given to the model are relevant measurements and building information known up to time t-1. The model predicts zone temperatures at time t. Those predictions are used for error calculation only—they are not involved in the predictions of zone temperatures at time t+1, *etc.*
- Pure simulation. With pure simulation, the model is given all relevant measurements and settings that are available at time t = 0. From that point on, the only information given to the model is the set of all future exogenous inputs (outside temperature, solar and wind data, electrical loads) and control inputs (aperture and fan settings). No further measured zone temperatures are used in the feature vectors. In their place, the model’s prior predictions are used. This method of testing the model is more rigorous than the one-step-ahead method, as the possibility of compounding prediction errors arises. The pure simulation method is how the model will be used in practice.

5.7.1.3 Data Sets: Training, Test and Validation

Some mention has been made of training and test sets. These terms are defined here.

When building a model, the intent is generally to create a tool that will prove useful when making inferences or predictions about scenarios that have not yet been encountered. One hopes that the model will “generalize” well to new situations. One common way to improve a model’s generalization ability is to partition the available data set into two portions: the training set and the test set. The points (feature vector/target pairs) in the training set are used directly in adjusting a model’s parameters, *i.e.*, “training” the model. The model’s performance is tested using data it has not been exposed to: the test set. If the performance on the test set is acceptable, one may have some confidence in the model’s ability to generalize (at least for the conditions present in the test set).

A phenomenon of “over-training” or over-fitting exists. In the case of a linear model, when the number of features (and associated parameters) increases, one risks training to the noise and developing a model with poor performance on the test set. The extreme example would be a look-up table, where if the features given to a model match those seen before, the prediction error is zero, but if the model is given new features, the prediction errors may be large. The same concern arises during the training of neural nets, where too many epochs of back-propagation can lead to over-training. As the number of parameters in a model increases, the training error drops nearly monotonically, while the test error drops initially, then begins to rise when the model begins to over-fit the data. An effective method for selecting model order is to pick that order associated with the minimum in the test error curve. This approach will be illustrated in a later section (8.3.3). A similar approach was used to select σ to use for KRLS in Chapter 3.

Finally, it may be necessary to divide the full available data set into three portions. In this case, the training set is used to directly adjust model parameters, the test set is used to help select appropriate features⁴³ and the final set—the validation set—is used to check model performance on data entirely separate from those used to develop the model.

In principle, one could use errors derived from either one-step-ahead simulations or from pure simulations to train models and to select model features. For models whose features are not pure measurements, but rather combinations of measurements, the latter approach is computationally demanding⁴⁴. For that reason, all training and automated feature selection will be performed using one-step-ahead simulations.

5.7.1.4 Regression with Singular Value Decomposition

Given a set of data, the information matrix for the data is created, with each of the N rows corresponding to the M features from a particular feature vector/target pair. The set of simultaneous equations to be solved is shown below:

⁴³ Note that selecting appropriate features encompasses the problem of selecting model order.

⁴⁴ A pure ARX model may undergo pure simulation very rapidly—one merely constructs the model as a linear filter into which all exogenous inputs are fed. Such is the approach in Ljung, L. (1991). System Identification Toolbox User's Guide, The MathWorks.

$$\begin{matrix} \left[\begin{array}{c} \text{feature vector}_1 \\ \text{feature vector}_2 \\ \vdots \\ \text{feature vector}_N \end{array} \right] & \left[\begin{array}{c} \tilde{x}_1 \\ \tilde{x}_2 \\ \vdots \\ \tilde{x}_M \end{array} \right] & = & \left[\begin{array}{c} \text{target}_1 \\ \text{target}_2 \\ \vdots \\ \text{target}_N \end{array} \right] \\ (NxM) & (Mx1) & & (Nx1) \end{matrix} \quad [5-8]$$

The task at hand is to solve for the regressors, \tilde{x}_i , which form an (Mx1) vector. Each of the M elements corresponds to one particular type of feature; they are the coefficients a, b, and c defined in Equation [5-1].

It was found that the least-squares solution of Equation [5-8] using traditional matrix inverse methods⁴⁵ frequently led to regressors with excessive magnitude (such as $\sim 10^{11}$). More reasonable and numerically stable regressors were obtained using a matrix inverse technique employing singular value decomposition (SVD), adjusted slightly from (Wunsch 1996). This approach has the benefit of smoothly coping with ill-conditioned matrices. The algorithm is provided here⁴⁶.

- 1) The information matrix, E, is normalized by dividing its elements by the absolute value of the largest magnitude element. This forms E_{norm} .
- 2) E_{norm} is decomposed by standard SVD: $E_{\text{norm}} = USV'$, where U is (MxM), S is (MxN), and V is (NxN).
- 3) The k singular values (diagonal elements of S) with magnitude exceeding a tolerance (in this work, 10^{-10}) are identified.
- 4) Those k singular values are used to form a (kxk) diagonal matrix S_k .
- 5) The first k columns of U are retained to form U_k , while the first k columns of V are retained to form V_k .
- 6) The inverse of E_{norm} is computed: $E_{\text{norm}}^{-1} = V_k S_k^{-1} U_k'$.
- 7) Finally, the inverse of E is found by dividing the inverse of E_{norm} by the absolute value of the largest magnitude element of E.

Unless otherwise noted, all regression operations from this point on are performed in this manner.

5.7.2 Initial Modeling Example: Sunspace

In this section, the procedure for developing and evaluating a model for a single zone will be presented. The Sunspace, with just a single operational mode, was selected for this purpose.

For the modeling example shown in this section, the same data set used in Section 5.6.1 was used again. The set of features with a conceivable link to the Sunspace behavior was selected. The

⁴⁵ mldivide.m and inv.m MATLAB commands (The Mathworks, Natick, MA).

⁴⁶ MATLAB implementation: mySVDinverse.m

chosen feature vector is shown above in Table 5-4⁴⁷. The data set was divided into three portions: a training set, a test set and a validation set. In this example, the test set was not involved in the training process.

As only one mode encompasses all the Sunspace data, the entire training set could be used to determine the regressors. The performance of the model in one-step-ahead simulations and pure simulations⁴⁸ is summarized below in Table 5-6. Note that the one-step-ahead errors on the test and validation sets are slightly better and worse, respectively, than the error on the training set. Only two pure simulations were performed, one including the entire data set, and the other the validation set.

Data Set	Points	1-Step-Ahead RMS Error (K)	1-Step-Ahead Max Error (K)	Pure RMS Error (K)	Pure Max Error (K)
Training Set	1-7,999; 24,000-43,800	0.124	1.013	-	-
Test Set	8,000-15,999	0.122	0.939	-	-
Validation Set	16,000-23,999	0.136	1.058	1.009	4.003
Complete Set	1-43,800	0.126	1.058	0.860	4.882

Table 5-6 Sunspace model performance data.

Shown on the following four pages are:

- 1) Figure 5-6: model performance on the complete data set, showing the locations of the three distinct subsets. Below the temperature plot are two error plots; the first shows the errors during pure simulation, while the second shows one-step-ahead errors at the same times.
 Note that the predicted and measured temperatures are nearly indistinguishable on this scale, the maximum error is 4.9/45 or 11% of the full range of the Sunspace temperature, while the rms error is approximately 2% of the full range.
- 2) Figure 5-7: model performance on the validation data set. Figure layout is identical to that in Figure 5-6. Two patterns are worth pointing out in the errors. The first is most apparent in the one-step-ahead error, where a clear distinction can be seen between daytime and nighttime error variances. Associated with the larger daytime one-step-ahead errors is a clear +/-/+/- pattern in the pure simulation error. The model predictions rise faster from the nighttime temperature lows than actual, then rise more slowly to the daily peak, which is often under-predicted. Predicted temperatures drop from the peak initially more slowly, then accelerate past the measured temperatures.
- 3) A closer view of the behavior just described is provided in Figure 5-8, where the solar radiation data and outside temperature are provided for comparison.
- 4) Figure 5-9 and Figure 5-10 show histograms of the two types of errors for the complete and validation data sets, respectively. Note the bias in the pure simulation errors for the validation set.

⁴⁷ Note that the Sunspace door was not included in the feature vector since it was shut for the duration of the data set.

⁴⁸ In these pure simulations, all features other than Sunspace temperatures were measured values.




Figure 5-6 Pure simulation of Sunspace temperatures showing the complete data set and subdivisions. Note the gaps in the data dividing the Batches. Top: Simulation and measured temperatures. Middle: Pure simulation errors (predicted-measured temperatures). Bottom: One-step-ahead errors. (SunspaceDemo1.m)




Figure 5-7 Pure simulation of Sunspace temperatures showing the validation data set. Note the gaps in the data dividing the Batches. Top: Simulation and measured temperatures. Middle: Pure simulation errors (predicted-measured temperatures). Bottom: One-step-ahead errors. (SunspaceDemo1.m)




Figure 5-8 Expanded view of August 20, showing, from top: Sunspace predicted and measured temperatures, with outside temperatures; Solar radiation on three faces; pure simulation error, one-step-ahead simulation error. (PlotSunspaceZoom.m)




Figure 5-9 Left: Histogram of pure simulation error over entire data set. Right: Histogram of one-step-ahead error over entire data set. (SunspaceDemo1.m)




Figure 5-10 Left: Histogram of pure simulation error over validation data set. Right: Histogram of one-step-ahead error over validation data set. (SunspaceDemo1.m)

Close inspection of the regular pattern of the residuals (pure simulation or one-step-ahead) reveals that the model is not fully capturing the thermal behavior of the Sunspace. We will return to the issue of time-varying error variance in Section 5.15. The pattern of the daytime pure simulation errors might be related to system nonlinearities (that were not captured by PHDRT). One such nonlinearity may be imperfect alignment between the orientation of the Sunspace glazing (particularly the sloped roof) and the pyranometers. Such an explanation appears plausible due to the time-varying nature of the errors.

Several authors (Teeter and Chow 1998; Jaakkola 2001) have suggested that it may be more effective to model a time-series variable T via predictions of $(T(t)-T(t-1))$ rather than of $T(t)$. Theoretically, there should be no difference in performance for linear models, and that result was confirmed with this data set⁴⁹.

Two neural nets⁵⁰ were trained on the same training data, using $T(t)$ and $(T(t)-T(t-1))$ as targets of the two models. Performance on the one-step-ahead training error was identical to that found for the linear models. The models with $T(t)$ and $(T(t)-T(t-1))$ as targets were progressively worse than the linear models in terms of the one-step-ahead validation error (0.136, 0.141). The neural net with $T(t)$ as target slightly outperformed the linear models and dramatically outperformed the other neural net model on the pure simulation of the validation set. The validation set rms errors for the two neural nets were 0.936 and 1.452). It was found on other data sets that the models with $(T(t)-T(t-1))$ as targets showed very poor performance.

The neural nets were tested here as a check to confirm that the conclusions of the PHDRT were reasonable. As no dramatic improvement of the predictive capability was found with the neural nets, this conclusion was not refuted. Additional comparisons with nonlinear mappings will be made below.

Two other features could have been included in the feature vector: a constant term, and a buoyancy-driven flow term (see Equation [5-5]). The first may be relevant due to the normalization performed on the feature vector/target pairs. The latter is physically relevant. The performance of the model when the buoyancy-driven flow term is included is listed below in Table 5-7. Note the improvement in the pure simulation performance relative to that shown in Table 5-6.

⁴⁹ SunspaceDemo1deltaT.m

⁵⁰ Two layers of three nodes each (hyperbolic tangent activation functions), followed by a linear output layer, trained using Levenberg-Marquardt back-propagation (trainlm.m). It is acknowledged that the number of nodes in the NN is smaller than customary, given the number of features. However, the use of more nodes (10 in the input layer, then five in the inner layer—still apparently a low number) led to overtraining and worse performance on the test set in one-step-ahead and pure simulations. (SunspaceDemo3.m)

Data Set	Points	1-Step-Ahead RMS Error (K)	1-Step-Ahead Max Error (K)	Pure RMS Error (K)	Pure Max Error (K)
Training Set	1-7,999; 24,000-43,800	0.123	0.989	-	-
Test Set	8,000-15,999	0.121	0.919	-	-
Validation Set	16,000-23,999	0.134	1.057	0.884	3.485

Table 5-7 Performance of the linear model used above, with the addition of five lag terms of the buoyancy-driven flow. (SunspaceDemo1ExtraFeatures.m)

The addition of the constant term has a minimal impact on this data set. The performance data of the model including this term (and the buoyancy terms) are given below in Table 5-8. The constant term was 1.1×10^{-4} , so these results are not surprising.

Data Set	Points	1-Step-Ahead RMS Error (K)	1-Step-Ahead Max Error (K)	Pure RMS Error (K)	Pure Max Error (K)
Training Set	1-7,999; 24,000-43,800	0.123	0.990	-	-
Test Set	8,000-15,999	0.121	0.919	-	-
Validation Set	16,000-23,999	0.134	1.057	0.884	3.474

Table 5-8 Performance of the linear model used above, with the addition of a constant term and five lag terms of the buoyancy-driven flow. (SunspaceDemo1ExtraFeatures.m)

Despite the performance improvement resulting from the inclusion of the buoyancy-driven flow feature, this feature will not be used unless otherwise noted. The benefit of developing a model with simple features is compelling—it is unrealistic to require the building operator/modeler to know that the Sunspace is the only zone where such a term is useful (and physically justified), and to determine the functional form the term takes.

The constant term is included in future models.

In this section (Section 5.7.2), it has been shown that a simple linear model can be used to make predictions of the Sunspace temperature days into the future with modest errors (provided one has perfect predictions of exogenous inputs). Furthermore, the temperature predictions do not diverge significantly over time, but rather maintain close agreement with the measured temperatures.

In the next section, the use of the simple linear model is extended to the other building zones, including those experiencing multiple operational (control) modes.

5.7.3 Multi-zone Modeling and the Performance Benefit of Implementing Mode Models

The purpose of this section is to illustrate how the simple linear modeling approach described in the previous section may be applied to the building in its entirety, and how model performance can be enhanced by dividing the data set into subsets corresponding to the control modes of the building. The section is divided into three subsections, the first (5.7.3.1) is devoted to a description of the data set used in the following subsections, as well as in Section 5.9. The

second subsection (5.7.3.2) describes the composition and performance of the multi-zone model, while the third (5.7.3.3) describes the multi-mode models.

5.7.3.1 Data Description

The table below contains information about the different data sets that were used in the following sections.

Data Set	Data Points	Corresponding Dates
Training Set	[1-10,999; 20,635-37,284; 37,816-54,580]	[Aug. 6-Aug. 16; Aug. 29-Sep. 10; Sep. 10-Sep. 22]
Test Set	[1-10,999; 20,635-37,284; 37,816-54,580]	[Aug. 6-Aug. 16; Aug. 29-Sep. 10; Sep. 10-Sep. 22]
Validation Set	[11,000-20,634; 42464-48224]	[Aug. 16-Aug. 22; Sep. 14-Sep. 17]
Complete Set	[1-54580]	[Aug. 6-Sep 22]

Table 5-9 Data sets used for modeling. (See notes in text.)

Several comments are required regarding the various data sets.

- 1) The training and test sets span identical time periods (data points). However, in Section 5.9 they are interleaved, with the training set containing the odd points and the test set the even. In Sections 5.7.3.2 and 5.7.3.3, the sets were identical, so the test set was not used. Generally, training and test sets are entirely distinct. In this instance, this luxury was not available due to the operation of the building. Specifically, certain operational/control modes were used during certain periods only. If data points in those periods fell in the test set and not the training set, poor performance would result that would not be indicative of the model's ability. The compromise made here was to use training and test points throughout the entire period, with both sets experiencing all modes equally.
- 2) The validation set is taken from the middle of the period spanned by the training/test sets. Those points were not used in any way in the training process.
- 3) Also removed from this period were the points on September 10 when the Sunspace door was opened. These points were not used in training the model, and were not used in establishing one-step-ahead errors.
- 4) Two types of pure simulations were run: one of the validation set alone, and one of the complete set. Due to the (understandably) large prediction errors resulting from opening the Sunspace door on September 10, a set of points was excluded from the error analysis of the complete set. The points excluded were [37,000-39,800], a greater period than that excluded from the one-step-ahead error analysis due to the lingering impact of the large disturbance.
- 5) The validation set does not contain all the control modes. It contains four of the modes associated with the Assembly and Mass temperatures (4, 5, 8, and 9) and four associated with the Attic temperatures (8, 9, 10, and 11).
- 6) The numbering of the points in Table 5-9 is consistent with up to five lag terms being used for the various features. Changing the maximum number of lags in the model affects how many points at the start of each batch must serve as lag terms for the first complete feature vector.

The features serving as model inputs are provided in the following table.

ID	Feature	Number of Lag Terms	ID	Feature	Number of Lag Terms
1	T_assembly	5	18	E Louver Door	2**
2	T_mass	5	19	W Louver Door	2**
3	T_attic1	5	20	Conf/Ass. Door	2**
4	T_attic2	5	21	Office Windows S	2**
5	T_sunspace	5	22	Office Windows W	2**
6	T_basement	5	23	Attic Slider	2**
7	T_out	5	24	Fan Slider	2**
8	Solar-N	5	25	Fan Speed	2**
9	Solar-E	5	26	Sunspace Door	2**
10	Solar-S	5	27	Attic2 Windows	2**
11	Solar-W	5	28	Assembly Windows	2**
12	Solar-Horiz	5	29	Basement Door	2**
13	sin(Wind-Dir)	0	30	(T_ass-T_att1)*FanSpeed	5***
14	cos(Wind-Dir)	0	31	(T_att2-T_att1)*FanSpeed	5***
15	WindSpeed	5	32	(T_out-T_ass)*FanSpeed	5***
16	Q_bsmt	5*			
17	Q_first	5*			

Table 5-10 Features used in generating models. Notes: [*] Q_bsmt used only for the model of the Basement zone; Q_first used for all zones but Sunspace and Basement. [] Controls were used as features in Section 5.7.3.2 only. [***] In Section 5.7.3.2, five lag terms of features 30-32 were available for models of all zones except Sunspace and Basement. In Section 5.7.3.3, five lag terms of 30 were available for the Attic 1 zone, five lag terms of 31 for Attic 2, and five lag terms of 32 for Assembly and Mass. In all cases, these features were available only when the fan was on at time t-1 (for making a prediction at time t). In Section 5.9, the same conditions holding in Section 5.7.3.3 were in effect. Finally, the ranges with which features 30-32 were normalized were [-10000, 11500], [-13500, 3000], and [-42000, 0] respectively.**

Note that the wind directional terms were not included as possible features. A few comments are warranted on the wind and how it could effect a model such as this. Physically speaking, the role of wind in influencing the zone temperatures in a building is to modify the pressures exerted at the exterior side of the building apertures. The pressure contribution is given by (ASHRAE 1997), Chapter 25:

$$P_{wind} = C_p(\theta, geometry)\rho \frac{V_{wind}^2}{2} \quad [5-9]$$

where C_p is the wind pressure coefficient, which, in the simplest case (Swami and Chandra 1988), is dependent upon wind angle relative to the building surface and the aspect ratio of the building. The C_p for a particular aperture would also be influenced by nearby trees and local geometry of the building. Invoking the orifice equation used in Chapter 3, the square root of the pressure difference across an opening is proportional to the flow, and the product of the flow and

the zone-outside temperature difference is proportional to the heat input to the zone. The net result is that the wind-induced thermal impact of an aperture is roughly proportional⁵¹ to:

$$\text{"Heat Input}_{wind} \propto (T_{out} - T_{in})f(\theta, geometry)V_{wind} \quad [5-10]$$

where f is the square root of C_p .

Consider first the case of including the wind as a candidate for the feature vector. Using the wind as a feature in a linear model is equivalent to identifying via regression the constant of proportionality in Equation [5-10] as well as the best constant term reflecting the time-average of $(T_{out} - T_{in})f(\theta, geometry)$. If there is a prevailing wind direction, finding a constant to replace f is plausible. When one recognizes that (at this particular location) the wind blows predominantly during daytime hours, it may be plausible to assume that an average ΔT could be found (that was useful).

Now consider the case of the wind angle terms. For these to be relevant in a linear model, one would have to assume a time-averaged wind speed (over all hours of the day) as well as a time-averaged ΔT . Furthermore, the relationship between the wind angles and f is highly nonlinear. Consequently, these wind angle terms are unlikely to be beneficial in a linear model and were not used in the sections indicated. They were included for use only as features in nonlinear models that could potentially capture interactions between features such as wind speed, wind direction and relevant temperatures.

In the following subsections, numerous plots will be presented over the periods of the validation set (the two portions, one of which is in August, and the other in September) and the complete set. For reference, three plots are given in this section to indicate the behavior of some of the exogenous inputs during the same periods. The control modes are shown as well. Note that the Sunspace and Basement zones were in Mode 12 for the entire time.

The complete period is shown in Figure 5-11, followed by the first validation set period in Figure 5-12 and the second in Figure 5-13.⁵²

⁵¹ Only approximately so since a pressure *difference* term appears within the square root.

⁵² (MainNoOptimNoModes.m)

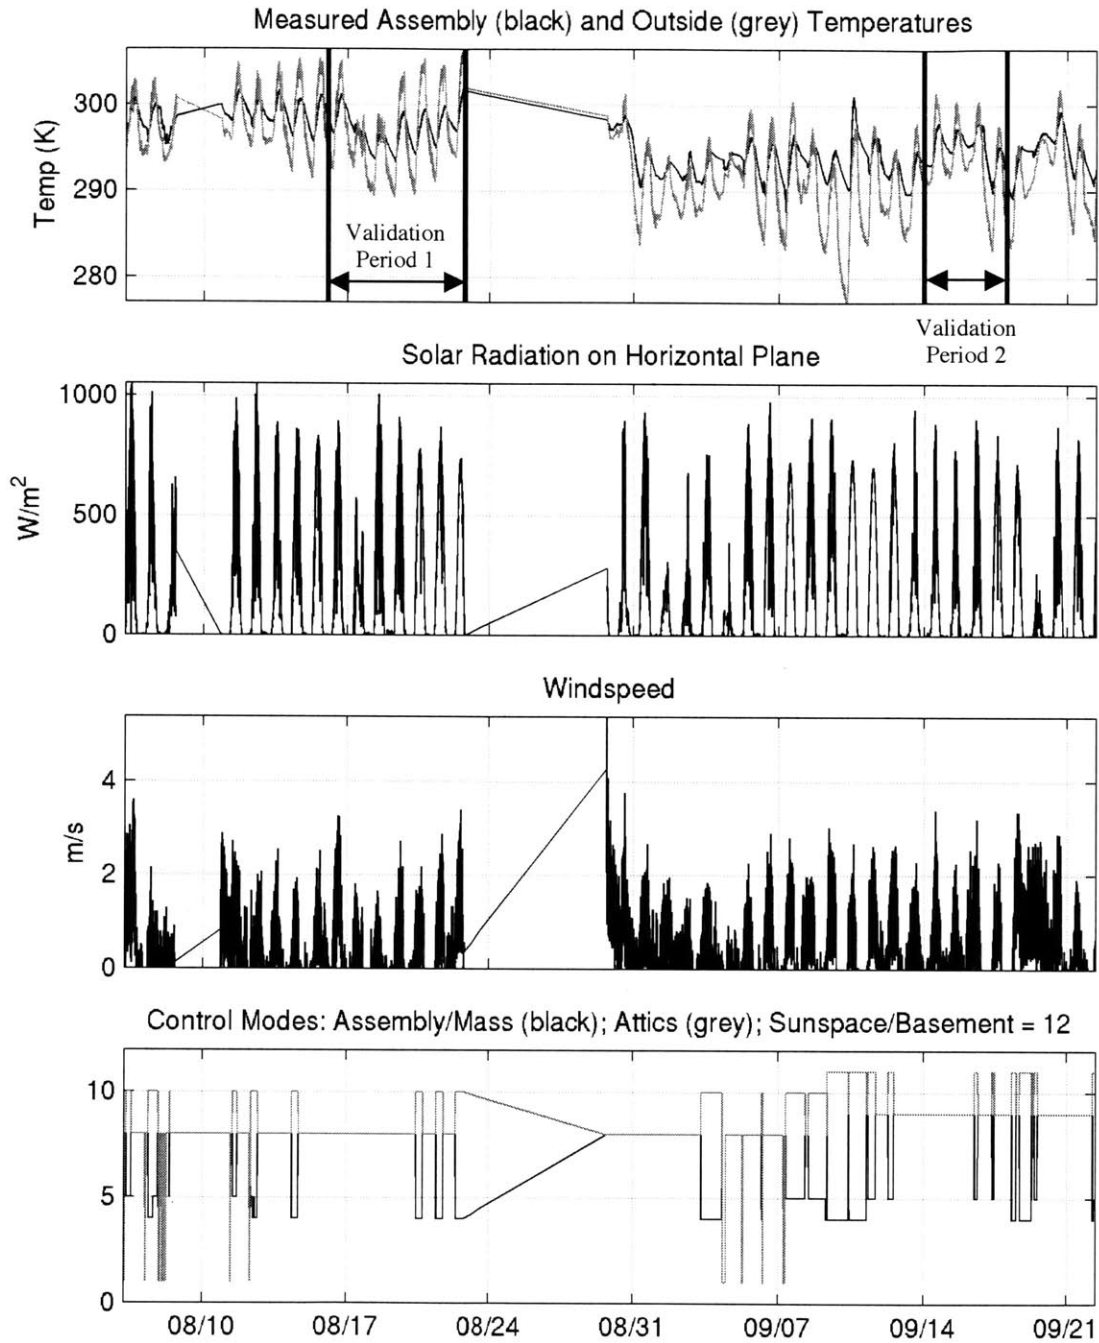


Figure 5-11 Selected exogenous inputs and control modes over the complete set. **First:** Assembly and Outside temperatures. Validation periods are indicated. **Second:** Solar radiation on the horizontal plane. **Third:** Wind speed. **Fourth:** Control modes.




Figure 5-12 Selected exogenous inputs and control modes over the first validation period. First: Assembly and Outside temperatures. Second: Solar radiation on the horizontal plane. Third: Wind speed. Fourth: Control modes. (Note that only two modes per zone were represented. Two additional modes were used in the second validation period.)

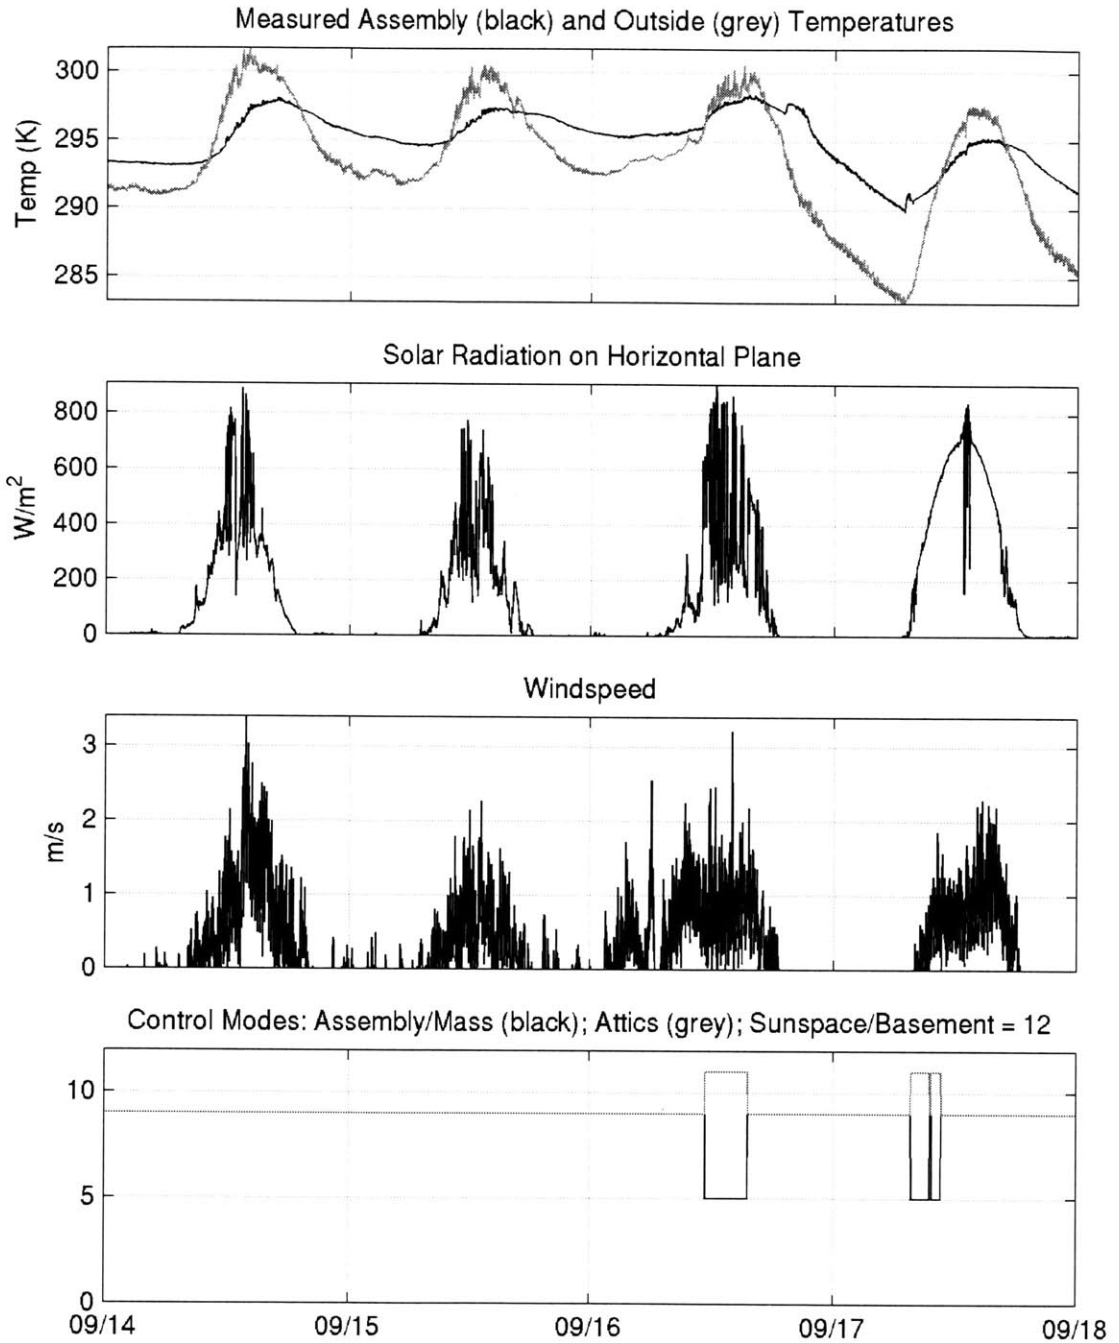


Figure 5-13 Selected exogenous inputs and control modes over the second validation set. First: Assembly and Outside temperatures. Second: Solar radiation on the horizontal plane. Third: Wind speed. Fourth: Control modes. (Note that only two modes per zone were represented, though they were different from those in the first validation period.)

5.7.3.2 Multi-zone Model

In this section, six zonal models were constructed using the features indicated in Table 5-10. A graphical representation of the model parameters selected via regression is given in the figure below. These parameters are the a, b and c coefficients in Equation [5-1].




Figure 5-14 Representation of model parameters. To reveal the minor components, all parameters were converted as follows: $\text{Plotted parameter}_i = \text{sign}(\text{parameter}_i) |\text{parameter}_i|^{1/2}$. The axis numbering is identical to the numbering in Table 5-10 (e.g., in the plot of basement parameters, the five parameters corresponding to the sixth feature, T_{Basement} , are located at 6.0, 6.2, ... 6.8). The parameter indicated by the arrow in each plot corresponds to the coefficient of that zone temperature at time $t - 1$. The constant term is plotted at 0.4 on the abscissa. (PlotFeatureParameters.m)

Note the dominance of the each zone's temperature parameters in the plot of that zone's parameters: Assembly temperature parameters dominate in the first plot, Mass temperature parameters dominate in the second, *etc.* Recall that the emphasis of the smaller parameters has been exaggerated by the plotting technique.

From the parameters shown in the first plot, it is clear that the Assembly and Mass temperatures are strongly coupled, and from the second plot, it is clear that the direction of influence is not balanced. Attic 1 can be seen to be coupled to the Mass. The Mass may serve as a smoothed surrogate for the Assembly temperature, to which it is directly coupled via airflow. Curiously, the Sunspace appears to be coupled to the Mass and Attic 1. The Mass connection may be related to conduction heat transfer through the building/Sunspace wall. The Attic 1 connection may merely indicate a strong correlation (but no causal relationship) between the two spaces.

The performance of the models composed of the above parameters is illustrated in the following six figures.




Figure 5-15 Multi-zone Model: measured (black) and predicted (grey) temperatures over the complete data set. (Features indicated by the arrows are discussed below in the text.) (MainNoOptimNoModes.m)




Figure 5-16 Multi-zone Model: prediction errors over the complete data set. Pure simulation errors are shown in black, while the one-step-ahead errors at the same time are overlaid in grey. (MainNoOptimNoModes.m)




Figure 5-17 Multi-zone Model: measured (black) and predicted (grey) temperatures over the first part of the validation data set. (Features indicated by the arrows are discussed below in the text.) (MainNoOptimNoModes.m)




Figure 5-18 Multi-zone Model: prediction errors over the first part of the validation data set. Pure simulation errors are shown in black, while the one-step-ahead errors at the same time are overlaid in grey. (Features indicated by the asterisks are discussed below in the text.) (MainNoOptimNoModes.m)

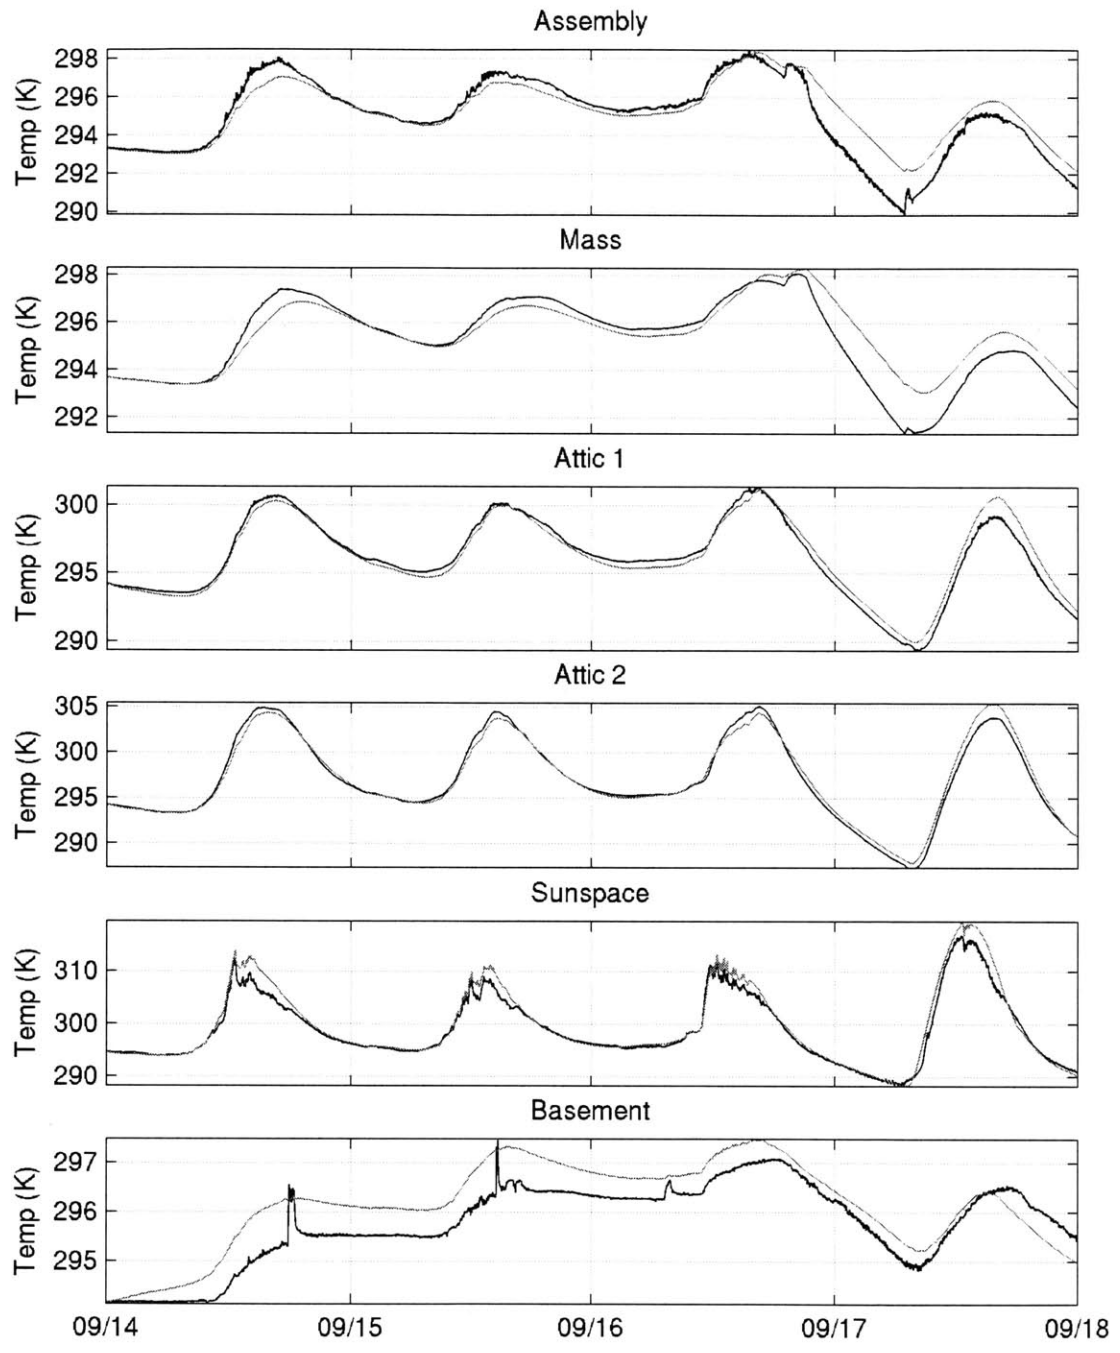


Figure 5-19 Multi-zone Model: measured (black) and predicted (grey) temperatures over the second part of the validation data set. (MainNoOptimNoModes.m)




Figure 5-20 Multi-zone Model: prediction errors over the second part of the validation data set. Pure simulation errors are shown in black, while the one-step-ahead errors at the same time are overlaid in grey. (MainNoOptimNoModes.m)

Zone	Mode	Training Set			Validation Set			Complete Set	
		Number of Points	rms error	max error	Number of Points	rms error	max error	rms error	max error
1	12	38653	0.048	0.520	15396	0.044	0.356	0.047	0.520
	Pure Simulation Results -->						0.717	2.364	0.637
2	12	38653	0.007	0.198	15396	0.007	0.065	0.007	0.198
	Pure Simulation Results -->						0.666	2.113	0.583
3	12	38653	0.031	0.482	15396	0.030	0.340	0.031	0.482
	Pure Simulation Results -->						0.454	1.614	0.417
4	12	38653	0.035	0.363	15396	0.034	0.324	0.035	0.363
	Pure Simulation Results -->						0.495	1.668	0.520
5	12	38653	0.124	1.257	15396	0.123	1.218	0.124	1.257
	Pure Simulation Results -->						1.224	4.844	1.176
6	12	38653	0.035	1.504	15396	0.027	0.689	0.034	1.504
	Pure Simulation Results -->						0.484	1.152	0.535

Table 5-11 Multi-zone Model: Performance measures. The performance of the model for each zone in three data sets is provided in the three “column” sections of the table: training, validation and complete sets. (A column for the test set was omitted since it was identical to that of the training set.) From left to right are, for each set: the number of points in the set, and the one-step-ahead rms error and maximum magnitude errors. In the second row of each zone are shown the pure simulation performance measures of the model in the validation and the complete sets. (Errors given in °C.) (ExportModelInfo.m)




Figure 5-21 Multi-zone Model, Assembly: Left: Histogram of pure simulation errors over the complete set. Right: Histogram of one-step-ahead errors over the complete set.




Figure 5-22 Multi-zone Model, Assembly: Left: Histogram of pure simulation errors over the first part of the validation set. Right: Histogram of one-step-ahead errors over the first part of the validation set. (MainNoOptimNoModes.m)




Figure 5-23 Multi-zone Model, Assembly: Left: Histogram of pure simulation errors over the second part of the validation set. Right: Histogram of one-step-ahead errors over the second part of the validation set. (MainNoOptimNoModes.m)

Much information has been presented over the preceding pages. Several points are noteworthy:

- 1) Qualitatively, predictive performance of the models is generally good over the periods examined. No extreme deviations occurred other than when the Sunspace door was opened.
- 2) The model has successfully captured the “average” behavior of the building, but not the particular behavior. For example, consider the Assembly temperature plots in Figure 5-15 and Figure 5-17 where they are marked with arrows. These correspond to times when the building was completely shut (Mode 4; see mode trajectories in Figure 5-11 and Figure 5-12). At these times, the “average” building behavior yields predictions that are uniformly higher than measurements at these times. Merely having the control settings available as features does not necessarily lead to a model that makes effective use of them. This model would not be useful for making predictions of how the building would respond to different control inputs. More evidence supporting this claim is that the parameters on all the control features are zero in Figure 5-14.

As another example, see the asterisks (*) in Figure 5-18, marking periods for the Assembly and Attic 1 when the building is completely shut. In the case of Attic 1, on average, air from the building cools the attic. In this case it cannot since the Attic and Fan sliders are shut, so “average-case” temperature predictions are lower than measured.

- 3) Some asymmetry in the distribution of pure simulation errors is apparent in Figure 5-21 through Figure 5-23.

The numerical measures of the model performance over the various data sets are listed above in Table 5-11. This table will be useful as a basis of comparison with the analogous Table 5-12 in the next section. For example, the pure simulation rms error for the Assembly temperature predictions was found to be 0.72 and 0.64°C, for the validation and complete sets, respectively. The comparable one-step-ahead errors were 0.04 and 0.05°C.

5.7.3.3 Implementation of Multi-mode Models

In an attempt to develop a model more capable of predicting building behavior, the data were divided up into modes as described in Section 5.5. As above, the training and test sets shared the same time periods, so the test set was not relevant in this case. Performance of the model on the validation set and the complete set is provided.

As noted in Table 5-10 above, no control features were used in the mode-based models in Section 5.7.3.3. This is indeed a shortcoming of the model framework, but there are several important reasons for constructing the model in this way. Recall from Section 5.5 that mode 8 was defined in the following way:

- At least one of the louvered doors was open AND
- At least one of the sliders (Fan/Attic) was open AND
- Attic 2 windows were open

Since the control settings are not available to the mode 8 model, no adjustment in temperature predictions is possible for the case where both attic sliders are open versus just one, or for the case where additional apertures, such as the Conference Room door, are opened.

Including these options was considered, but dismissed for several reasons. To include the Conference Room door setting as a feature, for example, it would be necessary to form a composite term of (door setting)*($T_{\text{Assembly}} - T_{\text{out}}$)*(airflow due to buoyancy/wind). This was undesirable from a model complexity standpoint. More important, though, is that for such a term to be useful, it would be necessary to know the flow rate and direction of the air passing through the aperture (recall the discussion of flow regimes in Chapter 3). Obtaining that information for any but the simplest building configurations and boundary conditions is not practical.

Consequently, such terms were not included.

If apertures such as the Conference Room Door are indeed important in governing the thermal behavior of the building, then perhaps a mode should be defined to reflect that. The same could be said for the Fan and Attic Sliders. If they are to be used independently, and the impact of opening one versus the other versus both is measurable, then each configuration should be described by a dedicated mode. In this work, a compromise was made to keep the number of modes at a manageable level. Certainly, better accuracy could have been achieved, but note that the number of data points required to identify model parameters increases as more modes are introduced.

Shown below in Figure 5-24 and Figure 5-25 are the (adjusted) model parameters for each mode of each zone. Note the distinct difference between the parameter envelope for Mode 1 versus the other modes in Figure 5-24. This distinction is somewhat artificial, since some features⁵³ were deemed irrelevant during fan operation due to the dominant effect of the fan on the thermal behavior of the building. For the other modes, it is difficult to detect a regular pattern distinguishing the parameter weights from one mode to another.

⁵³ For the Assembly and Mass zones, when the fan was on, features [3-5,13-15,30-31] were excluded; for the Attic 1 zone, [4-6,13-15,17,31,32] were excluded; and for the Attic 2 zone, [1,2,5,6,13-15,17,30,32] were excluded.

The performance measures for the multi-mode model are shown in detail in Table 5-12. Comparison with Table 5-11 reveals a dramatic improvement in performance. The pure simulation rms error on the validation set for the Assembly temperature drops from 0.72 to 0.34°C, while the rms error on the complete set drops from 0.64 to 0.40°C. Bear in mind that the simulations were run much longer than required in practice; the rms error for 24-hour predictions would be even smaller than reported. More important than the global error measures is that the model predictions are adjusted depending on the mode of operation of the building; there is a better likelihood that trial control strategies may be investigated with some confidence.

The improvement in performance is visually apparent in the set of nine figures (Figure 5-26 through Figure 5-34) analogous to those seen in the last section: the pure simulations and associated errors for the complete set and validation set (2 parts), followed by histograms of the pure simulation and one-step-ahead errors for the two sets⁵⁴. This discussion will be continued after the figures are presented.

⁵⁴ MainNoOptim.m

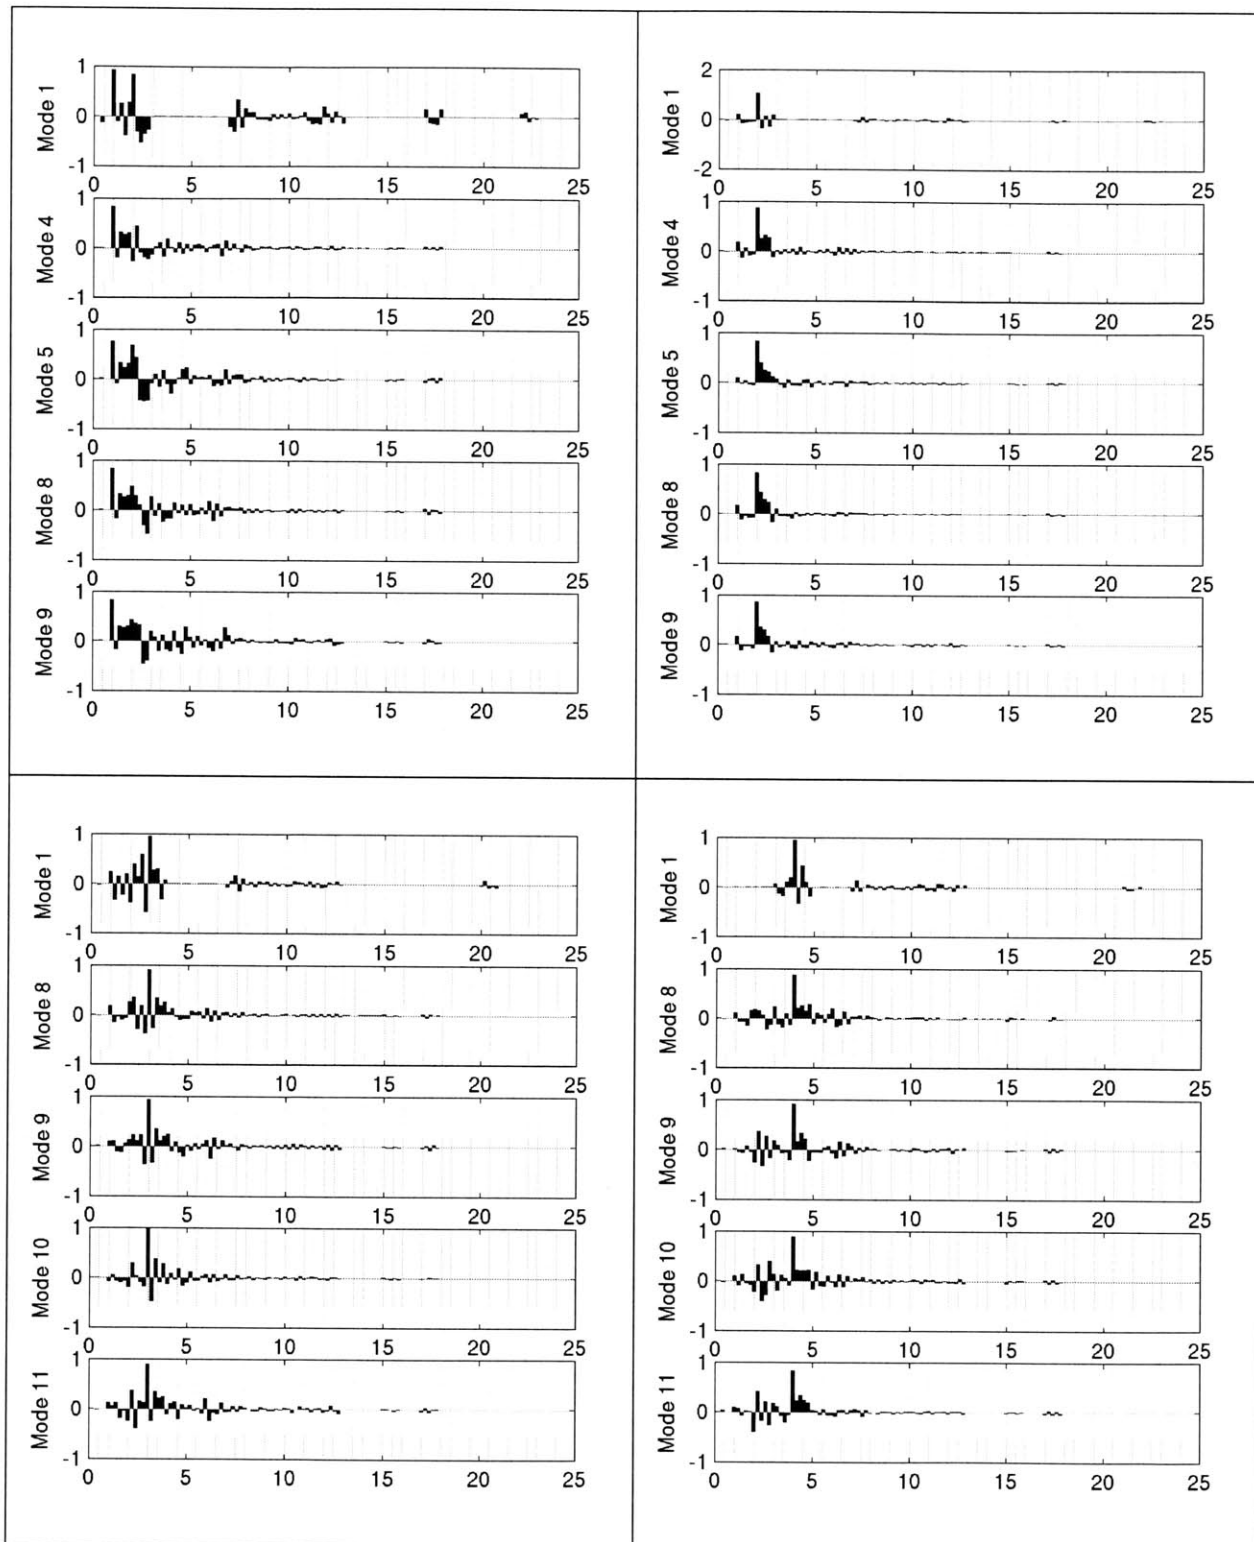


Figure 5-24 Model parameters (adjusted as before) for each mode. From upper left, clockwise, Assembly, Mass, Attic 2, Attic 1. Note that features 30-32 were renumbered to 20-22 to make the plots more readable. (PlotFeatureParameters2.m)




Figure 5-25 Model parameters (adjusted as before) for each mode. Left: Sunspace. Right: Basement. (PlotFeatureParameters2.m)

Zone	Mode	Training Set			Validation Set			Complete Set	
		Number of Points	rms error	max error	Number of Points	rms error	max error	rms error	max error
1	1	912	0.094	0.577					
	4	6241	0.037	0.455	1566	0.052	0.250		
	5	5042	0.059	0.344	419	0.056	0.214		
	8	18720	0.044	0.411	8069	0.037	0.290		
	9	7738	0.046	0.488	5342	0.050	0.353		
	all	38653	0.047	0.577	15396	0.044	0.353	0.047	0.577
Pure Simulation Results -->					0.338	1.490	0.404	1.621	
2	1	912	0.014	0.113					
	4	6241	0.007	0.191	1566	0.007	0.026		
	5	5042	0.007	0.055	419	0.006	0.020		
	8	18720	0.006	0.086	8069	0.007	0.052		
	9	7738	0.007	0.045	5342	0.008	0.064		
	all	38653	0.007	0.191	15396	0.007	0.064	0.007	0.191
Pure Simulation Results -->					0.338	1.501	0.395	1.501	
3	1	912	0.035	0.169					
	8	18720	0.029	0.285	8069	0.029	0.222		
	9	7738	0.036	0.410	5342	0.032	0.294		
	10	6517	0.023	0.196	1566	0.035	0.338		
	11	4766	0.037	0.295	419	0.024	0.073		
	all	38653	0.031	0.410	15396	0.030	0.338	0.031	0.410
Pure Simulation Results -->					0.413	1.412	0.404	1.458	
4	1	912	0.016	0.076					
	8	18720	0.037	0.353	8069	0.034	0.316		
	9	7738	0.020	0.137	5342	0.022	0.162		
	10	6517	0.043	0.218	1566	0.058	0.276		
	11	4766	0.027	0.130	419	0.035	0.131		
	all	38653	0.034	0.353	15396	0.034	0.316	0.034	0.353
Pure Simulation Results -->					0.375	1.248	0.372	1.334	
5	12	38653	0.124	1.261	15396	0.123	1.221		
	all	38653	0.124	1.261	15396	0.123	1.221	0.124	1.261
	Pure Simulation Results -->					1.206	4.779	1.175	6.097
6	12	38653	0.035	1.505	15396	0.027	0.690		
	all	38653	0.035	1.505	15396	0.027	0.690	0.034	1.505
	Pure Simulation Results -->					0.332	1.210	0.654	3.469

Table 5-12 Multi-mode Model: Performance measures. The performance of the model for each zone in three data sets is provided in the three “column” sections of the table: training, validation and complete sets (test set was irrelevant). From left to right are, for each set: the number of points in the mode, and the one-step-ahead rms error and maximum magnitude errors. For each zone, results for each mode are given in a dedicated row. The composite results for a given zone are shown in the row labeled “all”, while the final row of the zone contains pure simulation results for the validation and complete sets. (Errors given in °C.) (ExportModelInfo.m)

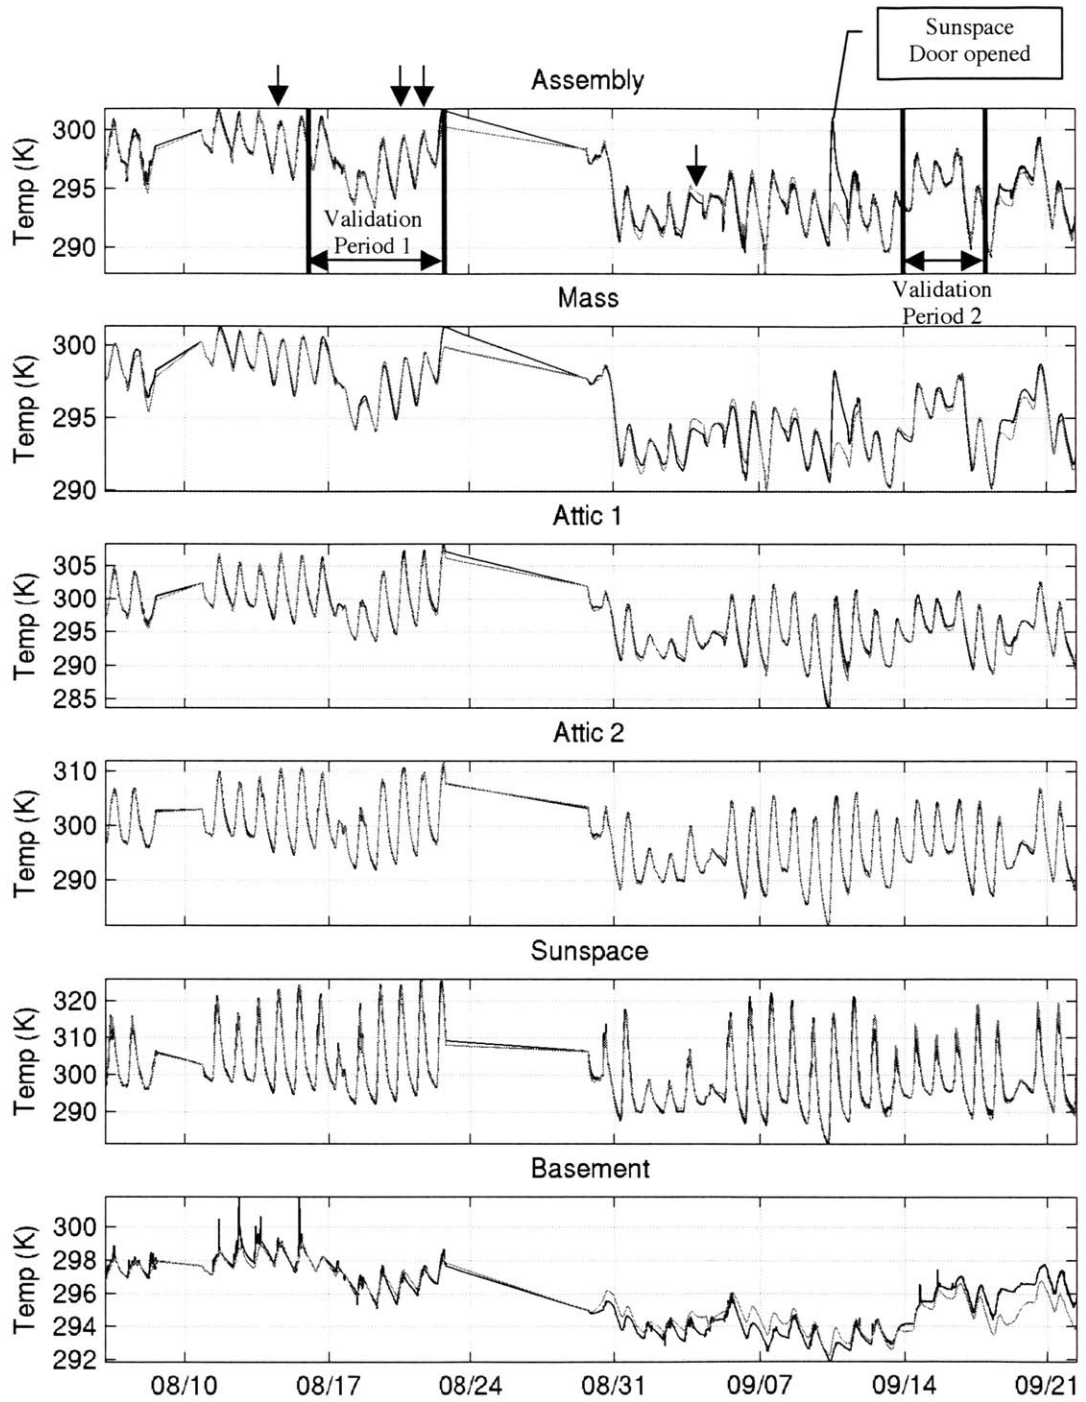


Figure 5-26 Multi-mode Model: measured (black) and predicted (grey) temperatures over the complete data set. (Features indicated by the arrows are discussed below in the text.)




Figure 5-27 Multi-mode Model: prediction errors over the complete data set. Pure simulation errors are shown in black, while the one-step-ahead errors at the same time are overlaid in grey.

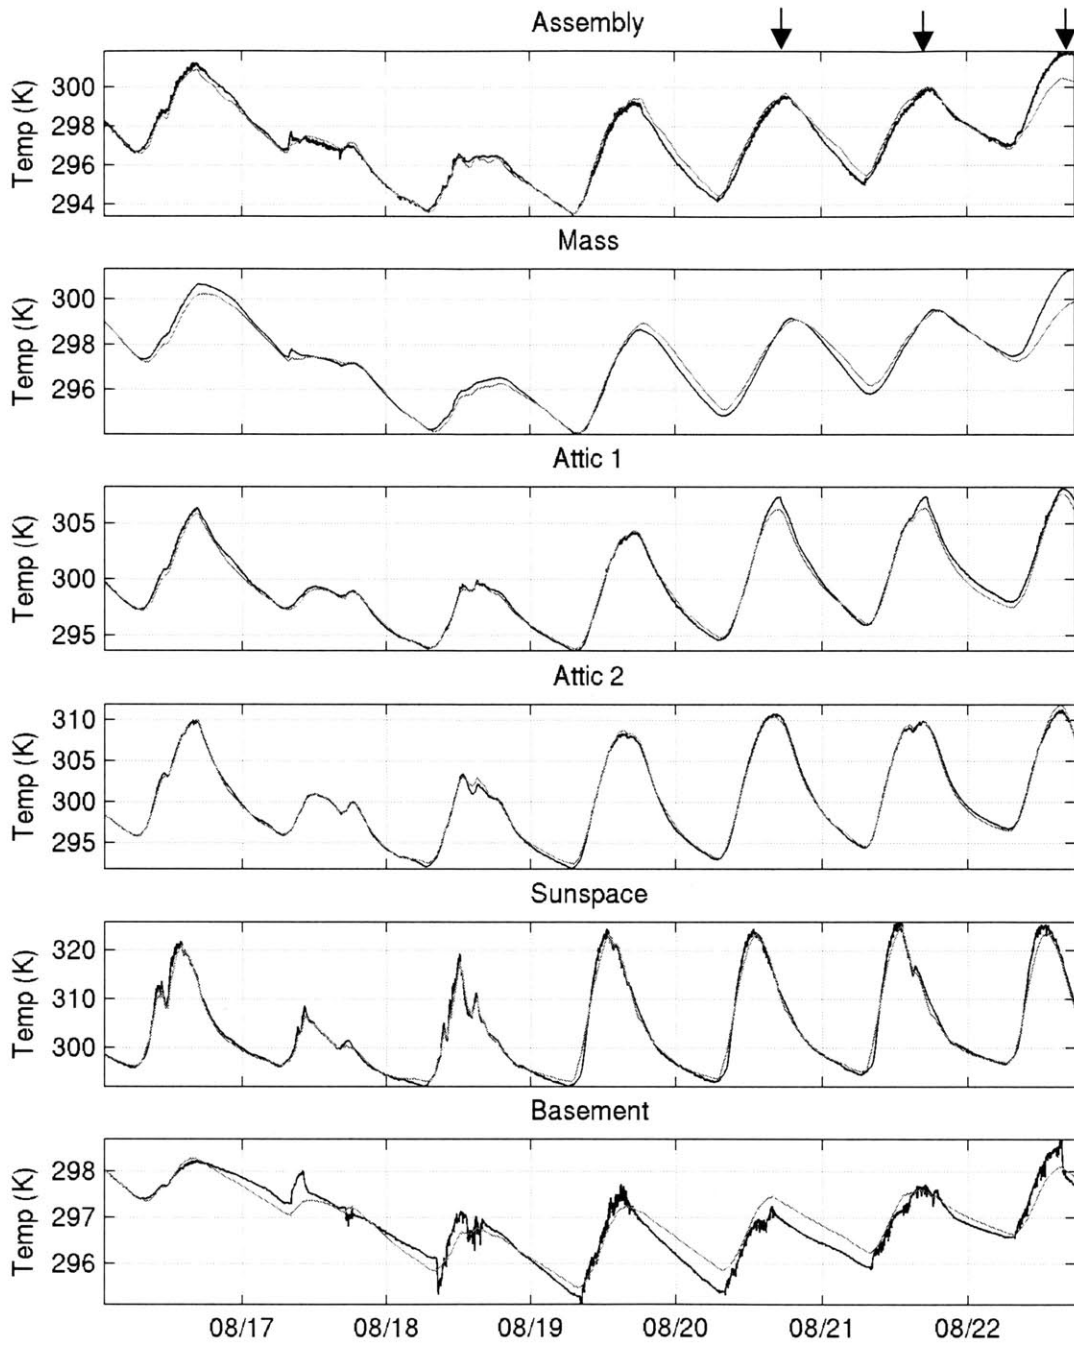


Figure 5-28 Multi-mode Model: measured (black) and predicted (grey) temperatures over the first part of the validation data set. (Features indicated by the arrows are discussed below in the text.)




Figure 5-29 Multi-mode Model: prediction errors over the first part of the validation data set. Pure simulation errors are shown in black, while the one-step-ahead errors at the same time are overlaid in grey. (Features indicated by the asterisks are discussed below in the text.)

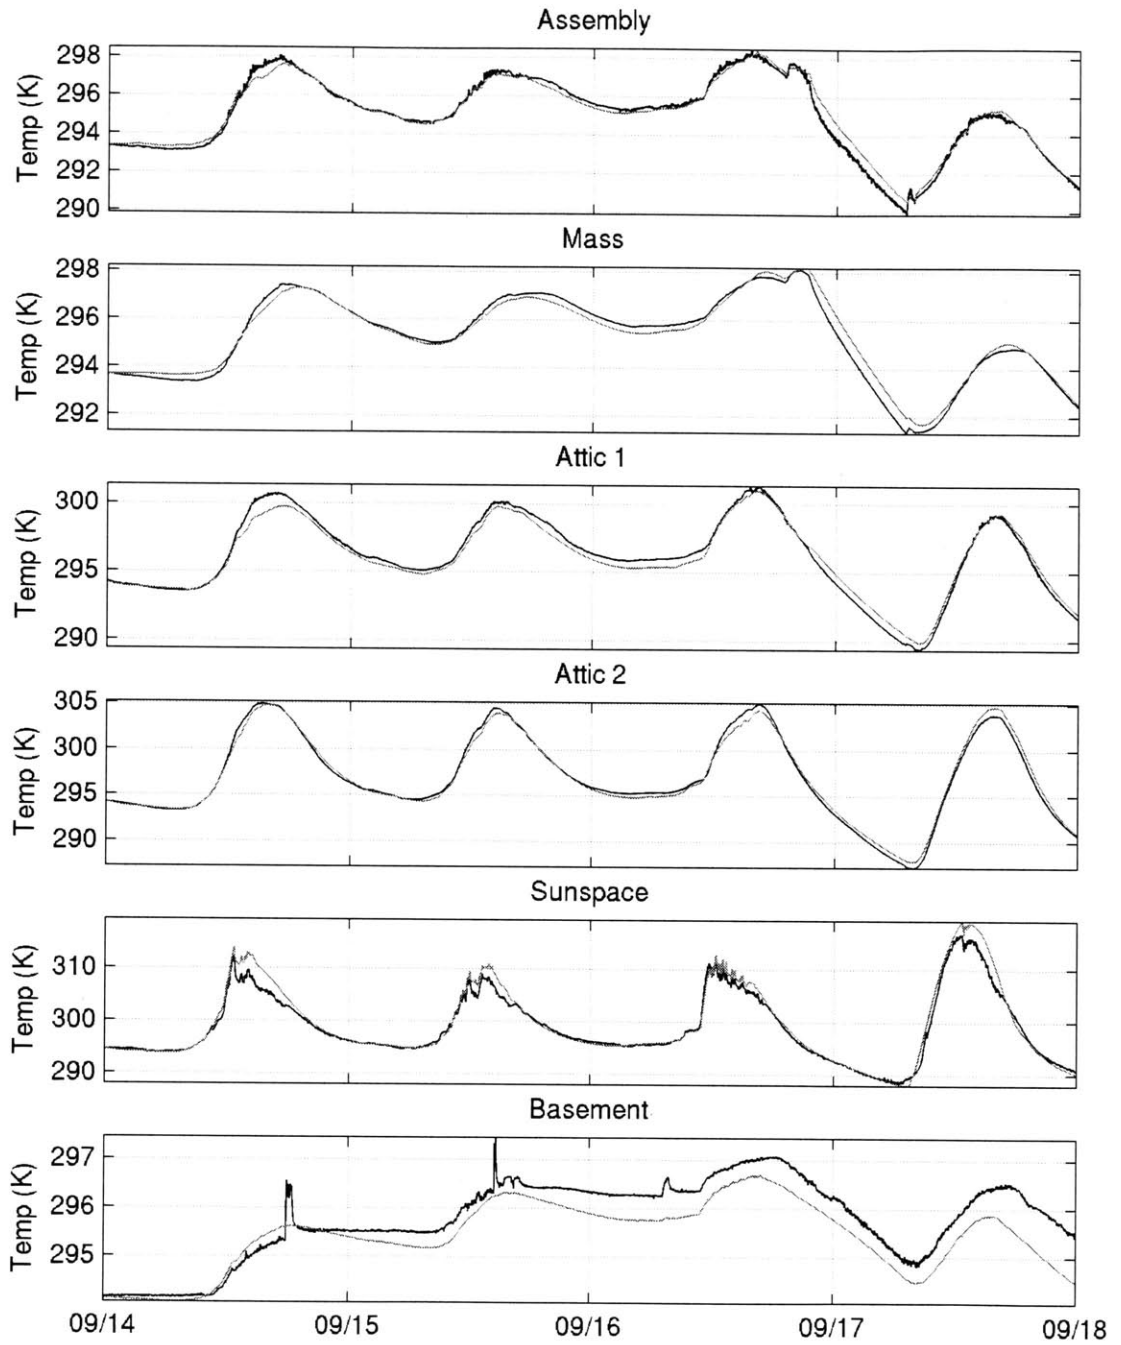


Figure 5-30 Multi-mode Model: measured (black) and predicted (grey) temperatures over the second part of the validation data set.




Figure 5-31 Multi-mode Model: prediction errors over the second part of the validation data set. Pure simulation errors are shown in black, while the one-step-ahead errors at the same time are overlaid in grey.




Figure 5-32 Multi-mode Model, Assembly: Left: Histogram of pure simulation errors over the complete set. Right: Histogram of one-step-ahead errors over the complete set.




Figure 5-33 Multi-mode Model, Assembly: Left: Histogram of pure simulation errors over the first part of the validation set. Right: Histogram of one-step-ahead errors over first part of the validation set.




Figure 5-34 Multi-mode Model, Assembly: Left: Histogram of pure simulation errors over the second part of the validation set. Right: Histogram of one-step-ahead errors over the second part of the validation set.

Recall the examples cited in the previous section (indicated by arrows on Figure 5-15 and Figure 5-17 and by asterisks on Figure 5-16 and Figure 5-18). These instances were exemplars of moments when a single model of a zone failed to adequately predict zone temperatures when the control mode was changed. Comparisons with the analogous figures (Figure 5-26 and Figure 5-28 (arrows) and Figure 5-27 and Figure 5-29 (asterisks)) show that in all cases, the new model incorporating mode information explicitly was able to predict the Assembly and Attic temperatures more accurately.

Note also the improvements in the second part of the validation set. Specifically, the single model under-predicted Assembly temperature on September 14 and did not capture the temperature drop over the night of September 16-17 (Figure 5-19). Noticeable improvement is observed in the multi-mode model predictions (Figure 5-30). These changes can be observed on the corresponding error plots, as well.

Other examples could be found, *e.g.*, with the fan operation on August 11 and 12 (compare error plots in Figure 5-16 and Figure 5-27). An apparent counter-example for the fan operation appears on August 8. However, the large negative error before the gap in the data on August 8 is associated with a period when just the Fan Slider was open, rather than both Fan and Attic Sliders. The building is in Mode 8 because *at least* one of the apertures was open. Since they are generally opened and closed together, the “average” Mode 8 model experiences greater airflow and cooling than the situation at this point. Hence the error.

That said, the model for Mode 1 could be improved with more data, especially with various fan speeds. The fan was used predominantly at full speed, so at the times when the speed was not 100%, the performance is not as accurate.

The (extremely) careful reader may have noted that the errors in the Sunspace temperature predictions were very different in September (validation set, part 2, Figure 5-31) from those in August (Figure 5-29). These errors can be linked to the operation of the fan in the Sunspace to charge the rockbed below the basement. This thermostatically-controlled fan was turned on on September 12. Since the model did not incorporate this information, the errors occur.

A comment is also warranted on the performance of the model on August 22 (see Figure 5-28 and Figure 5-29). There are several possible explanations for the large error observed. The first is model inadequacy, though such an explanation is not convincing. To understand why, consult Figure 5-12. August 22 is the third of three days with identical operational mode settings and roughly comparable weather conditions (though outside temperatures on the 22nd are 1-2°C higher at peak). The model predictions are good until the mode changes from 8 to 4. Note that on August 20 and 21, this change did not affect model prediction accuracy. The author was not on site on these three days. A different staff member adjusted the apertures and recorded them on the 22nd than on the previous days. It may be that a mistake was made. The observed temperature trajectory is more consistent with the building being open (such as in Mode 8). An alternative explanation is that the settings were made and recorded accurately, but that campers and camp staff opened other openings (such as Assembly and Conference Room doors). A clue that the building may be in a different operational mode may be the shape of the bimodal error distribution seen in Figure 5-33. There appear to be two distinct regions of operation. (The same could be said about the distribution in Figure 5-34, where the errors occur on the night of September 16-17. (The “bump” in temperature in the early evening hours was due to high electrical loads associated with lighting a public meeting. This rise was adequately predicted by the model, though the temperature trajectory after the meeting showed the curious deviation.)

Needless to say, aperture sensors would have been very helpful.

Now that the importance of dividing the data set into operational modes has been established, a few comments regarding the modeling assumptions are warranted. Since the Assembly zone is influenced by all other zones in some way, the choice was made to implement a multi-zone model, with all zone temperatures serving as candidate features. The Basement, Sunspace, Mass, and Attic 1 zone are all adjacent to the Assembly zone, and the temperatures of both Attics govern the flow of air up and out of the Assembly zone. It was desirable to establish the practical influence of the non-Assembly zones in predicting Assembly temperatures. Inspection of Figure 5-24 (upper left plots) reveals that the Assembly temperature model is apparently strongly influenced by the Mass temperature and moderately or weakly influenced by the other zone temperatures.

Two sets of calculations were performed to probe this matter further. In the first set of calculations, models were formed for all modes relevant for the Assembly zone as above, but with one exception: all other zone temperatures were removed as possible features. In the

second set of calculations, only the Assembly and Mass zones were preserved as features for the Assembly and Mass zones. In essence, these are the models that would have been created had only Assembly, or Assembly and Mass temperatures been available. Shown below in Table 5-13 are the features for the two sets of calculations.

ID	Feature	Number of Lag Terms	ID	Feature	Number of Lag Terms
1	T_assembly	5	18	E Louver Door	0
2	T_mass	0	19	W Louver Door	0
3	T_attic1	0	20	Conf/Ass. Door	0
4	T_attic2	0	21	Office Windows S	0
5	T_sunspace	0	22	Office Windows W	0
6	T_basement	0	23	Attic Slider	0
7	T_out	5	24	Fan Slider	0
8	Solar-N	5	25	Fan Speed	0
9	Solar-E	5	26	Sunspace Door	0
10	Solar-S	5	27	Attic2 Windows	0
11	Solar-W	5	28	Assembly Windows	0
12	Solar-Horiz	5	29	Basement Door	0
13	sin(Wind-Dir)	0	30	(T_ass-T_att1)*FanSpeed	0
14	cos(Wind-Dir)	0	31	(T_att2-T_att1)*FanSpeed	0
15	WindSpeed	5	32	(T_out-T_ass)*FanSpeed	5
16	Q_bsmt	0			
17	Q_first	5			

ID	Feature	Number of Lag Terms	ID	Feature	Number of Lag Terms
1	T_assembly	5	18	E Louver Door	0
2	T_mass	5	19	W Louver Door	0
3	T_attic1	0	20	Conf/Ass. Door	0
4	T_attic2	0	21	Office Windows S	0
5	T_sunspace	0	22	Office Windows W	0
6	T_basement	0	23	Attic Slider	0
7	T_out	5	24	Fan Slider	0
8	Solar-N	5	25	Fan Speed	0
9	Solar-E	5	26	Sunspace Door	0
10	Solar-S	5	27	Attic2 Windows	0
11	Solar-W	5	28	Assembly Windows	0
12	Solar-Horiz	5	29	Basement Door	0
13	sin(Wind-Dir)	0	30	(T_ass-T_att1)*FanSpeed	0
14	cos(Wind-Dir)	0	31	(T_att2-T_att1)*FanSpeed	0
15	WindSpeed	5	32	(T_out-T_ass)*FanSpeed	5
16	Q_bsmt	0			
17	Q_first	5			

Table 5-13 Features used to investigate influence of other zones on the prediction of Assembly temperatures. Top: First case, with no other zones used as features. Bottom: Second case, with the Mass used as a feature. Feature 32 was available only during Mode 1 (when the fan was on at time t-1 and a prediction was being made for temperatures at time t).

The performance of the models in pure simulations of the validations set (both portions) is presented in Table 5-14 below, where it is contrasted with the performance obtained using the full complement of features.

Temperature Features Used	Assembly Zone		Mass Zone	
	rms error	max error	rms error	max error
Assembly Temperature Only	0.62	2.24	-	-
Assembly and Mass Temperatures Only	0.37	1.18	0.37	1.10
All Zone Temperatures (original case)	0.34	1.49	0.34	1.50

Table 5-14 The influence of other zones on the performance of the Assembly temperature model. Pure simulations were performed for the entire validation set, and the errors in °C are reported here.

One may conclude from these results that a minimal penalty is paid if only the Assembly and Mass temperatures are available relative to the case where all are available. The benefit of using the Mass temperature is striking, as can be seen by comparing the first two rows of data. Note the minor improvement of the Mass temperature prediction error from the original case (third row) to the case with only the Assembly and Mass temperatures available.

It is not clear why the performance of the model that incorporates only the Assembly and Mass temperatures is as good as it is. It could be that the influence of the Attic temperatures on the Assembly zone defies linear modeling, so an average influence for each mode is the best the linear model can do. The average influence would therefore be unaffected by the availability of the Attic temperature measurements. It should be noted that this hypothesis contradicts the results found from using PHDRT.

The modified parameters of the model formed using both Assembly and Mass temperatures are shown below in Figure 5-35. Note the strong influence of outside temperature on the Assembly zone in Mode 1.

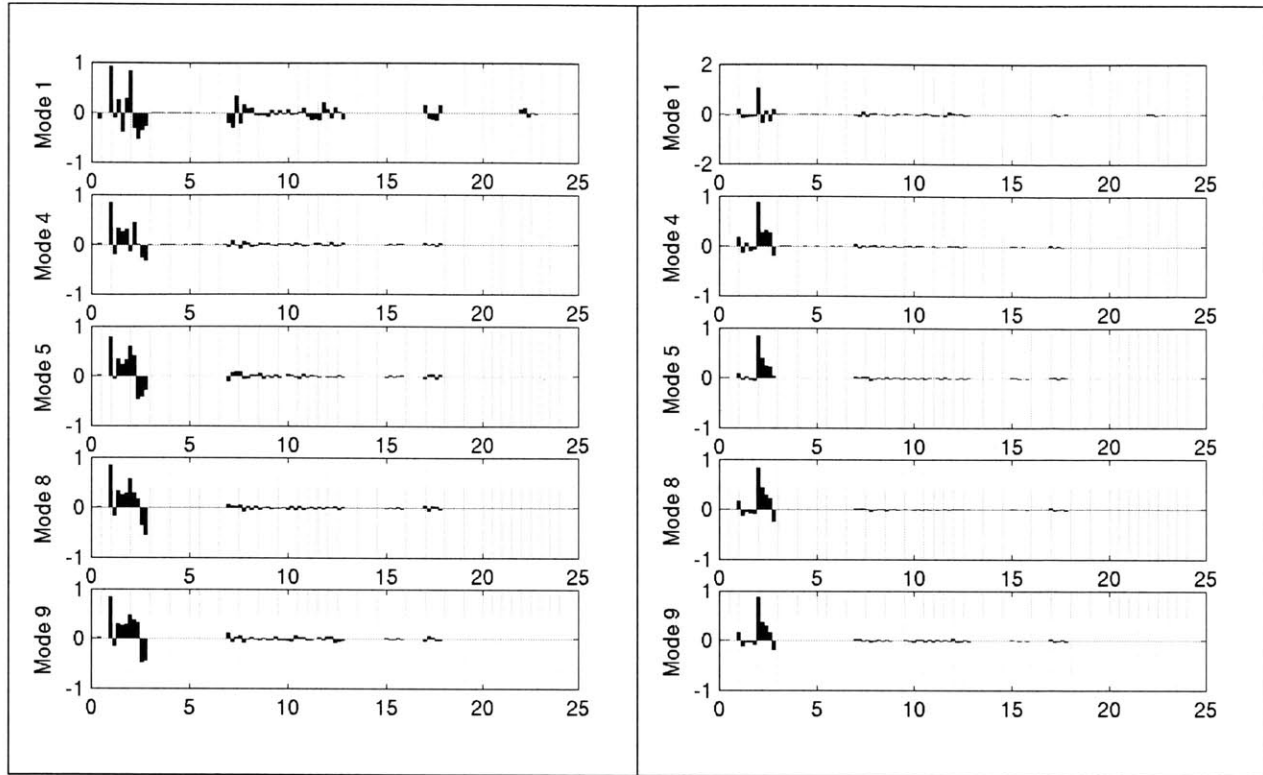


Figure 5-35 Modified model parameters. Left: Assembly zone. Right: Mass zone. The numbering on the abscissas matches the feature numbering in Table 5-13, with the exception of feature 32, which is plotted at 22 in these figures. (Very small bars are evident in some of the plots of features 3-6. This is a numerical and printing artifact.)

In this section, it has been demonstrated that subdivision of the data into modes had a dramatic positive impact on the predictive capability of the models. Of smaller benefit was the incorporation of all building zones into the model. In the next section the possibility of automatically selecting the most important features to be used as model inputs is explored. These could be zone temperatures or any of the measurements taken on site.

5.8 Automated Feature Selection

The general concept of the method outlined in this section has been called greedy optimization or forward/backward selection (Chiu 1996). The issue is that when many elements exist that may influence the calculation of some cost function (*e.g.*, the minimization of model prediction errors) it may not be clear which elements are most important. Some method is needed to identify those elements. The method outlined in this section is appropriate for any model structure, not just for linear models.

Specifically, in the case at hand, there are 32 types of features identified as possible model inputs. Others could be identified, such as time of day, day of year, *etc.* Furthermore, the optimal number of lag terms for each feature is not known. The best possible model will utilize only those features that are relevant and that improve model performance. This section details how those features can be selected. The technique described is suboptimal since local minima

exist in the performance-versus-feature-vector space. Another approach to the problem would be the via genetic algorithms. (See, *e.g.*, (Ferreira *et al.* 2003)).

The algorithm will be described briefly in words, then specifically in a set of steps to be executed. Note that the term “feature” refers to any input to a model, and that $T_{out}(t-1)$ and $T_{out}(t-2)$ are two distinct features even though they share the same feature type and number in Table 5-13.

1. Assume an initial number of lag terms common to all feature types (2 in this case).
2. Initial feature elimination (backward selection)
 - a. Identify the next (or first) feature type with at least one lag term available.
 - b. For this feature type, remove one lag term. In other words, if $T(t-1)$ and $T(t-2)$ are available features, eliminate the feature $T(t-2)$.
 - c. Using a training set, perform regression to compute all model parameters without the identified feature.
 - d. Compute the cost function: rms error on a separate test set.
 - e. Replace the eliminated feature, and return to Step 2a until all feature types have been identified.
 - f. Identify the feature whose removal led to the largest reduction of rms error on the test set. If no error reduction was found, proceed to Step 3. Otherwise, eliminate the identified feature from the feature vector and return to Step 2a using the new feature vector as the starting point.
3. Forward/Backward Selection (repeat this loop a fixed number of times, or until no performance benefit obtained)
 - a. Feature addition (forward selection)

Perform the same operations as outlined in Step 2, except rather than removing a feature from every eligible feature type, add one (perhaps up to a maximum). In essence, those features whose addition most benefits the model are added. Repeat this forward selection a set number of times or until no further benefit is achieved.
 - b. Feature elimination (backward selection)

Perform the identical operations outlined in Step 2. Repeat this backward selection a set number of times or until no further benefit is achieved. Then return to repeat Step 3. Again, Step 3 may be repeated a set number of times, or until no further benefit is obtained.

The use of the term “greedy optimization” becomes apparent after review of the algorithm—one always selects the feature associated with the biggest performance gain at every step of the process.

The structure of Step 3 of the procedure is intended to help the algorithm maneuver past local minima in the cost function surface. Note the underlying assumption that if $T(t-3)$ is relevant, then $T(t-2)$ and $T(t-1)$ are as well. This assumption dramatically reduces the search space of the problem. If it is important to relax this constraint, the same basic algorithm could be used, with only a few minor modifications. (Essentially, every place “feature type” exists, replace it with

“feature”. Also, the maximum number of lag terms becomes 1.) Especially this unconstrained problem may be more efficiently addressed by a genetic algorithm solution.

Finally, it should be noted that the set of features produced by this algorithm are optimized for use in a particular model structure (linear, in this case) and for performance on a particular measure. The sub-optimal set of features could change if either the cost function or the model structure were altered.

In the following section this algorithm is applied to the problem of developing a model for predicting zone temperatures at Broadmoor.

5.9 Application of Automated Feature Selection

The greedy feature optimization described in the previous section was implemented into the modeling process, using the Broadmoor data divided into control modes.

The specific inputs to the algorithm are provided for reference. An initial population of two lag terms was given to each of the feature types used (refer back to Table 5-10 for a listing). Each of the feature types was permitted a maximum of five lag terms. The initial feature elimination (Step 2, above) was performed, followed by seven loops of (two loops of Step 3a and one loop of 3b). If no improvement was found before the seven loops were completed, the algorithm terminated.

The training set and test set were described in Section 5.7.3.1. Recall that to evenly sample the different modes in the training and test sets, the training points and test points were selected as the odd and even points, respectively, of the full training/test period defined in Table 5-9. Several variations of this technique were used: 1) every first point of four was used for the training set and every fourth point of four was used for the test set 2) the points of the data sets were divided into six-hour batches, and the odd batches comprised the training set and the even batches the test set. In the cases examined, there were no distinct advantages to any of these approaches.

The resulting model performance is cataloged in Table 5-15, below. Comparison with the results of the model formed without optimization (Table 5-12) reveal that no striking performance gain was made in this instance.

Zone	Mode	Training Set			Test Set			Validation Set			Complete Set		
		Number of Points	rms error	max error	Number of Points	rms error	max error	Number of Points	rms error	max error	rms error	max error	
1	1	456	0.098	0.462	456	0.096	0.520						
1	4	3121	0.037	0.307	3120	0.038	0.449	1566	0.052	0.258			
1	5	2521	0.060	0.313	2521	0.059	0.335	419	0.056	0.227			
1	8	9360	0.044	0.374	9360	0.044	0.453	8069	0.037	0.298			
1	9	3869	0.046	0.473	3869	0.047	0.389	5342	0.050	0.378			
1	all	19327	0.048	0.473	19326	0.048	0.520	15396	0.044	0.378	0.047	0.520	
1		Pure Simulation Results -->							0.345	1.539	0.416	1.616	
2	1	456	0.015	0.117	456	0.015	0.108						
2	4	3121	0.007	0.165	3120	0.007	0.173	1566	0.007	0.027			
2	5	2521	0.007	0.055	2521	0.007	0.048	419	0.006	0.020			
2	8	9360	0.007	0.086	9360	0.007	0.055	8069	0.007	0.051			
2	9	3869	0.007	0.038	3869	0.007	0.046	5342	0.008	0.065			
2	all	19327	0.007	0.165	19326	0.007	0.173	15396	0.007	0.065	0.007	0.173	
2		Pure Simulation Results -->							0.335	1.403	0.411	1.401	
3	1	456	0.037	0.163	456	0.037	0.165						
3	8	9360	0.029	0.283	9360	0.029	0.265	8069	0.029	0.224			
3	9	3869	0.036	0.344	3869	0.036	0.376	5342	0.032	0.291			
3	10	3259	0.024	0.182	3258	0.023	0.187	1566	0.035	0.333			
3	11	2383	0.037	0.263	2383	0.037	0.305	419	0.023	0.080			
3	all	19327	0.031	0.344	19326	0.031	0.376	15396	0.031	0.333	0.031	0.376	
3		Pure Simulation Results -->							0.418	1.314	0.423	1.650	
4	1	456	0.016	0.058	456	0.017	0.089						
4	8	9360	0.037	0.350	9360	0.038	0.329	8069	0.035	0.319			
4	9	3869	0.020	0.138	3869	0.020	0.129	5342	0.022	0.155			
4	10	3259	0.043	0.217	3258	0.044	0.195	1566	0.058	0.266			
4	11	2383	0.028	0.124	2383	0.026	0.123	419	0.037	0.135			
4	all	19327	0.034	0.350	19326	0.035	0.329	15396	0.034	0.319	0.034	0.350	
4		Pure Simulation Results -->							0.376	1.463	0.380	1.446	
5	12	19327	0.125	1.117	19326	0.124	1.225	15396	0.124	1.254			
5	all	19327	0.125	1.117	19326	0.124	1.225	15396	0.124	1.254	0.125	1.254	
5		Pure Simulation Results -->							1.238	4.849	1.203	6.013	
6	12	19327	0.035	1.387	19326	0.036	1.508	15396	0.027	0.691			
6	all	19327	0.035	1.387	19326	0.036	1.508	15396	0.027	0.691	0.034	1.508	
6		Pure Simulation Results -->							0.594	1.407	0.863	2.439	

Table 5-15 Model performance resulting from greedy feature optimization. (See notes in Table 5-12 for an explanation of table entries. Note the inclusion of a “column” for the test set.)

The parameters for the optimized model are provided (modified as before) below in Figure 5-36 and Figure 5-37. Comparison with the comparable plots for the un-optimized model (Figure 5-24 and Figure 5-25) reveals that the new parameter set is more sparse. Particularly the plots of Mode 1 show strong evidence of the optimization procedure, where the number of features is greatly reduced. For example, the Mode 4 plot for the Assembly temperature shows model dependence on Assembly and Mass temperatures to the almost complete exclusion of all other features (there is a modest contribution from the Basement temperature and the outside temperature (feature types 6 and 7) as well as from the South-facing pyranometer (10) and the electrical loads (17). This is expected, since the building is entirely closed up in this mode. In the un-optimized model, dependence on all zone temperatures is observed.

Unexpected, however, is the result that the Attic 2 temperature (feature number 4) appears in the model for Mode 5 of the Assembly temperature. In this mode, the Attic and Fan Sliders are both shut. Closer examination reveals that parameters multiplying the Attic 2 temperatures sum to approximately 0.001, so unless there is an abrupt change in the Attic 2 temperature, little impact will be made on the model predictions. (In contrast, the parameters for the Assembly and Mass zones sum to 0.87 and 0.11; those for outside temperature sum to 0.01.)

While the exercise described in this and the preceding section did not yield any significant improvement in the predictive power of the models, it did provide some insight or confirmation of what features were most important for model performance. In other buildings, this type of tool may prove to be more essential to improve model performance. We will find this to be the case for the building studied in Chapter 7.

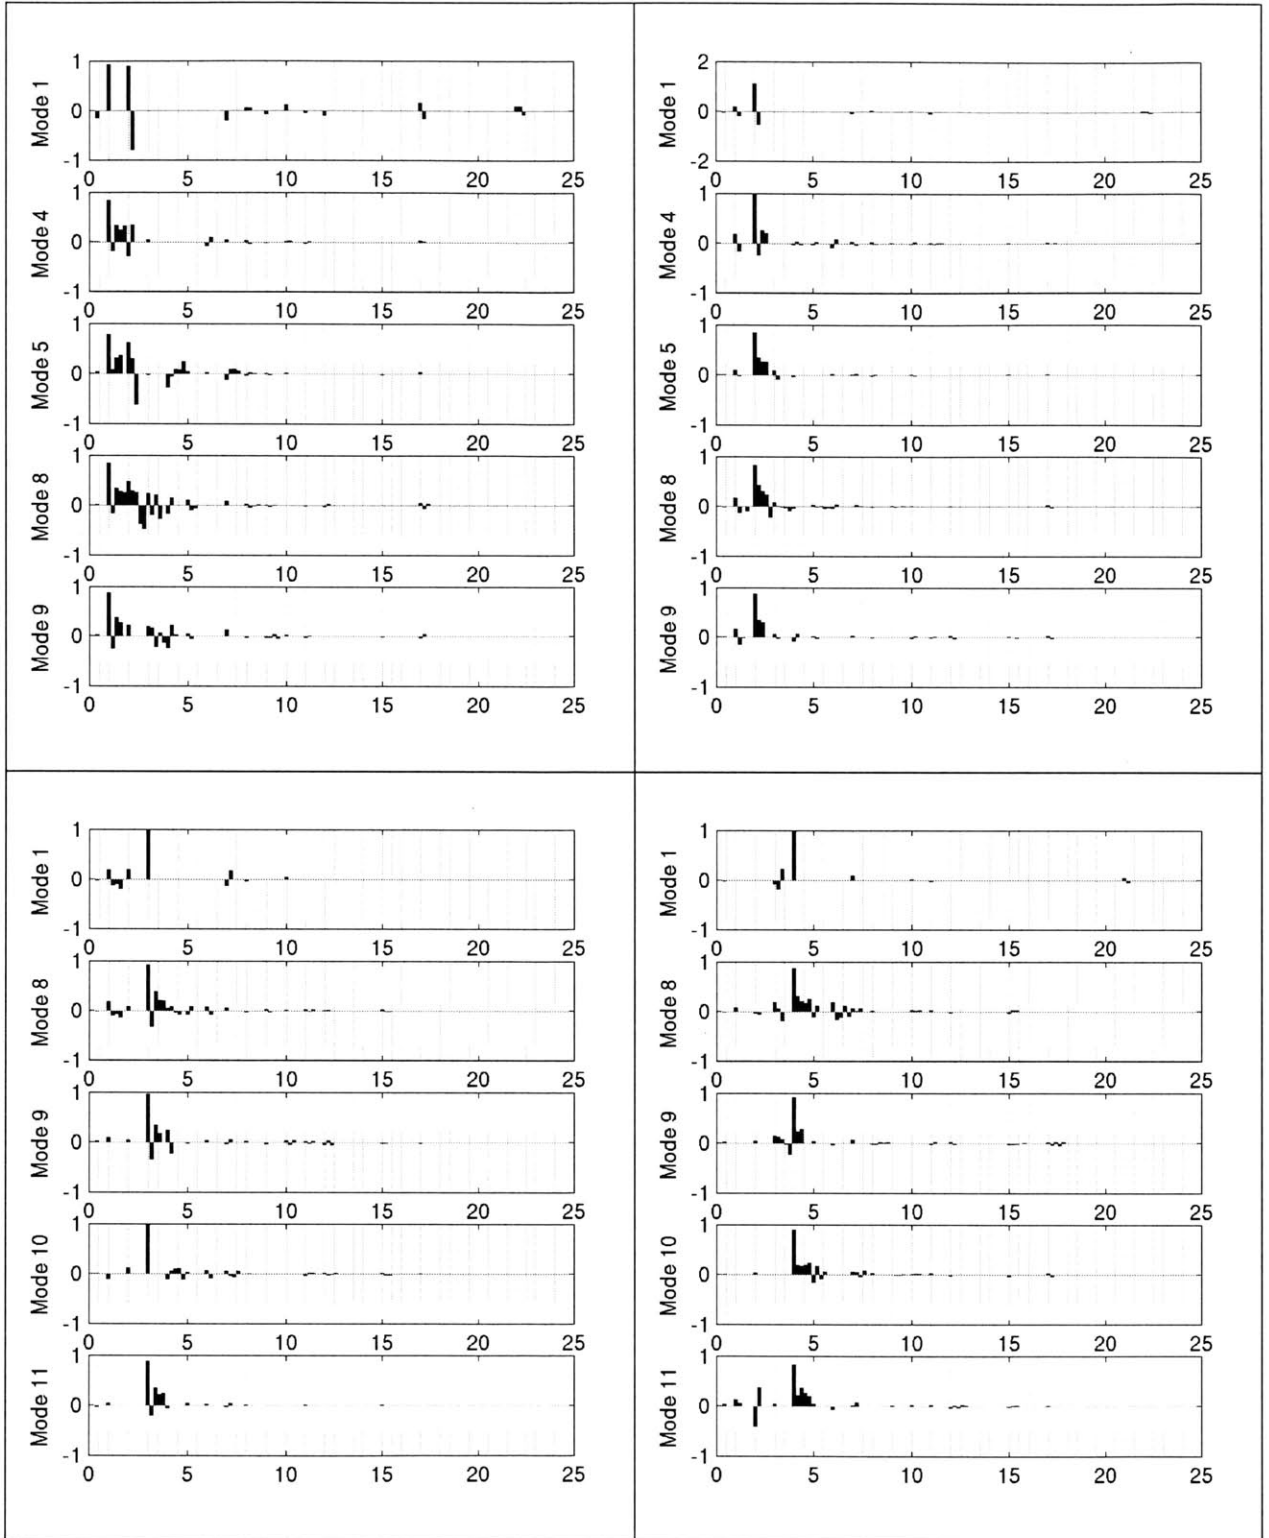


Figure 5-36 Optimized model parameters (adjusted as before) for each mode. From upper left, clockwise, Assembly, Mass, Attic 2, Attic 1. Note that features 30-32 were renumbered to 20-22 to make the plots more readable.




Figure 5-37 Optimized model parameters (adjusted as before) for each mode. Left: Sunspace. Right: Basement.

5.10 Automated Mode Selection

Since the feature optimization yielded relatively little performance benefit compared with the dramatic benefit associated with dividing the data set up into modes, attention was focused once again on the possibility of optimizing the process for selecting modes. The concept is appealing, since one may not know *a priori* how to partition the data. Once the goal of efficient automatic mode selection is in place, the entire modeling architecture is complete, with essentially all of it occurring without user intervention. This section details initial steps towards that goal.

The algorithm that has been developed will be presented, then followed by comments. The algorithm provides a structured way to group, or cluster, common data points together in such a way that the overall performance of the model is improved and, if desired, the model complexity is reduced. For the purposes of this algorithm, the expression “data set” is used to refer to those data points associated with a particular control mode. Each mode has its own data set.

1. Encode all control settings into a binary string, with a 1 indicating that the aperture is open or the device is on and 0 the converse. The string has length N , where N is the number of different controlled devices. This string, or code, represents which devices are active in a particular control mode. Multi-level settings are ignored in the string—partially open is treated identically to fully open.
2. Scan all available data to discover the complete set of unique modes (binary strings). N_{unique} is the number of unique modes.
3. Use the desired model structure (*e.g.*, linear) to develop a model using the data set associated with each mode.
4. Identify all modes with (one-step-ahead) rms error below a tolerance level (*e.g.*, 1×10^{-10}). These modes are populated too sparsely to create a useful model. N_{sparse} is the number of such modes.
5. For each mode identified in Step 4:
 - a. For each of the $(N_{\text{unique}} - N_{\text{sparse}})$ remaining data sets, compute the total number of bits in its associated binary string(s) that differ from that of the mode in question. If a data set is associated with multiple strings, select the smallest total bit difference of all such strings. (This situation occurs when modes have been merged.) Sort the $(N_{\text{unique}} - N_{\text{sparse}})$ sets in order of increasing total bit difference.
 - b. Select those data sets sorted in Step 5a whose minimum total bit differences are less than or equal to the overall minimum bit difference + 1. These are the $N_{\text{neighbors}}$ “neighbors” of the mode in question.

- c. Temporarily combine the data points from the mode to be eliminated with those of each of its neighbors.
 - d. For each of the $N_{\text{neighbors}}$ aggregate sets, generate a model and compute the one-step-ahead rms error on the data points *associated with the particular mode to be eliminated*.
 - e. Identify the neighbor associated with the lowest rms error calculated in Step 5d. Permanently aggregate, or merge, the data points from the mode to be eliminated with this neighbor. Eliminate the mode in question, and permanently associate its binary string with those of the selected neighbor.
 - f. Return to Step 5a until all N_{sparse} modes have been eliminated.
6. At this stage, all that must be done has been done. A collection of “mode clusters” has been generated. If it is desirable to work with fewer modes than currently present, continue to Step 7, otherwise terminate.
 7. Scan the remaining “mode clusters” to identify the one with the smallest one-step-ahead rms error on its own data set.
 8. With the best performing mode cluster:
 - a. For each of the other mode clusters, compute the minimum total number of bits in its associated binary string(s) that differ from any of the strings of the best mode cluster. If a mode cluster is associated with multiple strings, select the smallest total bit difference of all such strings. (This situation occurs when modes have been merged.) Sort these mode clusters in order of increasing total bit difference.
 - b. Select those mode clusters sorted in Step 5a whose minimum total bit differences are less than or equal to 2. These are the $N_{\text{neighbors}}$ “neighbors” of the mode cluster in question.
 - c. Temporarily combine the data points from the best mode cluster with those of each of its neighbors.
 - d. For each of the $N_{\text{neighbors}}$ aggregate sets, generate a model and compute the one-step-ahead rms error on the data points *associated with the best mode cluster*.
 - e. Identify the neighbor associated with the lowest rms error calculated in Step 8d. Permanently aggregate, or merge, the data points from the best mode cluster with this neighbor. Eliminate the best mode cluster, and permanently associate its binary string(s) with those of the selected neighbor.
 - f. Return to Step 8a until the desired number of mode clusters has been obtained.

Comments:

Step 1: Multi-level settings could be incorporated, if desired, by discretizing the settings and assigning binary variables to each setting. However, it would be valuable to link together all the binary variables associated with a particular device. There is no need to pursue this with the fan, since there is already a built-in feature, $(\text{fan speed}) \cdot \Delta T$, that accounts for fractional speeds.

Step 2: In practice, N_{unique} is likely far smaller than 2^N .

Step 4: Many of the modes will be populated by a small number of data points—so small that the number of parameters in the model may exceed the number of points. The error is calculated using the entire data set. The proposed tolerance level will not be exceeded when there are too

few points, since the data will be fit perfectly. Note that if one had a very large data set, one could simply construct N_{unique} models and bypass this algorithm.

Step 7: The idea here is that all “mode clusters” should exhibit approximately the same level of performance (this is debatable, and could be grounds for an altered algorithm). For example, a weighting by the number of points within a mode cluster might be advantageous. The “best” cluster can be merged with another, increasing its error somewhat. It was decided that if an error increase was required (as it is when clusters are merged), the best-performing cluster was a reasonable candidate for the slight degradation in performance.

An interesting aspect of the proposed algorithm is that the clustering described is performance-based. Proximity- or distance-based clustering is common; k-means clustering, c-means clustering are examples of common techniques. In this case, distance is considered—candidates to merge with the best mode cluster are its nearest neighbors—but the criterion used to merge mode clusters is the performance metric rms error.

In the next section, the algorithm will be implemented, and the performance of the resulting model will be compared with the others described in this chapter.

5.11 Implementation of Automated Mode Selection

The modeling procedure was comprised of two stages: 1) automated mode selection, and 2) model formation. This section is divided into two portions: the first devoted to the automated mode selection, and the second to the modeling procedure, results and discussion.

5.11.1 Illustration of Automated Mode Selection

The data set used for the selection of modes was nearly the entire available data set, points 1 – 60765, or August 6 through September 26. As before, the portion where the Sunspace door was opened was excised from the set (points 37285 – 37815). The features used are shown below in Table 5-16. Note that the Basement Door was not used in this exercise since those data were unreliable. The control settings listed were used strictly for generating the binary strings, and were not used as model features.

The general algorithm proposed in Section 5.10 was modified slightly to account for the particular situation at Broadmoor. No mode selection was performed for the Sunspace and Basement zones. In both cases, a single mode described the operation of the zones. A constraint was placed on the algorithm: it was required that all data points gathered while the fan was running be assigned to a common mode. When this constraint was not imposed, pure simulations diverged at times when the fan was turned on. It is believed that the rather limited number of data points associated with fan operation was linked to the problem. Ideally, with more data, one could treat all the fan data as an isolated data set—one that could be potentially divided into modes according to which apertures were open.

ID	Feature	Number of Lag Terms	ID	Feature	Number of Lag Terms
1	T_assembly	5	18	E Louver Door	2**
2	T_mass	5	19	W Louver Door	2**
3	T_attic1	5	20	Conf/Ass. Door	2**
4	T_attic2	5	21	Office Windows S	2**
5	T_sunspace	5	22	Office Windows W	2**
6	T_basement	5	23	Attic Slider	2**
7	T_out	5	24	Fan Slider	2**
8	Solar-N	5	25	Fan Speed	2**
9	Solar-E	5	26	Sunspace Door	2**
10	Solar-S	5	27	Attic2 Windows	2**
11	Solar-W	5	28	Assembly Windows	2**
12	Solar-Horiz	5	29	Basement Door	0
13	sin(Wind-Dir)	0	30	(T_ass-T_att1)*FanSpeed	1***
14	cos(Wind-Dir)	0	31	(T_att2-T_att1)*FanSpeed	1***
15	WindSpeed	5	32	(T_out-T_ass)*FanSpeed	1***
16	Q_bsmt	5*			
17	Q_first	5*			

Table 5-16 Features used in mode selection. Notes: [*] Q_bsmt used for the Basement zone only; Q_first used for all zones but Sunspace and Basement. [] Controls were used to generate binary strings only, they were not used as model inputs. The Basement Door was excluded due to unreliable data. [***] Features 30-32 were used exclusively for models of the Attic 1 (30), Attic 2 (31) and Assembly and Mass (32). In all cases, these features were available only when the fan was on at time t-1 (for making a prediction at time t).**

Using the data over the period defined above, 51 distinct modes were identified. Nine of these were associated with the fan, so 42 “pure” modes and one mode cluster for the fan remained after grouping the fan modes together.

The first clustering of modes was performed, according to Steps 1-6, merging all modes with too few data points with larger clusters. This procedure yielded a total of 25 mode clusters, one of which was the fan cluster. Clusters were merged one at a time, according to Step 7 and 8, until only 5 clusters remained. Figure 5-38 shows the composite one-step-ahead rms prediction error for the Assembly zone as the number of clusters is reduced. As expected, the composite training error increased as the number of mode clusters decreased.




Figure 5-38 Composite one-step-ahead rms prediction errors for the Assembly zone versus the number of mode clusters. (MainModeAutomation.m)

In Table 5-17 through Table 5-19 are shown the binary strings associated with each mode cluster for the cases where 25, 15 and 5 clusters remain. The mode clusters shown were determined for the Assembly zone. For comparison purposes, the reader may want to consult Table 5-2 for the definitions of the manually defined modes used in the remainder of this thesis.

Examination of Table 5-17 through Table 5-19 reveals that the goals set for the algorithm were met only partially. For example, inspect the elements of mode cluster 10 in Table 5-17. Compare the top with the bottom binary string. The top string represents the situation where all apertures are shut except the East Louver Door (original mode 5). The bottom string represents the situation where both Louver Doors, the Conference Room Door, the Office Windows, both Sliders and the Attic 2 Windows were open (original mode 8). These are very different situations, with the greatest difference being that, in the first case, the Assembly zone and Attic zones were isolated, while, in the second case, they were fully connected. One can find other such inconsistencies in the tables.

There are several explanations for the algorithm's counterintuitive selection of modes to cluster together:

1. The most critical factor is that there were few data points in many of the original 51 modes. Consider this example: if the only data points existing for a mode are during a period when there is a small temperature difference between the Assembly and the outside, then it is impossible to distinguish the performance of the building in that mode

from the performance in another mode. Given the low ΔT , all the modes behave in a similar fashion. When modes are clustered together for the sole reason that they behave similarly in one particular temperature regime, then performance of that mode cluster is expected to be poor under different temperature regimes. In other words, model generalization is poor due to lack of sufficient data.

An improved approach to mode clustering may be to entirely ignore those modes populated by fewer than a specified number of points. If those modes were encountered during simulations, then the regressors from the nearest-neighbor mode cluster could be used to make predictions.

2. The criterion for defining the neighbors of a mode cluster could be improved. Specifically, a mode cluster was considered to be a neighbor if the *minimum* total bit difference was below 2 in Steps 7 and 8. With two modes in a cluster, they will differ by at most two bits. With three, the largest bit difference among the modes is as large as four bits, *etc.* While the concept of neighboring modes or mode clusters is helpful, it could be refined.
3. There is no incorporation of basic building information. For example, it is trivial to identify that there are two kinds of apertures in the building: those connecting zones to the outside (such as the Louver Doors), and those connecting zones to one another (Attic and Fan Sliders). One method of incorporating this information into the clustering algorithm would be to impose a constraint that all modes in a cluster must share at least the same basic attribute regarding the inter-zonal apertures. Specifically, if there exists any connection between zones (if at least one of the Attic or Fan Sliders is open) for one mode, then the same must be true for any mode merging with that mode. (Note that such a requirement was imposed upon fan operation.)

Mode Cluster/# of points in cluster	E Louver Door	W Louver Door	Conf/Assembly Door	S Office Window	W Office Window	Attic Slider	Fan Slider	Fan	Sunspace Door	Attic 2 Windows	Assembly Windows
Mode 1											
22527	1	1	0	0	0	1	1	0	0	1	0
Mode 2											
881	1	1	0	0	0	0	1	0	0	1	0
Mode 3											
931	1	1	0	0	0	1	1	1	0	1	0
	1	1	0	0	1	0	1	1	0	1	0
	1	1	1	0	1	0	1	1	0	1	1
	1	1	0	0	0	0	1	1	0	0	0
	1	0	1	0	1	0	1	1	0	1	0
	1	1	1	0	0	0	1	1	0	1	0
	1	1	1	0	0	1	1	1	0	1	0
	1	1	1	0	1	0	1	1	0	1	0
1	1	0	0	0	0	1	1	0	1	0	
Mode 5											
384	1	1	0	0	1	0	1	0	0	1	0
Mode 10											
154	0	1	0	0	0	0	0	0	0	0	0
	0	1	0	0	0	1	1	0	0	0	0
	1	0	0	0	1	0	0	0	0	1	0
	0	1	0	0	1	1	1	0	0	1	0
	1	1	1	0	0	0	1	0	0	1	0
	1	0	1	0	1	0	1	0	0	1	0
1	1	1	0	1	1	1	0	0	1	0	
Mode 11											
242	1	1	1	0	1	1	1	0	0	1	1
Mode 12											
640	1	1	1	0	1	0	0	0	0	1	0
Mode 13											
314	1	1	1	0	0	0	0	0	0	1	0
Mode 14											
2337	1	0	0	0	0	0	1	0	0	1	0
	1	1	0	0	0	0	0	0	0	1	0
Mode 18											
2177	1	1	0	0	1	1	1	0	0	1	0
Mode 21											
344	0	0	0	0	1	0	0	0	0	1	0
Mode 22											
3980	0	0	0	0	0	0	0	0	0	1	0

Mode Cluster/# of points in cluster	E Louver Door	W Louver Door	Conf/Assembly Door	S Office Window	W Office Window	Attic Slider	Fan Slider	Fan	Sunspace Door	Attic 2 Windows	Assembly Windows
Mode 25											
117	1	1	1	0	0	1	1	0	0	1	1
	1	1	1	0	0	1	1	0	0	1	0
Mode 26											
195	0	0	0	0	0	1	1	0	0	0	0
	1	1	0	0	0	1	0	0	0	1	0
	1	0	0	0	0	1	1	0	0	1	0
	0	0	0	0	1	1	1	0	0	1	0
	1	0	1	0	1	1	1	0	0	1	0
1	1	1	0	1	0	1	0	0	1	0	
Mode 27											
457	1	1	0	0	1	0	0	0	0	1	0
Mode 30											
137	1	0	0	0	0	0	0	0	0	0	0
	1	0	0	0	0	0	0	0	0	1	0
	0	1	0	0	0	0	0	0	0	1	0
	0	0	0	0	0	1	0	0	0	1	0
0	1	0	0	0	1	1	0	0	1	0	
Mode 38											
3486	0	0	0	0	0	1	0	0	0	0	0
	0	0	0	0	0	0	0	0	0	0	0
Mode 41											
11371	1	1	0	0	0	1	1	0	0	0	0
Mode 42											
1779	1	1	0	0	1	1	1	0	0	0	0
Mode 43											
2127	1	1	0	0	0	0	0	0	0	0	0
Mode 44											
243	1	1	1	0	0	0	0	0	0	0	0
Mode 45											
522	1	1	0	0	1	0	0	0	0	0	0
Mode 46											
3080	1	1	0	0	0	0	1	0	0	0	0
Mode 47											
1699	1	1	0	0	0	1	0	0	0	0	0
Mode 51											
110	1	1	0	0	1	0	1	0	0	0	0

Table 5-17 Results of Steps 1-6: 25 mode clusters for the Assembly zone. The binary strings associated with each mode cluster are provided. The mode clusters were given the label of the first mode in the cluster (which other modes later joined). The bits in the binary string are set to 1 if the particular aperture is open at all, or if the fan was on at any speed. (MainModeAutomation.m; DisplayModes.m)

Mode Cluster/# of points in cluster	Apertures										
	E Louver Door	W Louver Door	Conf/Assembly Door	S Office Window	W Office Window	Attic Slider	Fan Slider	Fan	Sunspace Door	Attic 2 Windows	Assembly Windows
Mode 3											
931	1	1	0	0	0	1	1	1	0	1	0
	1	1	0	0	1	0	1	1	0	1	0
	1	1	1	0	1	0	1	1	0	1	1
	1	1	0	0	0	0	1	1	0	0	0
	1	0	1	0	1	0	1	1	0	1	0
	1	1	1	0	0	0	1	1	0	1	0
	1	1	1	0	0	1	1	1	0	1	0
	1	1	1	0	1	0	1	1	0	1	0
	1	1	0	0	0	0	1	1	0	1	0
Mode 11											
4648	0	0	0	0	0	1	0	0	0	0	0
	0	0	0	0	0	0	0	0	0	0	0
	0	0	0	0	1	0	0	0	0	1	0
	1	1	0	0	1	0	1	0	0	0	0
	0	0	0	0	0	1	1	0	0	0	0
	1	1	0	0	0	1	0	0	0	1	0
	1	0	0	0	0	1	1	0	0	1	0
	0	0	0	0	1	1	1	0	0	1	0
	1	0	1	0	1	1	1	0	0	1	0
	1	1	1	0	1	0	1	0	0	1	0
	1	1	1	0	0	1	1	0	0	1	1
	1	1	1	0	0	1	1	0	0	1	0
	0	1	0	0	0	0	0	0	0	0	0
	0	1	0	0	0	1	1	0	0	0	0
	1	0	0	0	1	0	0	0	0	1	0
	0	1	0	0	1	1	1	0	0	1	0
	1	1	1	0	0	0	1	0	0	1	0
	1	1	1	0	1	1	1	0	0	1	0
1	1	1	0	1	1	1	0	0	1	1	
Mode 12											
640	1	1	1	0	1	0	0	0	0	1	0
Mode Cluster/# of points in cluster	Apertures										
	E Louver Door	W Louver Door	Conf/Assembly Door	S Office Window	W Office Window	Attic Slider	Fan Slider	Fan	Sunspace Door	Attic 2 Windows	Assembly Windows
Mode 5											
1402	1	1	0	0	0	0	1	0	0	1	0
	1	0	0	0	0	0	0	0	0	0	0
	1	0	0	0	0	0	0	0	0	0	1
	0	1	0	0	0	0	0	0	0	0	1
	0	0	0	0	0	1	0	0	0	1	0
	0	1	0	0	0	1	1	0	0	1	0
	1	1	0	0	1	0	1	0	0	1	0
	1	1	0	0	1	0	1	0	0	1	0
Mode 13											
314	1	1	1	0	0	0	0	0	0	1	0
Mode 18											
2177	1	1	0	0	1	1	1	0	0	1	0
Mode 22											
3980	0	0	0	0	0	0	0	0	0	1	0
Mode 27											
457	1	1	0	0	1	0	0	0	0	1	0
Mode 41											
11371	1	1	0	0	0	1	1	0	0	0	0
Mode 42											
1779	1	1	0	0	1	1	1	0	0	0	0
Mode 43											
2127	1	1	0	0	0	0	0	0	0	0	0
Mode 44											
2580	1	0	0	0	0	0	1	0	0	1	0
	1	1	0	0	0	0	0	0	0	1	0
	1	1	1	0	0	0	0	0	0	0	0
Mode 45											
522	1	1	0	0	1	0	0	0	0	0	0
Mode 46											
3080	1	1	0	0	0	0	1	0	0	0	0
Mode 47											
24226	1	1	0	0	0	1	1	0	0	1	0
	1	1	0	0	0	1	0	0	0	0	0

Table 5-18 Results of Steps 7 and 8: 15 remaining mode clusters for the Assembly zone. The binary strings associated with each mode cluster are provided. The mode clusters were given the label of the first mode in the cluster (which other modes later joined). The bits in the binary string are set to 1 if the particular aperture is open at all, or if the fan was on at any speed. (MainModeAutomation.m; DisplayModes.m)

Mode Cluster/ # of points in cluster	E Louver Door	W Louver Door	Conf/Assembly Door	S Office Window	W Office Window	Attic Slider	Fan Slider	Fan	Sunspace Door	Attic 2 Windows	Assembly Windows	
Mode 3												
931	1	1	0	0	0	1	1	1	0	1	0	
	1	1	0	0	1	0	1	1	0	1	0	
	1	1	1	0	1	0	1	1	0	1	1	
	1	1	0	0	0	0	1	1	0	0	0	
	1	0	1	0	1	0	1	1	0	1	0	
	1	1	1	0	0	0	1	1	0	1	0	
	1	1	1	0	0	1	1	1	0	1	0	
	1	1	1	0	1	0	1	1	0	1	0	
	1	1	1	0	0	0	1	1	0	1	0	
	1	1	0	0	0	0	1	1	0	1	0	
Mode 42												
13150	1	1	0	0	0	1	1	0	0	0	0	
	1	1	0	0	1	1	1	0	0	0	0	
Mode 43												
31551	1	1	0	0	0	1	1	0	0	1	0	
	1	1	0	0	0	1	0	0	0	0	0	
	1	1	0	0	1	0	0	0	0	0	0	
	1	1	1	0	1	0	0	0	0	1	0	
	1	1	0	0	1	0	0	0	0	1	0	
	1	1	0	0	0	0	1	0	0	1	0	
	1	0	0	0	0	0	0	0	0	0	0	
	1	0	0	0	0	0	0	0	0	1	0	
	0	1	0	0	0	0	0	0	0	1	0	
	0	0	0	0	0	1	0	0	0	1	0	
	0	1	0	0	0	1	1	0	0	1	0	
	1	1	0	0	1	0	1	0	0	1	0	
	1	1	0	0	1	1	1	0	0	1	0	
	1	1	0	0	0	0	0	0	0	0	0	
	1	1	0	0	0	0	0	0	0	0	0	
Mode 44												
11522	0	0	0	0	0	1	0	0	0	0	0	
	0	0	0	0	0	0	0	0	0	0	0	
	0	0	0	0	1	0	0	0	0	1	0	
	1	1	0	0	1	0	1	0	0	0	0	
	0	0	0	0	0	1	1	0	0	0	0	
	1	1	0	0	0	1	0	0	0	1	0	
	1	0	0	0	0	1	1	0	0	1	0	
	0	0	0	0	1	1	1	0	0	1	0	
	1	0	1	0	1	1	1	0	0	1	0	
	1	1	1	0	1	0	1	0	0	1	0	
	1	1	1	0	0	1	1	0	0	1	0	
	0	1	0	0	0	0	0	0	0	0	0	
	0	1	0	0	0	1	1	0	0	0	0	
	1	0	0	0	1	0	0	0	0	1	0	
	1	1	1	0	0	0	1	0	0	1	0	
	1	1	1	0	1	1	1	0	0	1	0	
	1	1	1	0	1	1	1	0	0	1	1	
	1	1	1	0	0	0	0	0	0	1	0	
	0	0	0	0	0	0	0	0	0	1	0	
	1	0	0	0	0	0	1	0	0	1	0	
	1	1	0	0	0	0	0	0	0	1	0	
	1	1	1	0	0	0	0	0	0	0	0	
	Mode 46											
	3080	1	1	0	0	0	0	1	0	0	0	0

Table 5-19 Results of Steps 7 and 8: 5 remaining mode clusters for the Assembly zone. The binary strings associated with each mode cluster are provided. The mode clusters were given the label of the first mode in the cluster (which other modes later joined). The bits in the binary string are set to 1 if the particular aperture is open at all, or if the fan was on at any speed. (MainModeAutomation.m; DisplayModes.m)

5.11.2 Model Formation, Results and Discussion

Despite the perceived shortcomings of the results of the mode clustering algorithm, models were created for each mode cluster. The training set and validation set were identical to those detailed in Table 5-9. Pure simulations were performed for several cases: 1) no modes were clustered, with the exception of the fan modes, 2) 25 clusters were formed, 3) 15 clusters were formed, and 4) 5 clusters were formed. The performance of the four models on the entire validation set is listed below in Table 5-20. The different modes encountered in the validation set were: 1, 22, 41, 42, 43, 45, 46, 47. The associated binary strings may be located in Table 5-17.

Number of Mode Clusters	Validation Set					
	Assembly	Mass	Attic 1	Attic 2	Sunspace	Basement
43	0.340	0.299	0.376	0.478	1.188	0.320
25	0.354	0.318	0.368	0.469	1.188	0.322
15	0.414	0.385	0.354	0.468	1.181	0.320
5	0.372	0.355	0.499	0.437	1.187	0.355
Model with Original Mode Definitions	0.338	0.338	0.413	0.375	1.206	0.332
Multi-zone Model with no Modes	0.717	0.666	0.454	0.495	1.224	0.484

Table 5-20 Performance of pure simulations on the validation set. Numbers shown are rms errors given in °C. Models were generated using automatically clustered modes. For comparison, pure simulation performance using the original mode definitions from Section 5.5 is provided. Those numbers were copied from Table 5-12. Also provided are the comparable performance numbers from Table 5-11 for the multi-zone model defined in Section 5.7.3.2. (MainModeAutomation.m)

The performance of the models created for the automatically clustered modes is very close to that obtained using the original mode definitions, and is much better than the case where no modes were used at all. Therefore, it is possible to generate mode clusters automatically (if imperfectly) and obtain much of the performance achieved when “expert knowledge” is used to cluster the modes.

Note that the pure simulation rms errors do not increase monotonically as the number of mode clusters decreases. The key to understanding this observation is the notion of generalization ability. Assuming that modes belong together in a cluster, the greater number of data points in that cluster the more likely the model built for that cluster will be able to incorporate a wide range of operating temperatures, solar radiation, *etc.* Consequently, the model will be able to generalize better and perform well in conditions not used for model training. Returning to the observation, as the number of mode clusters dropped from 15 to 5, the number of points in each cluster increased, and the resulting models were better able to make predictions on the full set of conditions experienced during the validation set. In this case, the improved generalization outweighed the loss of information associated with cluster merging.

Also provided are the results of pure simulation on the complete data set:

Number of Mode Clusters	Complete Set					
	Assembly	Mass	Attic 1	Attic 2	Sunspace	Basement
25	0.380	0.354	0.409	0.442	1.169	0.620
15	0.605	0.522	0.516	0.520	1.165	0.621
5	0.674	0.580	0.497	0.518	1.166	0.707
Model with Original Mode Definitions	0.404	0.395	0.404	0.372	1.175	0.654
Multi-zone Model with no Modes	0.637	0.583	0.417	0.520	1.176	0.535

Table 5-21 Performance of pure simulations on the complete set. Numbers shown are rms errors given in °C. Models were generated using automatically clustered modes. For comparison, pure simulation performance using the original mode definitions from Section 5.5 is provided. Those numbers were copied from Table 5-12. Also provided are the comparable performance numbers from Table 5-11 for the multi-zone model defined in Section 5.7.3.2. (MainModeAutomation.m)

It is evident from the preceding table that the performance on the complete set was relatively poor for the cases where 15 and 5 mode clusters were used. The times when the performance was especially poor were those when the building was completely, or almost completely closed during the night. Examination of Table 5-18 and Table 5-19 reveals that the original Mode 4 (building closed) is clustered with other modes corresponding to original Modes 8 or 9 (building open). This does not have an enormous impact on daytime predictions (such as those occurring in the first part of the validation set when the building was closed), but it does affect the predictions of nighttime behavior. Temperatures significantly lower than measured were predicted for those times due to the influence of the data points from the other modes in the mode cluster. (For example, see the top two strings of mode cluster 11 in Table 5-18, and the top two strings of mode cluster 44 in Table 5-19.)

Ideally, the mode clustering would be carried out using pure simulation errors rather than one-step-ahead errors as feedback. With the current simulation code, this is inefficient and impractical. Additionally, those errors would be determined using a separate validation set containing a full representation of all modes. Were these suggestions followed, improved generalization ability would result and the pure simulation errors would indeed increase monotonically as the number of mode clusters decreased.

The concepts and results presented in this and the preceding section (Sections 5.10 and 5.11) are meant to serve as a proof of concept that automatic mode generation is both possible and fruitful. Further work on the topic is warranted, especially regarding additional constraints placed upon what modes are eligible to be merged together to form clusters.

5.12 Model Testing

It has been demonstrated in the previous sections that the models developed exhibit reliable and accurate performance on the training and validation sets. In other words, the model has been proven to reconstruct the true behavior of the building with moderate error. What is demonstrated in this brief section is that the model (using original mode definitions and optimized features) makes plausible predictions of the building's thermal behavior when

different control modes are imposed. Good performance under these conditions is required for predictive control, where the performance of the building over the next 24 hours, for example, must be evaluated for a variety of control mode sequences.

Shown in Figure 5-39 are the predicted Assembly temperatures associated with five different control strategies over a portion of the first part of the validation set. The “Fan/Shut” curve refers to the control strategy that operates the fan at full speed whenever $T_{\text{Assembly}} > T_{\text{outside}}$ and shuts the building entirely (Mode 4) otherwise. This curve achieves the lowest temperature trajectory of the group. This is the maximum-cooling, maximum-energy strategy. Note that the peak temperatures are 2-3K below those associated with the actual strategy (see Figure 5-12 for details of that strategy).

The next strategy depicted is the “Shut/Shut” strategy, which maintains the building in Mode 4 regardless of the Assembly-outside temperature difference. As expected, the temperatures associated with this strategy do not drop as far at nighttime as they do for the other strategies; nor do they rise as much during the day from the minimum.

The “Open/Shut” strategy keeps the building open (Mode 8) whenever $T_{\text{Assembly}} > T_{\text{outside}}$ and shut otherwise. This approach provides the maximum passive cooling possible. The peaks of the Assembly temperature curve associated with this strategy fall between those of the “Fan/Shut” and the “Open/Open” strategy.

The “Open/Open” strategy involves keeping the building open at all times (Mode 8). As expected under this strategy, daytime highs exceed those associated with the “Open/Shut” strategy. Hot outside air is allowed to enter the building during the daytime hours, raising the maximum temperature by 1-2K.

It is interesting (and reassuring) to note the good agreement between the “Open/Open” curve and the “Measured Assembly Temperature” curve on August 22. The good agreement is consistent with the conjecture made above that the true control settings on that day may not have been accurately reported. If this hypothesis is true, then the error associated with this portion of the first validation set should be removed.

A caveat is warranted at this point. In addition to the strategies shown above, another strategy was attempted, namely running the fan at half-speed whenever $T_{\text{Assembly}} > T_{\text{outside}}$ and shutting the building otherwise. This approach led to fan operation during the daytime hours. Since all the training data for the fan mode was obtained at night or early morning, the model was inappropriate for use during hours with sun. Furthermore, relatively few data points were available with low fan speeds. Temperature predictions diverged in this case. It is critical to ensure that training data are available for conditions comparable to those found under test or validation conditions.





Figure 5-39 Demonstration of model performance under five different control strategies. In the order listed in the legend: 1) Fan at full speed whenever $T_{\text{Assembly}} > T_{\text{outside}}$ (Mode 1); building closed entirely otherwise (Mode 4). 2) Outside temperature. 3) Actual Assembly temperature measured. See mode use schedule in Figure 5-12. 4) Building shut always (Mode 4). 5) Building open whenever $T_{\text{Assembly}} > T_{\text{outside}}$ (Mode 8); closed entirely otherwise (Mode 4). 6) Building open always (mode 8). (CombineModelsforControl.m)

5.13 Performance Comparison with a Physically Based Model

A previous study of the building at Broadmoor (Chang 2002) included the development of a physically based model of the building. The model consisted of two mass nodes (air and thermal mass) and contained inputs such as outside temperature, wind speed and direction, estimated solar radiation (based on a clear-sky radiation model coupled with cloud cover information from a nearby airport) and building usage. No measured electrical data were available, nor were measured solar radiation data. With the author's permission, several plots are provided of the model's performance over three of the weeks studied. These reproduce the author's Figures 220-222 showing the predicted Assembly temperatures compared with the measured Assembly and outside temperatures.







Figure 5-40 Physically based model performance (three plots reproduced with permission from Chang 2002)).

To compare the performance of Chang’s model with that of the model developed in this thesis, the rms error and maximum magnitude error were computed for the aggregate of the three data

sets shown in Figure 5-40. The rms error was 0.74°C and the maximum magnitude error was 3.32°C. For comparison, the pure simulation errors for the Assembly zone listed in Table 5-15 (complete set) were 0.42°C and 1.62°C. The improvement shown is even more striking when it is recalled that the major building apertures were fixed throughout the 2002 study. In other words, what has been shown here is a significant reduction of prediction error for a building operated in a more complex fashion. One should bear in mind that the prediction accuracy is affected by both the modeling approach and the availability of site-specific data (including solar radiation and electrical usage).

5.14 Performance Comparison with Nonlinear Models

This section is included to provide additional confirmation that the use of linear models is appropriate for the task at hand.

The data were divided into modes as before, and the optimized features were used to train a neural net for each mode. Three nodes were used in the input layer and three in the hidden layer. Both were hyperbolic tangent activation functions. The output layer was linear. The training set consisted of the odd points of the training set as defined above. It is possible to incorporate a validation set directly into the training process, so a new validation set was created: the even points of the original training set and all the points of the original validation set. This merging of sets permitted full representation of all modes in the new validation set, in addition to the presence of points wholly unrelated to those in the training set.




Figure 5-41 Selected exogenous inputs and control modes over the first 3000 points of the training set. (MainOptimNNmodes.m)




Figure 5-42 Performance of neural net model over the first 3000 points of the training set; measured (black) and predicted (grey) temperatures. (MainOptimNNmodes.m)

The neural net was trained for 100 epochs using `trainlm.m`⁵⁵ or until errors on the validation set increased. The resulting model (set of models—one for each mode of each zone) was exercised in a pure simulation of the first 3000 points of the training period. Shown in Figure 5-41 are some of the exogenous inputs, the measured Assembly temperature and the control modes over the same period. Performance of the model is depicted in Figure 5-42, where it is evident that severe instability exists. Unacceptable oscillatory outputs occur, generally associated with mode changes. This type of performance was observed when neural networks with different numbers of nodes were used as well as when the KRLS algorithm was used. This type of problem appears to be ill-suited for those particular nonlinear models.

5.15 The Effect of Wind

Multiple attempts were made to incorporate the effect of the wind direction into the zone models. These attempts were not successful. The Sunspace was selected for further study due to the fact that all ΔT 's were known and due to the fact that only a single mode was relevant. The data were divided into subsets based on the direction of the wind: N, E, S, W quadrants. The wind speed multiplied by ΔT was used as a model input for each of the four models dedicated to the wind direction quadrant. The performance of the model based on the four subsets of data performed slightly worse than the simple model⁵⁶.

Another approach to understanding the role of the wind on the temperature of the Sunspace was taken. Shown below in Figure 5-43 is a series of five plots over the period from August 11 to August 22. In the upper plot are shown the Sunspace and outside temperatures. Next are the one-step-ahead model prediction errors for the Sunspace zone. Next is what was called the “standard deviation” of the Assembly temperatures. This measure is simply the rms error of a quartic fit of the Sunspace temperature versus time. The fit was performed over a window of $t \pm 10$ minutes⁵⁷. The residuals of the fit were used to find the rms error for that window. The rms error for the window centered at time t was defined to be the “standard deviation(t)”. The window was slid over all points in the data set. This measure gives an indication of the fluctuations in temperature around the trend line. Next on the plot is the wind speed multiplied by the temperature difference, $(T_{\text{Sunspace}} - T_{\text{outside}})$, followed in the bottom plot by the solar radiation incident on the East, South and West pyranometers.

⁵⁵ Code in the Neural Network toolbox, The MathWorks, Natick, MA.

⁵⁶ `MainNoOptimSunspaceWind.m`. Using NE, SE, SW, and NW quadrants generated even worse pure simulation errors (35% increase versus 8% increase).

⁵⁷ Linear through quartic fits were used, as well as larger window sizes. Little qualitative change in the figure was observed.




Figure 5-43 Probe of the effect of wind on Sunspace temperature. From the top: Sunspace and outside temperatures, one-step-ahead prediction errors, “standard deviation” of Sunspace temperatures, (wind speed)*($T_{\text{Sunspace}} - T_{\text{outside}}$), solar radiation incident on East, South and West pyranometers. (EvaluateWind2.m)

A visual inspection of the third and the fourth plots reveals an apparent correlation between the presence of a wind driving force (represented by the (wind speed) * ΔT) and the “standard deviation” of the sunspace temperature. The solar measurements are shown in the figure since one might imagine that solar radiation fluctuations cause the temperature fluctuations. Examination of the August 20 data, however, reveals that even when the solar data is smooth (clear skies), the Sunspace temperature still displays the same fluctuations as on other days when the solar radiation is fluctuating (*e.g.*, August 14 and 15 have comparable wind profiles, but do not have clear skies). It is safe to conclude that solar radiation fluctuations are not the sole source of Sunspace temperature fluctuations.

The link between the wind and the Sunspace temperature fluctuations is supported by Figure 5-44, below, showing the frequency distribution of the “standard deviation” for times when no wind is present and for times when it is present. For times when wind is present, the “standard deviation” has a larger mean and a much larger tail.

It is interesting to observe that the one-step-ahead Sunspace temperature prediction errors follow the same pattern—when a strong (wind speed)* ΔT is present, prediction errors increase.

The purpose of this discussion is to propose that a link exists between the wind and temperature fluctuations (noise) in the Assembly zone. (See Figure 5-45, below.) The Sunspace was

examined here first because it was convenient to remove the potential confounding influence of electrical loads, which also generally fluctuated during daylight hours (the times when wind was generally present). Since temperature fluctuations are likely the result of air movement past the sensor, one may gather that the wind helps to create air movement in the building, which is beneficial for occupant comfort, but that it appears to have no significant thermal effect (since it was not possible to construct a model where the influence of wind-related features was important).




Figure 5-44 Frequency plot of the "standard deviation" of Assembly temperature for the times when Top: wind speed = 0, and Bottom: wind speed > 0. (EvaluateWind2.m)




Figure 5-45 Proposed impact of wind on Assembly temperature fluctuations. From top: Assembly and outside temperatures, "standard deviation" of Assembly temperatures (calculated as above for the Sunspace), wind speed, $(\text{wind speed}) \cdot (T_{\text{outside}} - T_{\text{Assembly}})$. Note that, in this case, a stronger link between wind speed and "standard deviation" exists than that between $(\text{wind speed}) \cdot (T_{\text{outside}} - T_{\text{Assembly}})$ and "standard deviation". (EvaluateWind3.m)

5.16 On-line Modeling

All of the discussion of developing building models in this chapter was focused on the task of off-line, or batch modeling. If it is possible to assign modes at the outset of the building's operation, then a recursive least squares (RLS) technique (Wunsch 1996) or a Kalman filter could be used to establish model parameters. Of course, one would have to select the appropriate number of lag terms to use, but several models could be developed in parallel, and the one exhibiting the best performance could be used for predictive control purposes.

An alternative approach would involve repetitive batch processing of all the data gathered to date. For example, at the end of each day, the day's data could be incorporated into the existing models. This need not involve processing all the data again (Wunsch 1996).

A concern one might have is that the model may encounter conditions very different from those seen to date in the on-line training. To some extent, this is unavoidable. However, presumably the weather will not change so dramatically from one day to the next that there are no existing data points in the vicinity of the new ones. With a linear model, this concern is much less of an issue than with a nonlinear model.

Of greater practical concern is the issue of the availability of an adequate number of data points to construct a model. If predictive capabilities are required as soon as the building becomes operational, then a simplified mode structure must be employed, where modes are clustered together. As more data become available and it is possible to train models for the individual modes, then the mode clustering can be eliminated. Of course, this approach presumes knowledge of the appropriate modes.

5.17 Summary

A great deal of material was presented in this chapter. This section provides a summary and discussion of some of the main findings.

Potentially relevant features were identified for the purpose of generating models for the six building zones at Broadmoor. The operation of the building was categorized into control modes, each corresponding to a particular grouping of control settings. The concept was that the building's behavior could differ significantly from one control mode to another, and that a separate model for each mode might be beneficial.

The PHDRT algorithm presented in Chapter 3 was implemented on the Sunspace zone, selected because its operation was governed by a single control mode. The results of the study were unexpected: no significant benefit was achieved through the use of the algorithm. A full explanation for this finding was not found, but rather the practical ramifications of the finding were taken into account in the modeling approach developed thereafter.

The bulk of the remainder of the chapter was devoted to developing an accurate model for the building zone temperatures. The model developed was linear in the features, and nearly all the features were simple measurements.

Again, modeling was initiated with the Sunspace, due to its simplicity. Successful modeling of this simple zone gave confidence that the linear modeling approach was worth exploring for the building as a whole. When applied to the six building zones, generally good agreement was found between model predictions and measured temperatures.

However, it was found that a significant improvement in model accuracy was found when the building data were divided into subsets according to which control mode was in effect. Furthermore, the resulting model provided the possibility of making accurate predictions of the building behavior when various control settings were made. That option was not available with the simple, multi-zone model, which merely gave temperature predictions that were averaged over all control settings. A demonstration of how to select the building modes automatically was provided.

Using multi-mode models, the importance of incorporating all of the building zones temperatures into every zone model was investigated. For the Assembly zone, for example, it was found that model predictions suffered little if the only other zone temperature used as a model feature was the Mass zone temperature. The Sunspace, Basement and Attic temperatures were less critical in generating an accurate Assembly model. However, the performance deteriorated when the Mass temperature was removed as a possible feature.

The possibility of optimizing the set of features used as model inputs was explored, with the presentation of a greedy feature optimization algorithm. In this instance, the algorithm did not lead to significant performance improvement, but rather furnished insight into which features were the most relevant.

The models developed were exercised using a range of control strategies to verify that model predictions were plausible. Good performance under various control strategies is critical for effective implementation of model predictive control. The model's performance was also compared with that of the physically-based model developed by (Chang 2002), where it was found that the proposed model exhibited substantial accuracy improvements despite the fact that the building operation was considerably more complex. Comparison was made with a nonlinear model, for further confirmation that a linear modeling approach was justified. Typical behavior of nonlinear models (neural nets, or KRLS models, for example) on this data set included strong oscillations and instability of predictions.

A section was devoted to the impact of wind on the thermal behavior of the zones at Broadmoor. The inclusion of wind-related features was not found to boost prediction accuracy. The role of wind appeared to be, rather, to increase air velocity fluctuations within the building (manifested as noise at the temperature sensors). These fluctuations tend to improve occupant comfort, but are not linked to any net thermal impact on the space. Possible explanations for the limited impact of wind are: 1) that the wind is generally present only during daytime hours, when inside-outside temperature differences were often relatively small and 2), that the building is not well oriented to take advantage of the prevailing winds. Note that the limited role of wind is beneficial at this site, since it is not advantageous to bring increased levels of outside air into the building during daytime hours in the summertime.

Finally, some guidance was provided regarding the implementation of the modeling strategy on an on-line basis.

The dominant question lingering at the close of this chapter is why the nonlinearity predicted by analytical models should be undetectable and unnecessary to incorporate into highly accurate models. The modeling architecture that makes use of control modes removes one source of nonlinearity: dramatic shifts in airflow patterns when the Attic and Fan Slider are closed, for example, need not be assimilated into one model. However, the fact remains that even within a particular mode (Mode 8, for example, with all major apertures open), nonlinear behavior is predicted by analytical models. One origin of the nonlinear behavior could be flow reversals (which were observed at Broadmoor), leading to abrupt changes of heat gain to the Assembly zone as the origin of incoming air switches from the outside to Attic 1. Since one aim of this project was to develop a framework for generating an accurate model suitable for use in temperature control, this topic was reserved for future study.

6 Broadmoor: Optimal Control Strategy Development

6.1 Introduction

With models now available to evaluate different control strategies at the Saltonstall Nature Center at the Broadmoor Wildlife Sanctuary, the next major step is to define clearly the goals for an ideal control strategy. Before taking that step, a brief discussion of control strategies described in the literature will be presented. The remaining contents of the chapter include: environmental goals to be achieved via optimal control at Broadmoor, a presentation of an efficient optimization algorithm incorporating physical insight (provided for 24-hour, 48-hour and 72-hour time horizons), a discussion of the benefits of sub-cooling space temperatures below the minimum acceptable occupied temperature, an extension of the concepts exploited during the optimization to permit optimal supervisory and local control in situations where cooling is possible during the occupied period, and, finally, a discussion of the incorporation of more complex mechanical modes.

6.2 Literature Review

Several broad control categories are reviewed in this section. The section will begin with a discussion of control strategies particularly designed for buildings with natural ventilation or mixed-mode cooling. Strategies for management of thermal mass (typically night cooling) will be reviewed, followed by a discussion of local control strategies. The same topics will be considered for more conventional buildings, to give a wider perspective of the tools that have been applied in the building and HVAC controls field. Finally, some of the factors influencing the quality of the control strategies are presented.

6.2.1 Thermal Mass Management: Naturally Ventilated and Mixed-mode Buildings

This subsection is devoted to the management of building thermal mass, or, in other words, the control of night cooling. The basic challenge uniting all the techniques discussed here is to find an effective and efficient way to cool the building at night via natural or mechanical ventilation in order to moderate indoor temperatures (and, in the case of mixed-mode buildings, reduce energy consumption) on the following day.

All of the techniques reviewed from the literature are heuristic; that is, they are based on fixed IF-THEN strategies that are not optimized (except perhaps by manual adjustment) for the particular building or climate. No model-based cooling strategies were found to be in operation in any actual building.

A number of sources provide case studies regarding existing buildings cooled by natural or mixed-mode ventilation. These provide basic descriptions of how the buildings are controlled. Much of this survey work was completed as part of the IEA Annex 35 (Hybrid Ventilation in New and Retrofitted Office Buildings). See, *e.g.*, (Martin 1995; Aggerholm 2001; Aggerholm 2002).

Multiple sources provide detailed night cooling decision trees (Martin 1995; Liem and van Paassen 1997; BRE 1998; van Paassen *et al.* 1998; BRE 1999; CIBSE 2000; Levermore 2000). To provide the reader with an illustration of some of the strategies, a selection is made from the presentation in (Levermore 2000).

“Average afternoon outside temperature control” was implemented in the Inland Revenue building in Nottingham, UK. Night cooling was used (vents were opened) if the following conditions were met:

1. The average outside temperature between noon and 5pm exceeded 18°C
2. The outside temperature exceeded 12°C
3. The inside temperature exceeded the outside temperature
4. The inside temperature exceeded $15.5 \pm 1.5^\circ\text{C}$

The last condition directs cooling to cease if the inside temperature dropped below 14°C, then to resume as soon as the temperature rose above 17°C.

Another strategy, implemented at the Ionica building in Cambridge, UK, was “Peak inside air temperature and degree-hours control”. Night cooling occurred between 9pm and 6am if:

1. The outside temperature exceeded 7°C
2. The inside temperature exceeded 14°C
3. The peak inside temperature during daytime hours exceeded 24°C
4. The inside temperature at 9pm exceeded 19°C
5. The inside temperature exceeded the outside temperature
6. The degree-hours limit was not reached

The meaning of the final condition may be explained as follows. The number of degree-hours during the occupied hours was calculated, with a base of 21°C. Positive and negative degree hours were included in the tally. During the night-cooling period, the number of *negative* degree hours (again, base 21°C) were calculated. Night cooling was halted as soon as this number exceeded the number of occupied degree-hours (multiplied by a number adjusted based on the amount by which the average inside temperature exceeded 21°C).

Finally, the strategy used at the PowerGen building in Coventry was “Average room temperature and peak outside air temperature control”. In this case, night cooling was used when:

1. The average room/zone temperature at the end of the day exceeded 23°C
2. The maximum outside temperature during the day exceeded 21°C
3. The room/zone temperature exceeded 18°C
4. If condition 3 had been violated, then night cooling resumed when the room/zone temperature rose above 20°C

In all cases described, provisions were made to shut the building in times of high wind and/or rain.

In the first two strategies detailed, the inside temperature may attain a level well below the daytime heating set point (the minimum allowable daytime temperature). In those cases, night cooling would have to cease before occupancy to allow the temperature to rise to a comfortable level without triggering the heating system. This issue was discussed briefly in (Martin 1995). It should be noted that it was claimed by Fletcher (1996) that cooling below the occupied period heating set point had not been shown to be beneficial in practice.

Another strategy of note was presented in (BRE 1999), where the set point used to halt night cooling was adjusted on a daily basis, depending on whether night cooling was used on a given day.

Some strategies described by other authors include the building thermal mass temperature in the “IF-THEN” conditions (Liem and van Paassen 1997; van Paassen *et al.* 1998). In the latter publication, the authors reported little difference in cooling effectiveness among the five techniques examined (using the number of hours during the cooling season exceeding 25°C as a metric).

For systems with fan assist, a mechanism is required for determining if and when to turn on the fans for night cooling to provide adequate cooling to prepare for the following day. This need was mentioned by Martin, but no guidance on the solution of this problem was given in any of the works reviewed.

Several authors approached the topic of night cooling through parametric studies. Kolokotroni (1997; 2001), used simulations to investigate the impact on energy consumption the following day of night cooling starting at different hours of the night (in a building with mechanical air conditioning). Also examined was the impact of the ventilation rate on the following day’s maximum temperature. Similar parametric studies were presented in (Geros *et al.* 1997; BRE 1998). Kammerud *et al.* (1984) also performed parametric studies on a simulated residential building at different U.S. locations to establish the impact on daytime cooling loads of the night ventilation rate as well as the minimum allowed interior temperature during night cooling.

6.2.2 Local Control: Naturally Ventilated and Mixed-mode Buildings

The aim of the research reviewed in this section was to develop a control strategy that, at the very least, provided the possibility of tracking a temperature set point, or maintaining a zone temperature between heating and cooling set points. The particular applications were naturally ventilated and mixed-mode buildings.

In the spirit of the heuristic strategies presented above, Martin (1995) outlined a detailed decision tree for use in a mixed-mode building. The cooling strategy is as follows:

1. If the inside temperature exceeds the cooling set point by 0.5°C, open all apertures
2. If the inside temperature exceeds the cooling set point by 1°C, turn fans on at low speed
3. If the inside temperature exceeds the cooling set point by 2°C, turn fans on at high speed or start using mechanical air conditioning

This type of strategy is consistent with the control strategy outlined by van der Aa (2002) that was in use at the Waterland school building in the Netherlands: if the concentration of CO₂ exceeds 700ppm, open one inlet grill; if it exceeds 1000ppm, open second inlet grill; if it exceeds 1300ppm, turn on the fan.

A similar example is provided by Levermore (2000) for the control of louvers on a rooftop windcatcher: close the louvers if the room temperature drops below 16°C; open in 10% increments for every 1°C rise in room temperature.

According to the publications reviewed, this was the only type of control used in actual buildings.

Van Paassen and Lute (1993) described a PI controller developed for a naturally ventilated test chamber where it was found that flow through the (single) window was a function of opening size only (and not wind speed or temperature difference). The controller adjusted the window opening size to maintain a room temperature set point.

Linker *et al.* (1999) developed a robust control strategy for temperature and CO₂ control in a greenhouse. A neural network was used to map ambient conditions, fan speed and aperture size to the airflow rate through the greenhouse. The control algorithm specified the desired flow rate, which was translated to actuator settings by the neural network. The authors demonstrated improved temperature control performance relative to a PI controller and relative to the robust controller that had no access to the flow rate model.

Dounis *et al.* (1996) proposed using fuzzy logic to control air quality (using CO₂ concentration as a proxy) in a naturally ventilated room with a controlled window. The authors constructed a Mamdani fuzzy inference system (see, *e.g.*, Jang *et al.* 1997) with no optimization of model parameters. The controller inputs were CO₂ concentration and the time rate of change of the CO₂ concentration. The controller was tested via simulation and found to maintain the concentration within a small deadband about the set point. (Fixed, open window settings provided concentrations well below the set point. It was interesting to note that the room temperatures under this regime were several degrees cooler than those obtained under the fuzzy control regime.)

Another application of fuzzy inference systems was presented by Bruant *et al.* (2001). In this case, a mixed-mode test chamber was cooled using a combination of ventilation air and mechanical air conditioning. Fan speed and air conditioner status were set by a Mamdani fuzzy controller whose parameters were optimized using genetic algorithms. A population of controllers underwent selection using simulated controller performance as a cost function. Specifically, the elements of the cost function were zone PMV, zone CO₂ concentration, energy consumption and oscillatory operation of the air conditioner. The optimized controller was used in experiments in the actual test chamber. Results were not conclusive and suggested further study.

Lute and van Paassen (1995) developed a predictive controller for use in a naturally ventilated test room with a controlled window and a heater. An ARMAX model of the space was used for

one-step-ahead predictions. A revised simplex method was used to minimize over the planning horizon a cost function comprised of a linear combination of the following elements: set-point tracking errors, energy consumption and actuator oscillations. Satisfactory tracking performance with little actuator oscillation was reported. Predictive control with a horizon as large as four 15-minute time steps was demonstrated for determining the optimal morning start time for the heater. More details are reported in (Lute 1992), which was unavailable to this author.

6.2.3 Thermal Mass Management: Conventional Buildings

In this section, the scope of thermal mass management is expanded from night cooling (Section 6.2.1) to the more general exploitation of the building's dynamics to achieve certain goals. Heat may be stored in or removed from the building fabric or storage tanks at one time to meet heating or cooling needs at a later time. Effective use of the available dynamics may lead to significant energy savings.

There are two main categories of research presented in this section. The first involves establishing an optimal sequence of set points (*i.e.*, an optimal supervisory controller), and the second involves some form of predictive control. The conceptual difference between the two groups is that, in the former case, the control variable is the set point, while in the latter it is a command to an actuator (or actuators).

Much set-point optimization research has been performed by Braun and coworkers. Braun (1990) used a direct search complex method (Gill *et al.* 1981) to establish optimal hourly set points for a simulated building zone over a 24-hour period. The goal was to minimize energy costs given comfort constraints. The set points included: zone temperature, coil discharge air temperature, chilled-water temperature and pump and fan speeds. The optimization was carried out in two levels; potential zone set-point trajectories were generated, then for each hourly zone set point and corresponding environmental conditions, the (steady-state) optimal plant set points were determined. Overall energy costs informed selection of the next potential set-point trajectory via the direct search complex method. Significant savings potential was demonstrated, with savings relative to standard night setback control in the range of 10-50%.

Morris *et al.* (1994) and Keeney and Braun (1996) extended the research above and used the complex method (Rao 1984) to minimize simulated central plant energy costs over 24 hourly time steps. The decision variables were the hourly set points. The authors found that the problem could be simplified to the selection of cooling rates for two nighttime periods and the time of transition between them, and the application of simple rules for the occupied period: cool to lower end of the comfort zone during off-peak hours, and cool to the upper end during peak hours. On average, the cost savings for the simplified case were 3% smaller than those for the optimal case. Another scenario was examined using a single, variable length pre-cooling period prior to occupancy. In this case, a 5% reduction in cost savings was found relative to the optimal case.

Braun *et al.* (2001) evaluated a discrete set of pre-cooling options using an inverse model developed with building measurement. The impact of the following supervisory control strategies was calculated on cost savings: Night Setup, Light, Moderate and Extended Precool, Maximum Discharge, and Two-, Four-, and Nine-hour Linear Rise.

Further work by Lee and Braun (2004) included a parametric investigation of occupied period set-point strategies. The goal was to limit the peak cooling load (and consequently the maximum electrical demand charge).

Wright and Loosemore (2001) used multi-objective genetic algorithms to generate Pareto-optimal solutions for the objectives of daily energy costs and comfort. The supply air set point for the nine occupied hours (8am-5pm) in a simulated building was optimized by the genetic algorithms.

Kummert (2000; 2001) demonstrated the use of quadratic programming to optimize the set point of a passive solar room with a radiator for backup heating. The objective function for the quadratic programming problem included terms for both energy consumption and comfort (via PPD).

A novel approach to establishing optimal central plant set points was presented by Curtiss and colleagues (Curtiss 1992; Curtiss *et al.* 1993b; Curtiss *et al.* 1994). A neural network was trained using set points and uncontrolled variables such as outside air temperature as inputs and plant energy consumption as the output. The set points were: chilled water supply temperature, supply air temperature and the number of condenser fans in operation. Bounds were established for each of the set points (for equipment longevity) and the trained network was used to select the group of set points within those bounds that minimized energy consumption for a particular set of uncontrolled inputs. The technique involved fixing the set-point inputs at unity, then training the input weights for those inputs only to minimize the output of the entire network. Note that this type of set-point selection is not dynamic in nature. Implied in the approach is that there exists for a given set of conditions (defined by outside temperature, current cooling load and sensible heat ratio at the cooling coil) an ideal group of set points to use. Also implied was that the local controllers used during the training period were the same as those used under optimal supervisory control. Ideally, these local controllers would be optimized in some sense (tracking performance, energy consumption, or a combination thereof).

The same goal of determining optimal set points for a given set of conditions was approached in a different manner by Ahn and Mitchell (2001). Rather than using a neural network to model the cooling plant energy consumption, a quadratic model was used (see also Braun 1988).

Multiple authors have published research on methods for controlling the operation of a thermal storage tank. A review of common approaches was provided by Henze (2003).

Kintner-Meyer and Emery (1995) presented a set-point optimization study for a simulated building with a thermal storage tank and thermal mass. By adjusting the seven system set points for each hour of the 24-hour period, the overall electrical cost was minimized using a nonlinear solver. The impact of the presence of the cold storage tank and of the building location on energy savings was investigated.

Other authors have used a predictive control approach to this problem. Kawashima *et al.* (1996) used a neural network cooling load predictor to determine when ice should be made or the chiller

operated in order to minimize electrical demand charges. This approach was predictive only in the sense of predicting how much cooling would be needed throughout the remainder of the day. Predictions regarding the evolution of the building dynamics or the impact of current decisions on later options were not made.

A method of predictive control using dynamic programming was demonstrated by Henze *et al.* (1997). The goal was to select the set of charge rates for a storage tank leading to the minimal electricity cost for a building operated under real-time electricity pricing. The time horizon used was the period for which the pricing remained fixed. At each time step, the first charge rate of the full horizon was implemented. The optimization was repeated at each time step, incorporating any new information available.

This work was extended to the case where control of building thermal mass was combined with that of cold storage (Henze *et al.* 2004b). In this case, the goal was to minimize electrical costs over a 24-hour period. The two relevant control variables were the room air temperature set point and the rate of charge of the thermal storage (positive or negative). Rather than optimizing the temperature set point over each of the 24 hours in the period, it was assumed to be constant over each type of building operation: unoccupied with off-peak rates, unoccupied with peak rates, occupied with off-peak rates and occupied with peak rates. This simplification reduced the dimension of the temperature set-point optimization problem from 24 to 5. An analogous simplification was not made for the storage charge rate variable. The overall optimization employed a two-stage iterative procedure. The first stage involved the optimization of the passive storage via the air temperature set point (the sequence of charge rates was fixed during this optimization). This optimization was performed using a quasi-Newton method. Using the results of the temperature set-point optimization, the charge rates were optimized via dynamic programming. The entire process was iterated until convergence. As above, the optimization was performed at every time step, implementing at each time only the first control settings.

Many other authors have implemented dynamic programming to optimize control of building systems. It was used to optimize control strategy for a heat pump/heat storage system (Rink *et al.* 1988; Zaheer-Uddin *et al.* 1988; Henze *et al.* 2004b), for a passive solar room with a floor heating system (Ferguson and Scartezzini 1992; Chen 2001), and for a conventionally heated room, using a neural network as a predictive tool (Morel *et al.* 2001). The related Pontryagin maximum principle (applicable for continuous systems) was used to optimize room temperature control (Bloomfield and Fisk 1977).

A further example of the implementation of predictive control was provided in (Mozer *et al.* 1996) where occupant discomfort and energy costs were minimized. A neural network was used to predict occupancy of the residence. Using a simple RC model of the building, optimal heating control signals for the 12-step horizon of 120 minutes were established by exhaustive evaluation of all feasible alternatives. The number of alternatives was reduced by considering only the most likely occupancy patterns.

6.2.4 Local Control: Conventional Buildings

In this subsection, research is summarized regarding the development of optimal and adaptive controllers for use in buildings and building systems. The unifying theme among all the work is

that the controllers were designed to track a set point, possibly in a manner that minimizes a composite energy consumption-tracking error cost function.

Research in two instances has drawn on optimal control theory to develop a linear quadratic regulator (LQR) for a single zone conditioned by a heat pump and storage tank (Zaheer-Uddin 1992) and a linear quadratic Gaussian (LQG) controller for a single zone (Kasahara *et al.* 1998). The latter group discussed the value of incorporating a feedforward preview of the set point to improve control at times when the set point is changing. In both cases, the cost function included energy consumption.

Townsend *et al.* (1986) implemented the Pontryagin maximum principle to optimize temperature set-point tracking in a single room with minimal energy consumption. The optimal control strategy arising from this approach was bang-bang control.

Several self-tuning, or adaptive, controllers have been proposed. Dexter and Haves (1989) proposed a first-order ARMAX plant model whose parameters were continuously updated, subject to a set of “jacketing” rules for robust estimation. Model-based predictive control (Clarke *et al.* 1987a; Clarke *et al.* 1987b) was implemented, with a receding horizon of 10 sample periods. Control action was shown to be strongly dependent on the estimated model parameters, whose accuracy degraded over time. The sampling period was also shown to have a strong impact on control performance.

Seem (1998) developed a self-tuning PI controller for HVAC applications. Key elements of the system’s closed-loop response to a set-point change or a disturbance were identified and used to select proportional and integral gains. A formula was provided for selecting those gains in terms of the current controller gains and a fourth-order polynomial in a quantity called the oscillation ratio. Gain adjustments (in the direction of the optimal gains determined by the equation above) were made gradually, weighted by significance of the set-point change or disturbance.

Curtiss and coworkers approached the local control problem using neural networks. In (Curtiss 1992; Curtiss *et al.* 1993a; Curtiss *et al.* 1993b), a neural network was used to model a heating coil. Model predictive control was performed using the neural network model recursively over the entire horizon. Identical control inputs at every time step were generated by a separate controller neural network. Tracking error at the end of the horizon (or, alternatively, the rms error over the entire horizon) was back-propagated through the cascaded coil model neural networks to optimize the outputs of the controller network. The optimized control was then implemented for the current time step, and the process repeated. Tracking performance was superior to that obtained from a PID controller.

The concept developed in the work above was extended to provide satisfactory control during the period when the neural networks were “learning” the system model (Curtiss 1996; Jeannette *et al.* 1998). The control strategy described above was used whenever the network predictions were accurate to within a given tolerance. Otherwise, PID control was used.

Curtiss *et al.* (1996) also developed a neural network compensator for a PID controller to minimize overshoot, settling time, tracking error and the standard deviation of the controller output. Performance improvement relative to PID-only control was demonstrated.

Several authors have used fuzzy system models to develop local controllers.

Ghiaus and Allard (2001) developed a model of a fan-coil unit using a Takagi-Sugeno fuzzy model (see, *e.g.*, Jang *et al.* 1997). Individual PI controllers were constructed for each of the three consequents (or rules). (It was not mentioned whether the control signal from multiple consequents were combined to form a composite control signal.) The controllers adjusted the control valve to track a room temperature set point. In effect, the fuzzy model assisted in determining the domains for gain scheduling.

Sousa *et al.* (1997a; 1997b) also generated a Takagi-Sugeno fuzzy model for a fan-coil unit. These authors presented a control algorithm that combined model predictive control (with optimal control sequences identified via a branch-and-bound technique) and internal model control (Economou and Morari 1986). The controller aim was to track a room air temperature set point by adjusting the control valve of the unit. Performance of the controller was comparable to that of a well-tuned PID controller.

6.2.5 Predictions of Exogenous Inputs

Actual performance of an optimal set-point trajectory or control sequence over an extended period, such as 24 hours, depends on the prediction quality of the exogenous inputs affecting the building. For a residential building, the critical inputs are outside temperature and solar radiation, while for a commercial building, electrical load prediction becomes important.

Several authors have discussed prediction methods for outside temperature and solar radiation (Chen and Athienitis 1996; Kawashima *et al.* 1996; Kummert 2001; Henze *et al.* 2004a), while load prediction methods have been discussed in (Kawashima *et al.* 1995; Kawashima *et al.* 1996; Henze *et al.* 1997) and others, including those participating in the “Energy Predictor Shootout” (Kreider and Haberl 1994).

The impact of the quality of temperature predictions on control performance was investigated by (Kummert *et al.* 2000; Henze *et al.* 2004a), while the impact of load prediction quality on control performance was assessed by (Henze *et al.* 1997; Yoshida and Yamaguti 2001). The impact of modeling uncertainties on control performance was not found in the sources reviewed.

6.2.6 Literature Review Summary

Several important conclusions may be drawn from the existing literature.

The first conclusion is that all documented night-cooling control strategies are heuristic in nature. Building-specific models were not used to inform the night-cooling strategies for a particular building or location. Such an approach is costly in terms of commissioning time (to establish rules for a particular building) or in terms of comfort or cooling energy. Presented below is an alternative solution that draws on knowledge contained in a building model.

The second conclusion is that the full battery of techniques available and—in some cases—implemented in conventional buildings has not been brought to the realm of control of mixed-mode buildings. Part of the reason for this situation appears to be the lack of suitable building models (for actual buildings) to use with dynamic programming or predictive control.

The third conclusion is that all of the (non-heuristic) control strategies proposed and implemented above were for systems whose temperature, for example, was controlled by varying a control input smoothly or in small steps. Control was not effected by mode switching (unless turning a heater on and off is considered to constitute such a strategy). In no case did the system under consideration display the type of mode-switching characteristics found at Broadmoor (and other full-scale naturally ventilated or mixed-mode buildings).

Several findings from the literature cited above are used in this work. In particular, the concept of focusing an optimization on several groups of hours in a planning horizon rather than on every single hour was adopted from Keeney and Braun (1996) and Henze *et al.* (2004b).

As will be discussed below, the type of thermal mass management investigated with Broadmoor is via predictive control, rather than via a hierarchical supervisory/local control structure. However, for other buildings, situations do occur where the latter approach may be advisable. This matter is discussed below.

Finally, prediction uncertainties (weather and electrical load) are ignored in this work. The reader is directed to the authors cited above for information on this topic.

6.3 Control Limitations and Statement of Goals

The goals for the operation of a naturally ventilated building or mixed-mode building are much the same as for a standard building: to maintain comfortable conditions for the occupants. While humidity plays an important role in occupant comfort, it will not be considered here. For the purposes of this work, dry-bulb temperature alone will serve as a proxy for occupant comfort. Controlling humidity and temperature simultaneously in a such a building remains an intriguing topic of future research.

Without a chiller at our disposal, it is important to recognize that there may be times when it is impossible to maintain comfortable temperatures within the building. As far as any control strategy is concerned, the impact of the lack of chiller is this: the only way to cool the building is to take strategic action when outside temperatures are cooler than those inside. As soon as they exceed those inside, the sole approach available is to attempt to maintain the cool temperatures already present. This inherent limitation greatly simplifies the control options, since outside temperatures exceed inside temperatures for the vast majority of daytime, occupied hours in the summer at Broadmoor.

The question arises of whether it would be reasonable to pursue natural or mixed-mode cooling in the shoulder seasons (Fall and Spring), when outside temperatures are more moderate. At Broadmoor, this approach is less important due to the low internal loads. (Heating control and maintaining warm enough temperatures become more of a concern.) For many office buildings, however, these times of year provide ideal conditions for substantial savings through mixed-

mode cooling. After Broadmoor-specific controls are discussed, this important topic will be revisited.

The two specific temperature control goals for Broadmoor are as follows:

Over the next 24-hour period,

1. Maintain the Assembly temperature between upper and lower set points (within a deadband) during all occupied hours.
2. Use a minimal amount of electrical energy (in this case, for fan operation).

The lack of actuators on apertures is ignored—it will be assumed that apertures can be opened and closed at any time. The goals were stated in order of priority; if it is possible to maintain comfortable conditions, it must be done. The role of the second goal is to dictate the selection of one of multiple feasible (comfortable) solutions. In other words, the optimization problem has been posed in this way:

$$\begin{aligned} \min_{u \in \{\text{control strategies}\}} J &= \text{Energy Consumption}(u) \\ \text{subject to:} & \\ T_{\text{Assembly}} &\leq \text{Max set point}(t) \\ T_{\text{Assembly}} &\geq \text{Min set point}(t) \end{aligned} \quad [6-1]$$

Another valid approach would be to weight the two objectives as follows:

$$\begin{aligned} \min_{u \in \{\text{control strategies}\}} J &= \frac{\alpha}{(\text{Normalization Factor}_1)} (\text{Energy Consumption}(u)) \\ &+ \frac{(1 - \alpha)}{(\text{Normalization Factor}_2)} \sum_{i=1}^{1440} (\beta_{\text{max}}(i)(T_{\text{Assembly}} - \text{Max Set point}(i))^2 + \\ &\quad \beta_{\text{min}}(i)(T_{\text{Assembly}} - \text{Min Set point}(i))^2) \end{aligned} \quad [6-2]$$

where

$$\begin{aligned} \beta_{\text{max}}(i) &= \begin{cases} 0 & \text{if } T_{\text{Assembly}} \leq \text{Max Set point}(i) \\ 1 & \text{otherwise} \end{cases} \text{ and} \\ \beta_{\text{min}}(i) &= \begin{cases} 0 & \text{if } T_{\text{Assembly}} \geq \text{Min Set point}(i) \\ 1 & \text{otherwise} \end{cases} \end{aligned}$$

The approach presented in Equation [6-2] is well suited for optimization algorithms such as simulated annealing or genetic algorithms, where imposition of hard constraints may be problematic. Note that the overall cost is a weighted sum (α is the weighting factor) of the

energy consumption and the squared departures from the deadband. As will be discussed below, it may be valuable to weight negative deviations differently from positive deviations.

Rather than forming a single objective function as the linear combination of several objectives, another approach is to pursue solution of the multi-objective problem, retaining a vector cost whose elements are the individual cost components in Equation [6-2]. Much has been written about this topic, *e.g.*, (Fonseca and Fleming 1998a; Fonseca and Fleming 1998b; Lu and Yen 2001; Wright and Loosemore 2001). Generating Pareto-optimal solutions (Pareto 1906) is more computationally demanding than the optimization associated with Equation [6-2], but one need not specify in advance the weighting factors, α and $(1-\alpha)$. As will be shown below, a Pareto front can give the user important information about tradeoffs between cost and comfort.

Note that other costs such as excessive actuator wear could be weighed at this stage, but such costs are more relevant for shorter-term control. This topic will be discussed below.

To summarize the goals for the control strategy, the aim is to develop a sequence of control settings to apply over the next 24 hours (*e.g.*, 5 pm to 5 pm) in order to maintain the Assembly temperature within the deadband and to minimize energy consumption.

6.4 Solution Approach

The nature of the problem suggests the use of dynamic programming (DP) (Bertsekas 2000), although the development of an optimal stationary policy becomes computationally intractable with the large number of exogenous inputs and zone temperatures involved in the model predictions. It may be feasible to develop a finite-horizon problem, although the large number of zones still creates a computational burden. Fortunately—from a computational perspective—relatively few control inputs are available.

Integer programming (IP) may provide a more efficient solution to the problem. The decision variables (control modes) are all integers, and the equations of the system dynamics would serve as constraints on the evolution of the states. The objective function to minimize would be a combination of energy consumption and tracking error (though the latter could also be imposed as a constraint).

The benefit of approaching the problem from the vantage point of DP, IP, or genetic algorithms and simulated annealing is flexibility—one need not make any assumptions about the nature of the exogenous inputs (as is done below), or of the electricity rate structure. These approaches would be well-suited for developing optimal strategies under irregular, or non-monotonic profiles of the exogenous inputs. The downside of using these techniques is primarily the associated computational burden.

Fortunately, the irregular conditions mentioned above are rare (at least for the case of Broadmoor), and we will find that some very basic assumptions yield dramatic benefits for the optimization process. The approach taken here is crafted after model predictive control (*e.g.*, Soeterboek 1992). However, what has been referred to as “the model” is in actuality a collection of linear models, one for each control mode. Traditional control inputs (specifically, inputs of

known amounts of cooling or heating) are not available; rather we have at our disposal a collection of modes from which to choose the best for a given set of conditions.

The end result of the control optimization is a sequence of control modes to be applied on an hourly basis over the 24-hour period. The time step is arbitrary, and could be finer, though with the uncertainties in predictions of exogenous variables, such precision in the control optimization appears unwarranted.

6.4.1 Further Definition and Reduction of the Control Problem

If we assume that the Attic 2 windows remain open for the duration, there are four modes from which to choose for each of the 24 hours in the optimization horizon. For reference, the modes used in this chapter are summarized in Table 6-1.

Mode	Description
1	Fan on
4	Attic and Fan sliders shut; all doors shut
5	Attic and Fan sliders shut; Louver doors open
8	Attic and/or Fan slider open; Louver doors open

Table 6-1 Summary of modes used in this chapter. Taken from Table 5-2 on page 118.

Theoretically, $4^{24} \sim O(10^{14})$ different sequences would need to be generated and used to simulate the building's thermal progression to establish which is optimal. Two practical considerations, however, dramatically reduce the number. These include:

1. During occupied hours, the fan may not be used since it was not trained at times when solar radiation was present. Furthermore, it was never used when the outside temperature exceeded the inside temperature. (Of course, this is no limitation—for cooling purposes, such use of the fan would never be advised.)
2. With a monotonically decreasing outside temperature at night (observed to be generally the case at Broadmoor), and without the concern of time-of-use electricity charges that may vary sporadically throughout the nighttime hours, it is sufficient to break the night into a few large segments. If this were a high-precision control problem, with highly accurate predictions of exogenous variables, then it would be suitable to subdivide the nighttime hours more finely. Morris *et al.* (1994) and Keeney and Braun (1996), in optimizing the use of thermal storage to minimize energy consumption in commercial buildings, found that breaking the night into two segments introduced a cost penalty of 3% or less relative to the case where each hour's settings were optimized.

Accounting for these practical issues leaves just four periods of the day to optimize: 5-7pm (early evening, with some solar radiation), 7pm – X and X - 7am (two nighttime periods of duration to be established), and 7am – 5pm. The variable X denotes the transition time between the first nighttime period and the second. The definition of nighttime is approximate, and could be adjusted based on local sunrise/sunset times or, more effectively, based on outside temperatures. Ideally, the nighttime period would commence as soon as the outside temperature dropped below the Assembly temperature and terminate as soon as it exceeded it in the morning. The following table presents the options available.

Period Number	Times	Modes Available
1	5pm - 7pm	4,5,8
2	7pm - X	1,4,5,8
3	X - 7am	1,4,5,8
4	7am - 5pm	4,5,8

Table 6-2 Periods for which optimal modes must be selected. Note that the variable X, denoting the transition time between nighttime periods, must also be selected. It may take one of 12 values.

The simplifications and practical accommodations have reduced the number of viable control sequences to $(3 \times 4 \times 4 \times 3 \times 12) = 1728$.

For the purposes of this discussion, it will be assumed that the occupied period is 8am – 5pm, with a deadband for the Assembly temperature defined by a minimum set point of 68°F / 20°C / 293.15K and a maximum set point of 78°F / 25.56°C / 298.71K. During the unoccupied period, no constraints are placed on the Assembly temperature.

6.4.2 Two Solution Methods

The optimization problem may be summarized as follows:

$\min J(u) = \text{Total Fan Energy Consumption over 24 - hour period}$

subject to :

$$T_{\text{Assembly}} \leq \text{Max set point}(t)$$

$$T_{\text{Assembly}} \geq \text{Min set point}(t),$$

where

$$u \in \begin{bmatrix} \{1,2,\dots,12\} & \text{(hour into 12 - hour nighttime period when X occurs)} \\ \{4,5,8\} & \text{(mode for hours 5pm - 7pm)} \\ \{1,4,5,8\} & \text{(mode for hours 7pm - X)} \\ \{1,4,5,8\} & \text{(mode for hours X - 7am)} \\ \{4,5,8\} & \text{(mode for hours 7am - 5pm)} \end{bmatrix} \quad [6-3]$$

Note that the first component of the decision variable, u, refers to the hour of the 12-hour nighttime period starting at 7pm. Therefore, hour 1 refers to a switch between nighttime periods occurring at 7pm. In this example, only a single nighttime period is used; the fourth component of the decision variable sets the mode for the entire night. If $X = 2$, then the first period runs from 7pm-8pm, and the second from 8pm-7am.

The first solution method selected to solve [6-3] relied upon genetic algorithms and simulated annealing. Both approaches were robust and provided optimal or near-optimal solutions within

about an hour on an 800Mhz computer⁵⁸. Such a calculation time requirement could fit into a practical control architecture; calculations could be performed starting around 4pm to determine the initial control settings to employ at 5pm. Several updates throughout the night could be performed to prevent the algorithm from proceeding in an entirely open-loop fashion. For example, new weather predictions and actual temperature measurements could be used to adjust control strategies.

While observing the progress of the two optimization algorithms, it became clear that some incorporation of “expert knowledge” could greatly improve the efficiency of the optimization. As soon as this knowledge was incorporated into the encoding of the decision variables, it was apparent that a much simpler and faster optimization approach would be appropriate. Consequently, the details of the genetic algorithm and simulated annealing solutions are not provided here, but rather the physical insight that led to the simple optimization approach is presented.

The first critical physical insight integral to the second solution method was that it was possible to rank the cooling inputs at every stage of the 24-hour period. Specifically, it was possible to establish the following ranking of cooling inputs, in order of decreasing cooling input (assuming $T_{\text{outside}} < T_{\text{Assembly}}$):

1. Fan operation at full speed (Mode 1)
2. Open Louvered Doors and open Attic/Fan Sliders (Mode 8)
3. Open Louvered Doors, Attic and Fan Sliders shut (cross ventilation, Mode 5)
4. Closed building (all doors and both sliders to the attic shut, (Mode 4))

The ordering is reversed if $T_{\text{outside}} \geq T_{\text{Assembly}}$. While there is nothing surprising about the ranking, including the basic information in the optimization algorithms is critical. This list need not be fixed. Consider the simplicity of evaluating the ranking of cooling inputs for a given day. See Figure 6-1, below.

The pattern displayed in Figure 6-1 is typical, but can be generated in several minutes’ time for a different day if desired. During occupied hours, the principal aim is to retain the cool temperature within the building as best as possible. The mode that best preserves the relatively cool temperatures inside—Mode 4—is therefore appropriate for use whenever $T_{\text{outside}} > T_{\text{Assembly}}$. A ranking was not specifically performed for the 5-7pm period, but by extension of the results of the 7am-7pm period, it appears that Mode 4 is best suited keeping the Assembly temperature as low as possible at times when the outside temperature is relatively warm. Note that none of the mode sequences depicted in the figure satisfies the deadband constraints.

⁵⁸ Each day’s simulation required about 60 seconds. Complete simulation results were stored, so if the algorithm requested duplicate evaluations of a decision variable, additional simulations were not required. Partial simulations were not stored. Doing so would greatly increase efficiency.




Figure 6-1 Trajectories of the Assembly temperature given selected mode sequences. Period from 5-7pm is Mode 4, then period from 7pm-7am shows Modes 1, 8, 5, and 4. All trajectories from 7am-5pm are Mode 4 with the exception of the dashed line (Mode 8) and the dotted line (Mode 5). In all cases, a single nighttime period was used to emphasize the relative cooling capacities of the various modes. Heavy horizontal lines indicate the occupied period deadband. Outside temperature is plotted in grey. (StrategyExamples.m)

To aid in the description of the second critical physical insight critical to the second solution method, consider Figure 6-2, below.

Shown in Figure 6-2 is a set of temperature trajectories generated when the fan is run for a single hour at different times of the night. The control setting for the remainder of the night is Mode 8. What is clear from this figure is again not surprising from a physical standpoint: the best time to operate the fan is when it is coldest outside. In other words, if one has the use of the fan for a limited time, there is no reason to use it at any time other than the final segment of the nighttime

period. This logic may be generalized: if more aggressive cooling is needed to satisfy set point constraints, first add the aggressive cooling to the final, coolest hours of the night.




Figure 6-2 Illustration of the impact of running fan for one hour. All trajectories share the following mode settings: 5-7pm, Mode 4; 7pm-7am, Mode 8 (except for one hour of Mode 1); 7am-5pm, Mode 4. The heavy dotted line (uppermost curve visible 6am and later) shows the case where no fan was used. All other Assembly temperature trajectories contain one hour of fan run time. Bold solid line, 6am-7am; solid thin line, 5am-6am; dotted line, 4am-5am; dash-dot line, 3am-4am; dashed line, 12am-1am. Heavy horizontal lines indicate the occupied period deadband. Outside temperature is plotted in grey. (Demo1HrFan.m)

These two physical insights simplify the optimization dramatically. First, a ranking is available of the relative cooling capacities of the various modes. Second, it is clear that as the nighttime hours progress, cooling capacity should never be decreased. Using this information, the nighttime mode sequence and mode switch time can be reduced to a single integer that conveys an overall “cooling capacity” of the mode *sequence*. The table below illustrates the code.

Code	Switch from Mode 8 to Mode 1 at the beginning of hour:	Code	Switch from Mode 5 to Mode 8 at the beginning of hour:	Code	Switch from Mode 4 to Mode 5 at the beginning of hour:
1	1 (all Mode 1)	13	1 (all Mode 8)	25	1 (all Mode 5)
2	2	14	2	26	2
3	3	15	3	27	3
4	4	16	4	28	4
5	5	17	5	29	5
6	6	18	6	30	6
7	7	19	7	31	7
8	8	20	8	32	8
9	9	21	9	33	9
10	10	22	10	34	10
11	11	23	11	35	11
12	12	24	12	36	12
				37	never (all Mode 4)

Table 6-3 Encoding of the nighttime mode sequence. Cooling capacity increases as code numbers decrease.

The ordering in the table reflects the findings in Figure 6-1. At another building, a different ordering may be found. Since it was also found in Figure 6-1 that Mode 4 was best suited for retaining cool temperatures during the periods when outside temperatures are warmer than Assembly temperatures, the optimization of the *entire* mode sequence can be reduced to the determination of the best code number. Further aiding the optimization is that the code numbers are monotonic in cooling capacity. While the execution of an optimization via GA or SA would certainly be aided by this encoded physical insight, a simple binary search is also fully capable of solving the problem.

Before proceeding with a stepwise presentation of the algorithm, one important underlying assumption should be stated again clearly. It was assumed that outside temperatures declined monotonically throughout the night to a global minimum, from which they rose during the day. A departure from this basic pattern could render the encoding in Table 6-3 inappropriate. In such a situation, it may be desirable to either perform a GA or SA optimization (perhaps with hourly increments during the nighttime period—expensive) or to accept that the mode sequence developed using Table 6-3 is suboptimal. Additionally, in a building strongly influenced by the wind, an irregular nighttime wind-speed profile, or even a monotonically decreasing wind speed could also render the encoding inappropriate.

It is assumed that the reader is familiar with a binary search. The basic elements of the complete algorithm used in this work are provided:

1. Perform simulations akin to those in Figure 6-1 to determine the relative cooling capacities of the various modes. If the scenario using the greatest cooling capacity does not provide maximum occupied temperatures (M.O.T.'s) below the deadband maximum,

terminate. The optimal solution is to use the greatest cooling throughout the nighttime period.

2. Using the results of Step 1, create a mode sequence encoding similar to that in Table 6-3.
3. Starting from code 1, increase the code by 12 repeatedly until the mode sequence leads to M.O.T.'s exceeding the deadband maximum. The code associated with this mode sequence serves as the upper bound of the space to be searched via binary search. Code 1 serves as the accompanying lower bound.
4. If, during Step 3, code 13 is found to be fully within the deadband, terminate. This is the optimal solution.
5. Perform a binary search between the upper and lower bounds of the code as determined in Step 3. Determine the warmest mode sequence (largest code number) whose M.O.T. falls below the deadband maximum. If this code number ≤ 13 , terminate. This is the optimal solution. Otherwise, proceed to Step 6.
6. Using the upper and lower bounds established in Step 3, perform a binary search to determine the coolest feasible mode sequence (smallest code number). Terminate. This is the optimal solution.

Notes:

Step 1: It may be necessary to terminate here due to insufficient cooling capacity. Since it is assumed that meeting the temperature constraints are more important than minimizing the energy consumption (indeed, that is why they were set as constraints), the mode sequence providing the smallest constraint violation is selected.

Step 4: Given a choice, for a passively cooled building, it is safest to implement the coolest feasible passive strategy. Such an approach provides some additional margin of comfort if the day is warmer than predicted. (Given outside temperatures greater than inside temperatures, it is a simple matter to raise the inside temperature during the day if the predictions err in the other direction.) For example, if it is found that code 13 yields a fully feasible solution, no code number below it will yield a fully feasible solution with a smaller (than zero) energy consumption. Code numbers above it may also lead to feasible solutions, but the corresponding M.O.T.'s will be higher.

Step 5: The reader may have noticed that the deadband minimum is not accounted for in Step 5. If the warmest mode sequence (highest code with M.O.T. below the deadband maximum) violates the deadband minimum, nothing can be done about it⁵⁹. This is so because of the monotonicity of the mode sequences as organized by the codes in Table 6-3. Any attempt to satisfy the deadband minimum constraint necessarily leads to a violation of the deadband maximum constraint. It is assumed that violations of the minimum constraint are preferable to those of the maximum constraint. Note that these constraint violations could be weighted (as in Equation [6-2]) if a compromise is desired.

Step 6: This step exists in order to implement the concept described in the note to Step 4.

⁵⁹ Although a potential work-around is provided in Section 6.5.




Figure 6-3 Trajectories investigated to evaluate the optimal mode sequence. For this 24-hour period, the optimal mode sequence was: Mode 4 from 5-7pm, Mode 8 from 7pm-4am, Mode 1 from 4am-7am and Mode 4 from 7am-5pm. The optimal fan run time was three hours. Up to 5 hours fan run time yielded a feasible trajectory (switching to Mode 1 at hour 8, or 2am). The range of feasible trajectories is outlined by dashed lines. Heavy horizontal lines indicate the occupied period deadband. Outside temperature is plotted in grey. (ExampleOptim1day.m)

The trial trajectories evaluated during the control sequence optimization for the 24-hour period of 5pm, August 19 through 5pm, August 20 are provided in Figure 6-3, above. The order of evaluation was code: 1, 13, 7, 10, 12, 11⁶⁰. Note that the solution of the optimization—code

⁶⁰ Codes 12 and 11 had to be evaluated to confirm that code 10 was indeed the warmest feasible solution.

10—was obtained after only six simulations⁶¹. Calculated in an analogous fashion was the coolest feasible solution: code 8.

The trajectories evaluated for the 24-hour period of 5pm, September 6 through 5pm, September 7 are provided in Figure 6-4, below.




Figure 6-4 Trajectories investigated to evaluate the optimal mode sequence. Note that the four curves labeled Mode 1, 8, 5 and 4 were all that was required to establish a feasible solution: code 37. The other trajectories were necessary for the computation of the minimum feasible solution: code 35. Curves for codes 35 and 37 are plotted with dashed lines. For this 24-hour period, the optimal mode sequence was: Mode 4 from 5pm to 5am; Mode 5 from 5am to 7am; Mode 4 from 7am to 5pm. The range of feasible trajectories is outlined by dashed lines. Heavy horizontal lines indicate the occupied period deadband. Outside temperature is plotted in grey. (ExampleOptim1day.m)

⁶¹ After preliminary simulations used for ranking cooling capacity, some of which could be reused.

6.4.3 Sensitivity Analysis and Pareto Optimality

The material in the previous subsection (6.4.2) was developed to solve the optimization problem posed in Equation [6-3] as faithfully as possible. The material in this section is provided to illustrate to the building operator the impact of the particular constraints imposed. For example, one might be interested determining whether, if the constraints are slightly relaxed, it would be possible to eliminate or significantly reduce fan usage.

For the same 24-hour period beginning at 5pm on August 19, a series of simulations was executed to establish the relationship between the number of hours of fan usage and the M.O.T. The trade-off curve, or Pareto front, is shown in the upper plot of Figure 6-5. The curve is comprised of Pareto-optimal points since, according to the findings in the previous subsection, it is impossible to attain a lower M.O.T. using a smaller number of fan hours⁶². One may learn from this plot that the deadband maximum would have to be increased to approximately 25.75°C to reduce the fan run time from three to two hours. From a different perspective, reducing the fan hours from three to two carries a temperature penalty of only 0.2°C. Also indicated is the minimum acceptable M.O.T. This is the M.O.T. associated with the coolest feasible strategy.




Figure 6-5 Sensitivity analysis. Above: Pareto plot showing hours of optimal fan use versus the associated maximum occupied temperature. Below: Plot of the minimum occupied temperature versus the maximum occupied temperature. Quadrants of the plot indicate feasible points. (ParetoExample.m)

⁶² For a given number of fan hours, it was determined that the most effective time to operate the fan was at the end of the nighttime period.

The lower plot of Figure 6-5 reveals how the minimum occupied temperature relates to the maximum occupied temperature. A linear pattern is evident. The maximum set point (vertical line) and minimum set point (horizontal line) define the feasible region.

Plots such as those in Figure 6-5 provide the building manager with information that could lead to more flexible, and perhaps less costly, operation of the building than that provided by pure adherence to the optimal strategy.

6.5 Sub-cooling

An additional strategy may be useful to explore if difficulty arises in meeting the constraints defined by the deadband. This concept has been explored by Armstrong (2004) and others. The nighttime period is divided into three segments, rather than two. The final segment also has variable length. The idea is that if it is possible to cool the building well below the minimum set point (hence the term sub-cooling), time will be required before occupancy for the temperature to warm up again to acceptable levels. The third period of the night provides this opportunity. The mode to use during this period is the one with the lowest cooling capacity: Mode 4.

Optimization proceeds in a step-wise fashion. First, the optimal mode sequence is computed assuming the third nighttime period has zero length. It is repeated for the cases where the third period has length one hour, two hours, *etc.*, until an overall optimal solution is found.

Consider the trajectories plotted in Figure 6-6, below. The dashed lines depict the warmest and coolest feasible trajectories when the third nighttime period has zero length. The solid lines depict the warmest and coolest feasible trajectories when the third nighttime period is one hour. Note the behavior of the lower solid line (the coolest feasible trajectory under sub-cooling operation) around 8am. The warming hour from 6-7am has allowed the Assembly temperature to rise more than it did when no sub-cooling/warming was performed. In other words, this technique permits tighter control of the Assembly temperatures during the occupied hours. Shown on the plot are dashed horizontal lines indicating the extent of a revised deadband whose constraints could be satisfied with this technique. Notice that the revised minimum set point would be violated by the mode sequence with code 8.

This increased performance flexibility is achieved at the cost of increased fan energy consumption. As discussed earlier, the optimal time to operate the fan is during the coolest hours of the night. Eliminating the possibility of running the fan during one of the coolest hours (6-7am) forces the fan to run at warmer hours, thereby increasing the number of runtime hours required. For comparison, the coolest and warmest feasible mode sequences—defined by codes 8 and 10—require five and three hours of fan time. When the sub-cooling/warming strategy is used, the coolest and warmest feasible mode sequences require eleven and five hours of fan operation.

It may have been noted that a one-hour warming period was already being employed: between 7 and 8am. In all examples shown, nighttime Assembly temperatures dropped well below the minimum occupied temperature, but then rose to acceptable levels during the hour between 7 and 8am. The key message from this section is that sub-cooling followed by warming provides a

potent tool for achieving comfortable occupied temperatures. Used in excess, as demonstrated in the example in Figure 6-6, it leads to unnecessary consumption of energy.

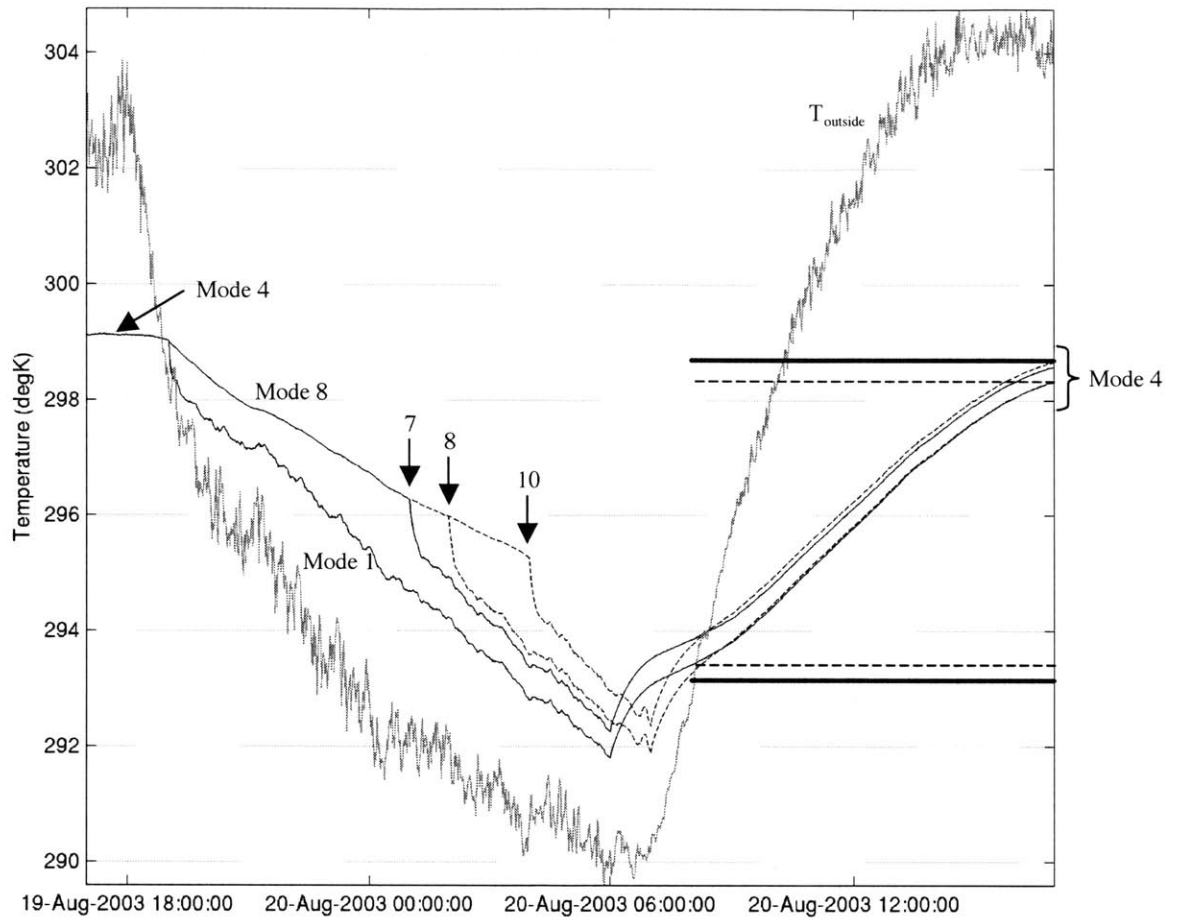


Figure 6-6 Trajectories investigated to evaluate the optimal mode sequence with the possibility of sub-cooling. For this 24-hour period, the optimal mode sequence was: Mode 4 from 5-7pm, Mode 8 from 7pm-4am, Mode 1 from 4am-7am and Mode 4 from 7am-5pm. The optimal fan run time was three hours. Up to 5 hours fan run time yielded a feasible trajectory (switching to Mode 1 at hour 8, or 2am). The maximum and minimum feasible trajectories are plotted with dashed lines. Trajectories associated with sub-cooling followed by a one-hour period of warming are plotted with solid lines. The maximum feasible sub-cooling trajectory is generated when the modes switch from 8 to 1 at 1am, then to 4 at 6am. The minimum feasible sub-cooling trajectory employs Mode 1 from 7pm-6am, then Mode 4 from 6-7am. Dashed horizontal lines indicate the more stringent constraints of a deadband that could be satisfied using the sub-cooling technique. Heavy horizontal lines indicate the occupied period deadband. Outside temperature is plotted in grey. (ExampleOptim1day_Subcool.m)

6.6 Multi-day Optimization

To this point, the optimization horizon examined has been 24 hours. Such a window is reasonable to consider for optimizing the mode sequence to prepare for the coming day. Another interesting situation to consider is when a multi-day weather forecast is available and a hot day is forecast in two days' time. Can anything be done on day 1 to prepare for the hot weather on day 2? Can the total energy consumption over the coming 48 hours, or even 72 hours, be minimized, while meeting deadband constraints on all days? It is these questions that are addressed in the following two subsections of this chapter: Sections 6.6.1 and 6.6.2.

6.6.1 Two-day Optimization

The strategy presented in this subsection for optimization of mode sequences over 48 hours relies on the same assumptions required in the one-day optimization (monotonically decreasing nighttime temperatures, and the possibility of ranking modes in terms of cooling capacity). An additional simple physical insight is required to make efficient progress with this more complex problem: the cooling strategy of day one should never be relaxed to prepare for the next day. In other words, the only impact a second day can have on the first day is to require a more aggressive cooling strategy on the first day. A more relaxed strategy on day 1 can only violate constraints on day 1 (unless the cooling strategy is entirely passive) and make it more difficult to achieve a feasible solution on day 2.

Bearing this in mind, the following two-day optimization algorithm is proposed:

1. Compute the warmest feasible mode sequence for day 1. Name the sequence and associated code A.
2. Compute the coolest feasible mode sequence for day 1. Name the sequence and associated code B.
3. Using the temperatures at the end of the A trajectory on day 1 as a starting point, optimize the day 2 mode sequence and calculate total energy consumption over the two days. Does a feasible solution exist for day 2?
 - a. Yes: Increase the aggressiveness of cooling on day 1 (maximum allowed is B).
 - i. Rerun day 2 using temperatures at the end of the more aggressive day 1 as a starting point. Compute the total energy consumption and compare with that found in Step 3. Is the total reduced?
 1. Yes: Return to Step 3a.
 2. No: Select from among all two-day sequences examined the one with the lowest total energy consumption. This is the optimal solution.
 - b. No: Given B on day 1, does a feasible solution exist for day 2? Compute the total energy consumption over both days.
 - i. Yes: Reduce the aggressiveness on day 1 (minimum allowed is the strategy one level more aggressive than A)
 1. Rerun day 2 using temperatures at the end of the less aggressive day 1 as a starting point. Compute the total energy consumption and compare with that found in Step 3b. Is the total reduced?
 - a. Yes: Return to Step 3bi.

- b. No: Select from among all two-day sequences examined the one with the lowest total energy consumption. This is the optimal solution.
 - ii. No: A solution meeting deadband constraints on both days is impossible. Assuming that it is better to violate the minimum temperature constraint than the maximum temperature constraint, discard the former. Using B and the code 1 strategy (maximum possible cooling) as bounds for a binary search, locate the warmest feasible solution for day 1 that leads to a feasible solution on day 2. Call this sequence C. Calculate the total energy consumption for C and for day 2.
 - iii. Increase the cooling aggressiveness on day 1, starting with the sequence one level more aggressive than C.
 - 1. Rerun optimization for day 2 using temperatures at the end of the more aggressive day 1 as a starting point. Compute the total energy consumption and compare with that found in Step 3bii. Is the total reduced?
 - a. Yes: Return to Step 3biii (until the most aggressive cooling level on day 1—code 1—is reached)
 - b. No: Select from among all two-day sequences examined the one with the lowest total energy consumption. This is the optimal solution.

When this strategy was applied to Broadmoor, it was found that the building contained insufficient thermal storage for there to be a significant lasting impact of one day's strategy on the next day. Consider Figure 6-7, below. To accentuate any impact of day 1's strategy on day 2, the minimum set point for the deadband was reduced to 290K/16.85°C/62.33°F. The warmest feasible strategy on day 1 (A) was code 12, with one hour of fan operation. The coolest feasible strategy (B) was code 1, with 12 hours of fan operation. The minimum-energy (warmest feasible) solutions for day 2 were computed using the temperatures at the end of day 1 as starting points. One optimization used the temperatures at the end of A, the other the temperatures at the end of B. The sequence following A was code 10, with three hours of fan operation, while that following B was code 11, with two hours of fan operation. Note that the temperature trajectories coincide at the end of day 2.

The choice is clear: implement code 12 followed by code 10 with a total of four hours of fan operation, or implement code 1 followed by code 11 with a total of 14 hours of fan operation. In this particular case, the more aggressive cooling on day 1 is wasteful; while a modest reduction of fan run time on day 2 can be achieved, the overall fan operation is prohibitive⁶³. Furthermore, recall that the thermal impact of aggressive cooling on day 1 was exaggerated by lowering the minimum set point to permit sequence B to be code 1.

⁶³ In another scenario with a hotter day 2, the more aggressive cooling on day 1 may be costly, but essential to maintain comfortable temperatures on day 2.

For the two-day optimization algorithm described above to be useful, it would need to be implemented in a building with significantly greater thermal storage capacity than found at Broadmoor. (A discussion of thermal time constants at Broadmoor follows in Chapter 8.)




Figure 6-7 Two-day optimization. For day 1, the warmest feasible solution (A) was code 12, with one hour of fan operation, while the coolest feasible solution (B) was code 1, with 12 hours of fan operation. Commencing day 2 optimizations at the end of day 1 A and B sequences yielded, respectively, code 10, with three hours of fan operation and code 11, with two hours of fan operation. Outside temperature is plotted in grey. (ExampleOptim2day.m)

6.6.2 Three-day Optimization

Increasing the optimization window to three days could be fruitfully applied to the operation of a building over a weekend to prepare for the coming Monday. The optimization routine is a logical extension of that presented for the two-day optimization. Rather than present an exhaustive algorithm for this scenario, the highlights and key concepts will be presented.

As before, we recognize that it will never be advantageous to reduce the cooling “effort” on one day for the purpose of reducing it on the next. The only possibly beneficial approach is to make the cooling “effort” more aggressive on one day to lower the energy consumption on a later date.

Note that the key zone temperature—the Assembly zone in the case of Broadmoor—at the end of the day is a function of the mode sequence for that day. If the warmest feasible sequence is implemented (A), then the temperature at day’s end will be at a maximum. If the coolest feasible sequence is implemented (B), then the temperature at day’s end will be at a minimum. These two scenarios provide bounds for the starting temperature of the following day.

Consider, first, the situation where there exists a feasible mode sequence to ensure occupied temperatures within the prescribed deadband over a three-day period. Implications of the preceding remarks are as follows:

1. The code for the day 1 strategy is bounded between [A, B].
2. The code for the day 2 strategy is bounded between [BA, AB], where BA denotes the code for day 2 associated with use of strategy B on day 1 followed by A on day 2, and AB the reverse.
3. The code for the day 3 strategy is bounded between [*BA, *AA], where *BA denotes the code for day 3 associated with: either A or B on day 1, B on day 2 and A on day 3. Since the strategy on day 2 determines the start temperatures for day 3, the strategy used on day 1 is not relevant here. Note that strategy B on day 3 is not worth considering since no preparations are required for day 4.

It is assumed above that the element limiting strategy B is the lower set point of the deadband rather than the cooling capacity of the fan. In other words, it is assumed that there exists adequate cooling capacity to depress the temperature below the minimum occupied set point. If this is not the case, then the temperature at the end of day 2 is indeed a function of the strategy employed on day 1. Consequently, the strategy on day 3 would be bounded between [BBA, *AA].

Before embarking on a three-day optimization, it is beneficial to determine whether the third day of the sequence is “challenging”, or in other words, whether any preparation for it is necessary. If a feasible, zero-energy strategy on day 3 can follow strategy A on day 2, then the third day can be ignored. Effort can be concentrated instead on a two-day optimization (assuming it is required).

Assuming a three-day optimization is necessary, then the optimization proceeds as follows:

1. For each of the three days, determine the bounding mode sequences, *i.e.*, find A and B for day 1; find A on day 2 assuming B was used on day 1, and find B on day 2 assuming A was used on day 1; similarly for day 3.
2. There may be few enough scenarios to implement them all and select the three-day mode sequence with the lowest total energy consumption.
3. If this is not the case, or if a more efficient solution is desired, then:
 - a. Assume A is used on day 1.

- b. Perform a two-day optimization (as outlined above) for days 2 and 3 using the final temperature of day 1 as a starting temperature. Determine the optimal mode sequence and total three-day energy consumption. Compare total with previously computed totals. If it exceeds the minimum, terminate. The optimal solution is the minimum of all those previously calculated. If it does not, then proceed to the next step.
- c. Increase the cooling on day 1 by one step (to a maximum of B). Return to Step 3b.

Consider now the situation where it is not possible to find a mode sequence over the three-day period satisfying the deadband constraints. In that case, at least one of the deadband limits must be relaxed. Again, assuming that the lower set point is the more expendable of the two, discard it and proceed with the algorithm sketched above. If it is preferable to introduce soft constraints (*i.e.*, not impose an infinite penalty for constraint violations), then a cost function similar to that shown in Equation [6-2] may be used, perhaps with different weighting on the deviations from the two deadband limits. If this is done, then the straightforward optimization approach outlined in this section would need to be replaced by more flexible optimization tools such as GA's or SA.

As for the one-day optimization, an optimal performance-cost plot similar to the upper plot in Figure 6-5 would be useful to provide some insight into the sensitivity of the optimization to the particular constraints. A plot of optimal total energy consumption over the three-day period versus the M.O.T. on day 3 is most relevant⁶⁴.

It should be clear how the concepts presented in this section could be further applied to longer optimization windows, should reliable forecasts be available and the building contain sufficient thermal storage capacity.

6.7 Practical Implementation of Optimal Control and Extension to the Supervisory Control/Local Control Architecture

An optimal mode sequence remains optimal only as long as the model predictions are accurate and the inputs to the model are accurate. It would be prudent to repeat the optimization process at several intervals throughout the 24-hour period to make adjustments for revised weather forecasts and to remove any zone temperature errors by replacing temperature predictions with actual measurements. For a building located in a climate such as Broadmoor's, all subsequent optimizations should occur during the nighttime, since little, if any, strategy change during the daytime is likely. Given the very modest calculation cost of the optimization (on the order of a few minutes), it would be reasonable to repeat the optimization every hour. The structure of the control strategy is then identical to that in model predictive control (with a fixed, rather than receding, horizon).

⁶⁴ Again, we assume that the third day requires advance preparation. If it does not, and a zero-energy strategy for the day is feasible (following strategy A on day 2), then day 3 should be removed from the optimization, and a two-day optimization should be performed instead. For a two-day optimization, plot the optimal total energy consumption over the two-day period versus the M.O.T. on day 2.

If the building is located where outside temperatures are generally cooler outside than inside (and the magnitude of internal loads warrants heat removal), then a much more flexible approach to thermal control is possible. The recommendations of this section have not yet been implemented, but rather are provided as guidance for further research.

An optimal mode sequence for the 24-hour period from 5pm – 5pm should be developed, using hourly or greater time steps. In other words, each mode is employed in time increments of at least an hour. As before, it may be beneficial to lump the nighttime period into several large segments, and to lump any periods of high outside temperature into larger segments as well.

As an example, assume that the hours between 3pm and 7pm are unsuitable for ventilation cooling due to high outside temperatures. Assume that the building is closed during those times (Mode 4). The remainder of the 24 hours may be used for cooling. The nighttime (unoccupied) hours are divided into two segments as before with the time of the mode switch being variable, as before⁶⁵. The cooling capacities of the nighttime mode sequences are ranked and assigned code numbers. The proposed optimization technique is iterative, with two principal components: the unoccupied hours and the occupied hours.

1. Establish for the nighttime hours the minimum code number (most aggressive cooling) that provides an acceptable temperature at the start of the occupied period. The sequence associated with this code number is designated B. (It is assumed that some fan operation is required to achieve this.)
2. Given the temperature at the start of occupied hours determined by B, select for each hour of the day the passive (zero-energy) mode that keeps the inside temperature closest to (but not below) the minimum extreme of the deadband.
3. If the upper deadband limit is not exceeded during occupied hours, reduce the aggressiveness of the nighttime cooling until either: no energy is used during the day or night; or additional cooling capacity is required during the daytime to maintain temperatures below the upper deadband limit. If the first case is true, then the optimization is complete. If the second case is true, proceed to Step 4.
4. Given the current nighttime mode sequence, determine the hours of the occupied period when Mode 1 (for fan use) is capable of lowering the inside temperature closer to the deadband minimum than any other passive modes (without violating the limit). Rank all such hours in order of increasing outside temperature⁶⁶.
5. Repeatedly change the mode appearing first on the ranked list to Mode 1 until the M.O.T. does not exceed the maximum limit. (Each time a single hour's mode is changed, the list must be reformulated, since, for example, even if the third and fourth occupied hours appear in first and second place on the list, use of Mode 1 during the third occupied hour may allow a passive mode to maintain the temperature at the deadband minimum in hour four.)

⁶⁵ A warming hour or hours may be used for the hour(s) preceding occupancy. The optimization should be repeated at least for the cases of zero and one warming hours.

⁶⁶ It is assumed here that the outside temperature rises monotonically throughout the occupied period.

6. Store the strategies and associated energy consumption for the entire 24-hour period. If the energy consumption exceeds the minimum obtained so far, select the minimum-energy strategy and terminate. Otherwise, proceed to Step 7.
7. Reduce the aggressiveness of the nighttime cooling (increase the code number by 1) and return to Step 4.

The resulting optimal trajectory could be applied directly in model-predictive-control fashion. However, it may be desirable to exert a tighter control on the inside temperature than can be achieved using hourly mode adjustments. In this case, consider the following approach.

The optimal temperature trajectory established above serves as the set point for a local controller. In other words, the algorithm above plays the role of an optimal supervisory controller. The receding optimization horizon for the local controller extends one hour into the future, with the hour divided into 10-minute segments, for example. The purpose of the local controller is to devise a strategy over the optimization window such that: the space temperature never exceeds the set point established by the supervisory controller, the space temperature never drops below the minimum deadband limit, and a minimal amount of energy is expended. The first two goals are imposed as hard constraints.

As before, it is possible to simplify the optimization procedure by recognizing that it is always preferable to cool more aggressively at earlier times than at later times (assuming that aggressive cooling at earlier times does not violate any constraints and that the outside temperature rises monotonically throughout the occupied period). Additionally, it was found above to be very fruitful to divide the nighttime hours into two segments with a variable transition time between the segments. Consequently, the same approach is advised for the local control optimization. One may then simply rank the cooling capacity of mode sequences *for the local control time horizon*. (It may also be desirable to limit actuator wear by imposing a penalty for mode sequences that involve a mode switch at time t and at a later time in the local control horizon.)

In essence, the local control algorithm is identical to the one-day optimization algorithm provided earlier, with several exceptions:

1. The upper deadband limit is replaced by the set point from the supervisory controller.
2. The horizon is 60 minutes, divided in increments of 10 minutes (for example).

To re-emphasize a point made during the discussion of the one-day optimization, it is always advantageous to use the most aggressive *passive* cooling to maintain temperatures as close as possible to the deadband minimum. If passive cooling is not possible throughout the local control horizon, then the warmest feasible solution should be selected.

In order to incorporate as much measured information as possible, it is recommended that the optimization be repeated every 10 minutes, implementing at each time the mode setting for the first portion of the optimal mode sequence. Such an approach is consistent with model predictive control.

6.8 Incorporation of More Complex Mechanical Cooling

In this chapter, the focus has been on cooling a building using a simple mixed-mode system: passive cooling via natural ventilation and mechanical cooling using a single-speed fan. Had there been sufficient data to reliably model the operation of a variable-speed fan, there is no reason that the methods presented in this chapter could not have been applied. The approach would have been identical. The only change to consider would be the incorporation of additional modes such as Mode 1a, Mode 1b, *etc.*, each devoted to a different fan speed. It would remain a straightforward matter to rank the cooling capacity of each mode under certain conditions, and, consequently, to generate a ranking of mode sequence cooling capacity⁶⁷.

Note that the same is true if one considers a building outfitted with mechanical air conditioning⁶⁸. In that case, an additional mode would be created for the air conditioning. As for the variable-speed fan, if variable cooling capacity is available, the air conditioner mode may be subdivided into sub-modes, each for a different capacity. The cooling capacity of each sub-mode could be ranked in the same way as those for other modes.

The inclusion of mechanical air conditioning provides the possibility of maintaining comfortable conditions for the duration of the occupied period. Therefore, the supervisory control/local control approach proposed above should be applied when air conditioning is available.

6.9 Summary

An approach has been presented in this chapter that draws upon physical insight to greatly simplify and accelerate the process of optimizing the scheduling of control modes to maintain comfortable conditions with minimal cost. The effectiveness of the algorithm was demonstrated with control strategy optimization for a single day. Also discussed was the value of generating and considering the Pareto front for the purpose of assessing the sensitivity of the optimization results to the particular constraints imposed. The concepts employed in the single-day optimization were applied to develop algorithms for multi-day optimizations. It was shown how these concepts could also be extended to develop a more traditional supervisory-control/local-control architecture in cases where outside temperatures permit ventilation cooling during occupied hours. Finally, the incorporation of more complex mechanical cooling was considered.

⁶⁷ With a variable-speed fan, it is conceivable that the mechanically induced cooling rates could be lower than those attained using natural ventilation. Such low-speed operation would be eliminated as an option for any such conditions.

⁶⁸ It should be noted that initial start-up transients may need to be accounted for in such a system.

7 Modeling of Houghton Hall

7.1 Introduction

The success of the modeling approach developed for Broadmoor provided motivation for its application to another building. A colleague, Christine Walker, was studying a mixed-mode office building in the U.K., and graciously offered to provide measurement data from the building. Walker *et al.* (2004) published another study of this building. It was known from the outset that much information about the building was not available, such as the opening schedules of its windows. It was hoped that it would be possible to proceed despite this shortcoming. As we will see, some good progress was made, but ultimately, more measurements would be required to create a high-performance model of the space capable of informing an optimal control strategy.

The goals for the work in this chapter were to assess the model quality achievable with the given information, and to determine if it could be used for optimal control, particularly of the exhaust fans in the building.

7.2 Building Description

Houghton Hall is located in Houghton Regis, northwest of Luton, U.K. (0.37°W , 51.87°N). The building, whose long axis runs SW-NE, has three levels. For simplicity, the NW side will be referred to as the North side, and the SW side as the South side. The building is shown below in Figure 7-1.



Figure 7-1 Houghton Hall, view of the South side. (Photo: Vincent and Gorbing)

Four of the building's office zones are conditioned with NV and MV: South Ground Floor, South and North First Floor, and North Second Floor. The remainder of the building is used for storage or for air-conditioned office and meeting space. It is isolated from the portion of the building under investigation. The building's central atrium was divided into zones for the purpose of this project: Atrium Ground Floor, Atrium First Floor, and Atrium Second Floor. See Figure 7-2 for a cross-sectional view of the building showing the seven zones.

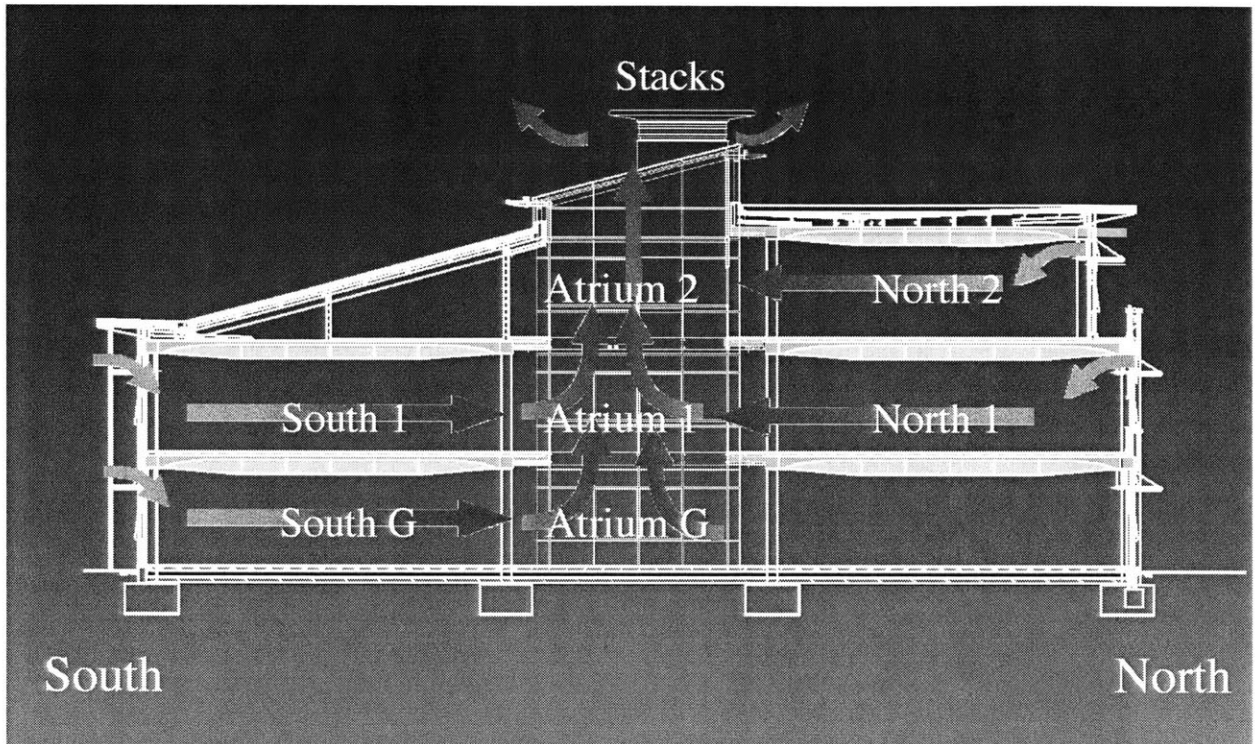


Figure 7-2 Cross-sectional view of Houghton Hall, showing intended airflow paths for buoyancy-driven ventilation. (Section: Vincent and Gorbng)

A set of five fans⁶⁹ was installed to boost airflow out of the five stack vents above the Atrium. Each stack was 3 m by 3 m, and extended approximately 1.6 m above the roof. The motor-driven louvers on the sides of all stacks operated in concert: all East louvers operated together, all West together, *etc.* These were adjusted based on current wind, rain and temperature conditions. The louvers were approximately 14.5m above the atrium floor. The constant-speed fans also operated as a group. Their operation was programmed based on the inside temperature at the elevation of the fans, but was manually overridden during the period of the study. Except for cases of high wind and/or rain, the louvers were programmed to open when the stack temperature exceeded 26°C, and the fans were programmed to turn on when it exceeded 27°C (regardless of outside temperature).

⁶⁹ The fans were not typical exhaust fans, but rather resembled ceiling fans. Measurements taken on site revealed that some air in the vicinity of the fans flowed downwards when the fans were on.

The office areas, totaling approximately 2600 m², were carpeted, with exposed concrete ceiling coffers of average thickness equal to 400 mm. The ceiling height was 3.6 m. The façades were approximately 45% glazed, with the remainder being masonry cavity wall. As shown in Figure 7-2, the office areas were all open to the atrium zones. On the North and South façades of the building were two sets of windows: small awning windows at ceiling level (1.45 m x 0.58 m) that remained open during the cooling season and awning windows (1.45 m x 1.15 m) that were adjusted by the occupants. (Some doors could also be opened to the exterior.) For security reasons, all but the small awning windows were shut when the building was unoccupied.

The occupied hours were approximately 8:30 – 19:00, Monday through Friday.

7.3 Building Measurements

The building was extensively monitored as part of Walker’s research (2004). Included in this section are the measurements relevant to the current study. These include temperature measurements inside and outside the building, solar radiation on the horizontal plane, wind speed and electrical usage.

7.3.1 Temperature Measurements

HOBO[®] temperature loggers were placed throughout the building at desk height. The loggers recorded average temperatures at 15-minute intervals. Temperatures in the four office areas were recorded by six loggers: three in a row near the façade, and three in a row adjacent to the atrium. Two were placed in the ground-floor atrium, four around the perimeter of the first-floor atrium, and four around the perimeter of the second-floor atrium.

The temperature sensors were grouped together according to the zone in which they were placed. Due to their placement, a number of the sensors received direct solar gains. To compensate for this problem, the sensor data were combined as follows:

Zone	Sensors Used	Notes
Atrium Ground	avg(129,130)	all sensors
South Ground	avg(131-136)	all sensors
North 1	avg(149-151)	sensors adjacent to façade only
South 1	avg(137-143)	all sensors
Atrium 1	min(144,145,636426,640788)	minimum of all sensors
North 2	avg(157-159)	sensors adjacent to façade only
Atrium 2	min(152,153, 636777,642439)	minimum of all sensors

Table 7-1 Sensor grouping to form zone temperatures. In the North 1 and North 2 zones, the sensor rows adjacent to the Atrium were not included in the average zone temperature since they received direct solar heating during the day. For Atrium 1 and 2, it was necessary to use the minimum of the four temperatures to select the sensor not affected by solar radiation at a given time.

The average outside temperature was recorded by a Campbell Scientifics, Inc. 10X data logger, also at 15-minute intervals.

Zone temperatures over a two-day period are provided in the following figures.




Figure 7-3 Atrium and outside temperatures, measured in °C.




Figure 7-4 Office and outside temperatures, measured in °C.

7.3.2 Solar Radiation

A LI-COR pyranometer was used to record average solar radiation incident on the horizontal plane for every 15-minute period. It was found that, for some zones, model performance could be improved by making use of solar radiation incident on the four faces of the building in addition to that on the horizontal plane. The procedure for converting the horizontal measurements to directional data is outlined in this section.

Guidance in the literature is widely available for the determination of angles of incidence of direct solar radiation on a surface of a given orientation at a given time (e.g., McQuiston and Parker 1994; ASHRAE 1997). The challenging part of the task encountered here is to extract the fraction of the total radiation striking the horizontal surface that is due to diffuse radiation. Duffie and Beckman (1991) provided a procedure for doing so. It is summarized here⁷⁰. (Another approach is provided by Orgill and Hollands (1997).)

A clearness index, k_T , is defined:

$$k_T = \frac{I}{I_0} = \frac{\text{Total solar radiation on a horizontal surface}}{\text{Total solar radiation on an extraterrestrial horizontal surface}} \quad [7-1]$$

The clearness index was calculated for each measurement given. The variable I_0 was calculated via:

$$I_0 = I_{SC} \left(1 + 0.033 \cos \left(\frac{360n}{365} \right) \right) \cos(\theta_z) \quad [7-2]$$

The variables I_{SC} and θ_z refer to the solar constant (1367 W/m²) and the solar zenith angle (readily calculated using the sources cited above). The variable n refers to the day of the year.

A correlation developed by Erbs *et al.* (1982) was provided to estimate the ratio of diffuse to total radiation incident on a horizontal surface:

$$\frac{I_{diffuse}}{I} = \left\{ \begin{array}{ll} 1.0 - 0.09k_T & \text{for } k_T \leq 0.22 \\ 0.9511 - 0.1604k_T + 4.388k_T^2 & \text{for } 0.22 < k_T \leq 0.80 \\ 0.165 & \text{for } k_T > 0.80 \end{array} \right\} \quad [7-3]$$

Once the amount of diffuse radiation striking the horizontal ($I_{diffuse}$) is calculated, the direct radiation striking the horizontal is calculated by subtraction: $I - I_{diffuse}$. The direct normal radiation, I_{DN} , is then calculated via:

⁷⁰ A spreadsheet automating the calculations is "Solar Calcs Luton.xls".

$$I_{DN} = \left\{ \begin{array}{ll} \frac{I_{direct\ on\ horizontal}}{\cos(\theta_z)} & \text{if } \cos(\theta_z) > 0.1 \\ I_{direct\ on\ horizontal} & \text{otherwise} \end{array} \right\} \quad [7-4]$$

The second line of Equation [7-4] is a compromise. For solar zenith angles close to 90°, the calculated direct normal radiation becomes very large. The second line was added to the calculation to prevent this from happening. The radiation predicted in the second case is a lower bound of the true radiation. One could argue that a different choice could be made.

Once the diffuse radiation on the horizontal and the direct normal radiation are available, then it is straightforward to calculate the total radiation (direct, diffuse and reflected) incident on any building surface by drawing on the cited source material.

One other issue to consider is the time at which one calculates the solar angles: the end of the sample period or, perhaps, at the midpoint. More appropriate would be the time-averaged solar radiation on each surface over the entire sample period. In this case, the solar angles would have to be time-averaged over the period. This issue becomes particularly relevant with longer sample times, such as once per hour, but it also may present problems around sunrise or sunset (*e.g.*, if one tries to calculate solar angles at the end of a sample period that happens to end just after sunset).

An illustration of the output is provided in Figure 7-5, below.




Figure 7-5 Calculated solar radiation incident on the NW, NE, SE and SW faces (actual directions) of Houghton Hall on August 27, 2003. The “Horizontal” line shows measured values.

7.3.3 Wind Speed

The wind speed was measured using an anemometer. No wind direction measurements were made. As can be seen below in Figure 7-7, wind was present for the bulk of the monitoring period.

7.3.4 Electrical Usage

Electrical usage was monitored using a Highland Technology/Enernet Corporation K20 power logger. The electrical usage profile was relatively steady. The combination of this observation and the fact that some data were missing led to the adoption of average daily profiles for the building's three levels. Average Monday, Tuesday, *etc.* profiles were generated by averaging the available data at every sample time for each day of the week. The resulting profiles are shown below in Figure 7-6.




Figure 7-6 Electricity consumption profiles for the three building levels.

The stack fan circuit was also monitored. Unfortunately, it was not possible to monitor the stack louver actuators. No record exists for when the louvers were opened or closed. Similarly, there is no complete schedule of other aperture openings and closings.

7.3.5 Data Summary

Measurement data on a 15-minute time interval was available for the first 19 quantities listed in Table 7-2, below. A total of 35 features was used for modeling, including terms such as Fan*temperatures and (Wind speed)*temperatures. For simplicity, it was decided to not employ terms such as Fan ΔT . In the case of the wind, it was not known which ΔT terms would be

appropriate. Also included in the table were the maximum number of lag terms permitted for each feature type during the feature optimization process.

ID	Feature	Number of Lag Terms	ID	Feature	Number of Lag Terms
1	Atrium Ground	5	18	Bldg. Occupied	0
2	South Ground	5	19	Fan	0
3	North 1	5	20	Fan*AG	1
4	South 1	5	21	Fan*SG	1
5	Atrium 1	5	22	Fan*N1	1
6	North 2	5	23	Fan*S1	1
7	Atrium 2	5	24	Fan*A1	1
8	Outside	5	25	Fan*N2	1
9	Solar_Horiz	5	26	Fan*A2	1
10	Solar_N	5	27	Fan*Tout	1
11	Solar_E	5	28	Wind*AG	5
12	Solar_S	5	29	Wind*SG	5
13	Solar_W	5	30	Wind*N1	5
14	Wind Speed	5	31	Wind*S1	5
15	Elec. Ground	5	32	Wind*A1	5
16	Elec. 1	5	33	Wind*N2	5
17	Elec. 2	5	34	Wind*A2	5
			35	Wind*Tout	5

Table 7-2 Summary of model features available. Included are the maximum numbers of lag terms that may be selected during the feature optimization process. “Bldg. Occupied” refers to whether the building was in use at a particular time. The Fan*temperature features (20-27) were allotted only a single lag term. Doing so avoided unstable predictions at mode transitions.

The data available span two time periods, for a total of 3911 data points:

1. 14:15, August 14, 2003 – 10:15, September 10, 2003
2. 09:15, September 21 – 06:15, October 5, 2003

It was found to be beneficial to use the little information available about the operation of the building to divide the data set into segments reflecting distinct operational modes. Features 18 and 19 were used to create four modes. The building usage feature was selected to serve as a surrogate measure of whether the windows were open or not. We were told that no windows other than the small awning windows were open during unoccupied periods. During occupied periods, the occupant-controlled windows may or may not have been open. The four modes were labeled as follows:

Mode Number	Description	Number of Points
3	Fan ON, Occupied	729
4	Fan ON, Unoccupied	1654
5	Fan OFF, Occupied	474
6	Fan OFF, Unoccupied	1054

Table 7-3 Listing of operational modes.

The following figure provides an overview of the outside temperature, the solar radiation incident on the horizontal plane, the wind speed and the operational modes for the duration of the building monitoring period.

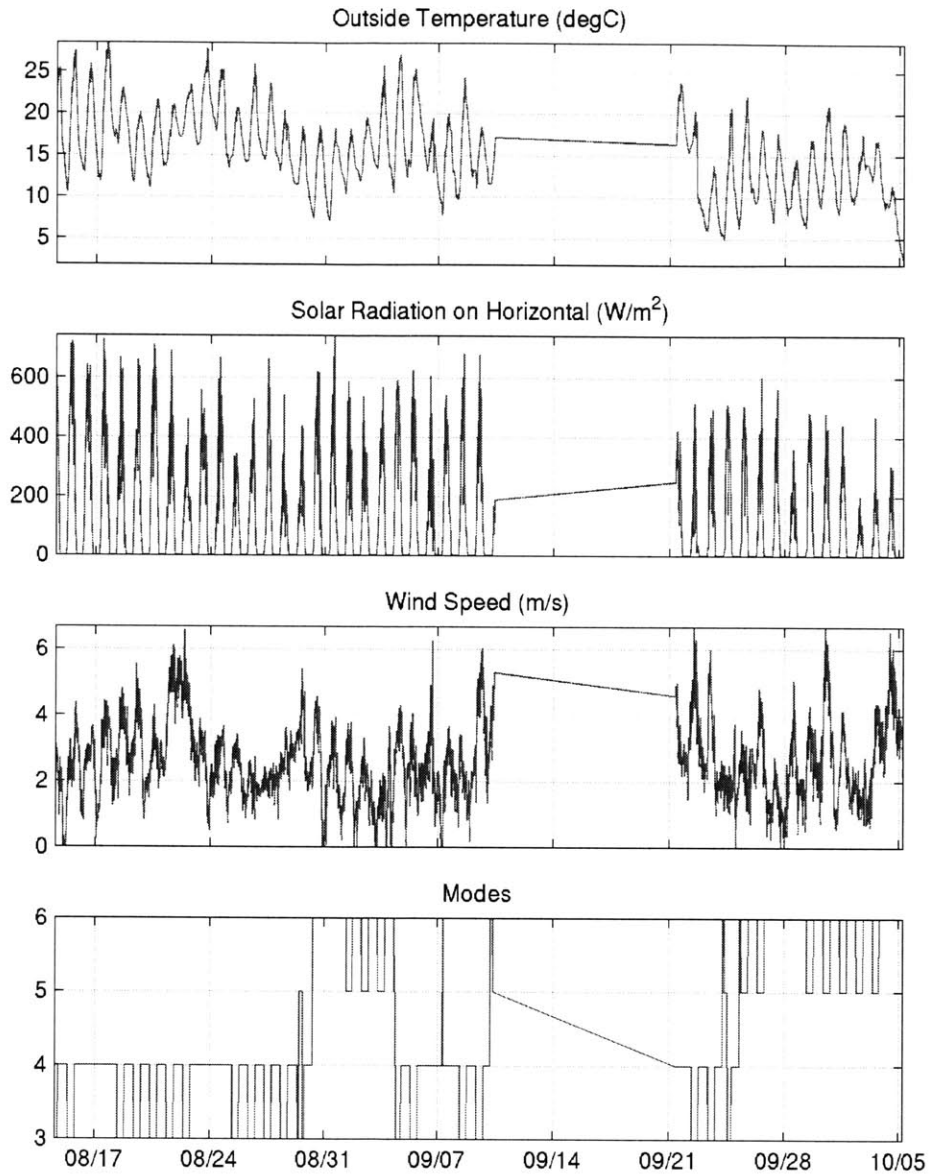


Figure 7-7 Outside temperature, solar radiation on the horizontal plane, wind speed and operational modes for the entire monitoring period.

It should be noted that the composition of the points assigned to each mode does not reflect the full range of temperatures experienced over the monitoring period. Consider Figure 7-8, below.




Figure 7-8 Histograms of the outside temperature for the four operational modes. The ordinates represent frequency of occurrence of each temperature. (TouHistograms.m)

Comparison of the histograms for each mode reveals that very few data were available with warm temperatures and NO fan. Also apparent from the second plot is that the fan was operated down to temperatures as low as 5°C. As mentioned earlier, the fan was operated manually—not according to its programmed schedule.

7.4 Model Development

In this section, models for the thermal behavior of Houghton Hall will be developed. The same framework used for Broadmoor will be employed.

7.4.1 PHDRT

The ground-floor and the second-floor atria were investigated as test zones for the PHDRT algorithm. The features used were the elementary features shown in Table 7-2: features 1 through 19. The number of lags for each feature was set to two. It was found in both cases that segmenting the data using the algorithm did not improve the one-step-ahead rms error on the test set. (In this case, the training set and test set were interleaved, with the odd points used for training and the even for testing.) As for Broadmoor, the algorithm was set aside, and the modeling proceeded with the different forms of linear regression discussed above.

7.4.2 Simple Model

In this section, a simple model was created for each zone using all the data (no modes were used). The features listed in Table 7-4 were used in addition to a constant term serving as the

intercept. For simplicity, all features types were allotted 5 lag terms, with the exception of 18 and 19.

ID	Feature	Number of Lag Terms	ID	Feature	Number of Lag Terms
1	Atrium Ground	5	18	Bldg. Occupied	2
2	South Ground	5	19	Fan	2
3	North 1	5	20	Fan*AG	5
4	South 1	5	21	Fan*SG	5
5	Atrium 1	5	22	Fan*N1	5
6	North 2	5	23	Fan*S1	5
7	Atrium 2	5	24	Fan*A1	5
8	Outside	5	25	Fan*N2	5
9	Solar_Horiz	5	26	Fan*A2	5
10	Solar_N	5	27	Fan*Tout	5
11	Solar_E	5	28	Wind*AG	5
12	Solar_S	5	29	Wind*SG	5
13	Solar_W	5	30	Wind*N1	5
14	Wind Speed	5	31	Wind*S1	5
15	Elec. Ground	5	32	Wind*A1	5
16	Elec. 1	5	33	Wind*N2	5
17	Elec. 2	5	34	Wind*A2	5
			35	Wind*Tout	5

Table 7-4 Model features used for simple modeling. Also included are the numbers of lag terms used.

The results of the modeling are shown in tabular form in Table 7-5, below.

Zone	One-step ahead rms error		Pure simulation rms error
	Train Set	Test Set	Complete Set
Atrium Ground	0.10	0.11	0.61
South Ground	0.07	0.08	0.63
North 1	0.08	0.09	0.72
South 1	0.07	0.07	0.68
Atrium 1	0.17	0.19	0.71
North 2	0.11	0.12	0.76
Atrium 2	0.20	0.22	0.67

Table 7-5 Results of one-step-ahead and pure simulations using a simple model for each zone (no modes). (MainN0modes.m)

These results will be compared with those obtained in the following sections. Plots of predicted zone temperatures will be shown in later sections, as well.

7.5 Model Using Modes, but no Feature Optimization

Several attempts were made to develop a model using modes, but with no optimization of the features used as inputs to the models. The features shown in Table 7-2 were used. Also tried were those same features, but with every instance of five lag terms replaced by two lag terms.

Performance was poor in pure simulations, with large errors sometimes occurring at mode switches.

The results of one-step-ahead and pure simulations are provided in Table 7-6, below, for the case where two lag terms were used for most features.

Zone	One-step ahead rms error		Pure simulation rms error
	Train Set	Test Set	Complete Set
Atrium Ground	0.10	0.15	0.94
South Ground	0.07	0.11	0.92
North 1	0.08	0.10	1.76
South 1	0.07	0.08	1.18
Atrium 1	0.17	0.30	1.60
North 2	0.11	0.13	2.06
Atrium 2	0.19	0.28	1.40

Table 7-6 Results for each zone of one-step-ahead and pure simulations using a model with modes, but no feature optimization. (MainModesNOoptim.m)

When the features in Table 7-2 were used, the number of parameters in the model was 156, including the constant intercept term. It was thought that the number of parameters to train may have been excessive for the number of data points available in each mode, so the model was created using just two lags per feature type. This change reduced the parameter number, but did not solve the problems sometimes encountered when mode switching occurred. As we will see in the next section, feature optimization, coupled with the used of modes, leads to improved performance relative to that found with the simple model.

7.6 Model Using Modes and Feature Optimization

Models were created for each zone, using the operational modes to segment the data set. The model features were optimized using the greedy feature optimization presented earlier. The maximum number of lags used for each feature type are shown in Table 7-2. The simulation performance is shown below in Table 7-7.

Zone	One-step ahead rms error		Pure simulation rms error	Pure simulation max error
	Train Set	Test Set	Complete Set	Complete Set
Atrium Ground	0.10	0.11	0.55	2.2
South Ground	0.07	0.07	0.53	2.3
North 1	0.08	0.09	0.57	2.1
South 1	0.07	0.07	0.57	1.9
Atrium 1	0.18	0.18	0.61	3.0
North 2	0.11	0.12	0.86	3.8
Atrium 2	0.20	0.21	0.77	3.7

Table 7-7 Results for each zone of one-step-ahead and pure simulations using a model with modes and feature optimization. The maximum magnitude errors are provided to show the size of errors associated with even the best-performing model. (Main.m)

In Figure 7-9, the pure simulation temperature predictions for the seven zones are compared with the measured temperatures. The prediction errors are shown in Figure 7-10.




Figure 7-9 Measured temperatures (black) and model predictions (grey) for the seven building zones. (Main.m)




Figure 7-10 Pure simulation prediction errors for the seven zones; and modes in effect. (Main.m)

A cursory inspection of the simulation results in Figure 7-9, Figure 7-10 and Table 7-7 reveals that the performance of the models for Houghton Hall are not as accurate as those for Broadmoor. A noticeable improvement from Table 7-5 to Table 7-7 is found for all but the second-floor zones, demonstrating the importance of both the use of modes and the use of feature optimization.

A possible explanation of the increased error on the second floor is proposed. Consider the North 2 plot for the weekend of September 27 – 28 in Figure 7-9. The period is indicated with an arrow. Typically, the occupant-controlled windows were closed during unoccupied periods. If windows in the second-floor office area were left open for the weekend indicated, the actual temperature profile for North 2 and Atrium 2 (and to a lesser extent, the other zones) would have been lower than predicted⁷¹. Had modes not been employed to create zone models, there would have been no indication that there was potentially unusual behavior on that weekend. Unfortunately, due to the monitoring setup, there is no way to verify the opening schedule of the windows.

Additional examination of the temperature predictions in Figure 7-9 reveals that the peaks and troughs of the diurnal temperature swings are frequently over- or under-predicted. A systematic trend in the over- and under-predictions could not be found. To investigate the issue, the one-step-ahead errors were plotted versus outdoor temperature. No trend was found. Once again, we must resort to the lack of data for an explanation. Recall that neither the louver opening schedule nor the occupant-controlled window opening schedule was available.

Over-prediction of peak temperatures could result from the following scenario. Consider the over-prediction of temperatures on August 26 (Figure 7-9). Reference to Figure 7-7 reveals that the outside temperatures were warmer than the zone temperatures in the ground-floor zones and the first-floor office zones (this, perhaps not coincidentally, is where the over-prediction occurred). If the occupant-controlled windows or the louvers were opened this day less than on average, then the measured inside temperatures would be cooler than expected for average operation. Such a scenario is consistent with the results. The change would not affect the upper zones to the same extent since the inside-outside temperature difference there was relatively small.

Consider, now, the under-prediction of the minimum temperatures during the early-morning hours of September 7. Wind speed is relatively low during this period, so the dominant flow regime was likely buoyancy-driven flow. In this scenario, air rises from the first floors and exits through the stacks and the second-floor small awning windows. Temperatures are under-predicted for all but the second-floor zones. If the stack louvers were closed during this period more than on average (it was cold that night), the model predictions would be expected to be lower than the measurements.

⁷¹ The reader may have noted that the predictions for the daylight hours on Friday were too high, and may have surmised that this error at the start of the weekend spawned the errors observed over the entire weekend. To address this issue, a simulation was begun using correct data from Friday night as the starting point. Given those correct initial conditions, the same over-prediction of temperatures still occurred on Saturday and Sunday.

It is clearly not ideal to make such conjectures. Access to a more complete data set could establish whether the modeling approach or the lack of completeness of the data set was to blame for the model errors. Unfortunately, we do not have the option.

Histograms of the one-step-ahead and pure simulation prediction errors are displayed in Figure 7-11 and Figure 7-12.




Figure 7-11 Histogram of one-step-ahead prediction errors (°C) for each zone. (PlotOutput.m)

One can see in the histograms of the one-step-ahead prediction errors that significant tails exist in the distributions, particularly for the atrium zones. It appears that there may be multiple distributions superimposed in the Atrium 1 and 2 distributions. A moderate peak appears in both in the vicinity of 0.25 – 0.30 °C. The origin of this behavior is not known, but may be linked to an additional mode of operation not accounted for in the models, perhaps related to the operation of the louvers.

The pure simulation histograms reveal a prediction bias of: 0.05, 0.04, 0.01, 0.02, 0.03, 0.10, and 0.13°C for the seven zones. The bias is greatest for the second-story zones.




Figure 7-12 Histogram of pure simulation prediction errors (°C) for each zone.
(PlotOutput.m)

7.7 Application of the Model

Up to this point, the discussion has centered on the ability of the model to recreate the measured temperatures. In this section, the model is used to try to establish the importance of running the fan. The question of the suitability of the model for optimal control will be discussed as well.

The approach taken in this section for assessing the effect of the fan was to compare the predictions of the model with the fan on with the predictions of the model with the fan off. When a pure simulation is performed, the predictions at one step impact those at later steps. For this reason, it is not meaningful to perform two pure simulations—one with the fan on and one with the fan off—and then make comparisons between the two trajectories. Only via one-step-ahead predictions can one make a fair assessment of the fan’s contribution to temperature changes.

Unfortunately, the data available do not provide many points where the fan was off and the outside temperature was relatively high. In other words, if we were to take the set of points when the fan was on, and then turn the fan off (by changing the modes used for predictions), the models for the “fan off” modes would be required to make temperature predictions in a region where they were not trained. (See Figure 7-8.)

The “fan on” modes cover a greater span of outside temperatures, nearly completely encompassing the range encountered by the “fan off” modes. In other words, the models for the

“fan on” modes were trained using nearly the full range of data seen by the “fan off” models. Therefore, if the modes are manually switched from “fan off” to “fan on”, the models will not be required to make predictions in regions where they were not trained.

The following procedure was implemented:

1. A one-step-ahead simulation was performed for the building as it was operated.
2. Another one-step-ahead simulation was performed, but every time Mode 5 was actually used, it was replaced by Mode 3, and every time Mode 6 was encountered, it was replaced by Mode 4. (This has the effect of turning on the fan whenever it was not already on.)
3. All data points for which a mode change was made in Step 2 were identified. (These are the points that can be used to make a comparison between model predictions with and without the fan on.)
4. For the points in Step 3, the prediction difference was formed: (Fan OFF prediction – Fan ON prediction).
5. Only those predictions were kept that occurred at times when the outside temperature was cooler than the zone temperature (this excluded very few points, since it was generally cool when the fan was off).
6. A histogram was made of this difference for each zone.

The histograms are shown below in Figure 7-13.




Figure 7-13 Histograms showing the frequency of occurrence of prediction differences between fan OFF and fan ON operation of the building. Mean and median prediction differences are shown for each zone. (FanONOFF.m)

It is difficult to see in the figure, but some extreme outliers exist in the prediction differences. Atrium 1 has prediction differences lower than -8°C and greater than 6°C . North 2 has differences lower than -8°C and greater than 8°C . These large differences are troubling, as it is difficult to believe that the temperature could change so much in a single sample period. Of greater concern, however, is the fact that all predictions were made at times when the outside temperature was lower than the zone temperature. In other words, even when it is cooler outside than inside North 2, the model predicts that the temperature within the zone could rise by more than 8°C in one sample period. This response is not at all consistent with the expected effect of the fan.

Setting aside the outliers, one may focus on the median difference of the fan OFF – fan ON predictions. Examination of the Atrium 2 zone reveals that the predictions with the fan ON are approximately 0.2°C cooler than those with the fan OFF (at times when the outside temperature is cooler than the inside temperature). This is the expected, or desired, consequence of running the fan. However, one finds that the median effect on the North 2 zone of running the fan is to *increase* the zone temperature by approximately 0.1°C relative to the case when the fan is OFF.

This unexpected result is found in all other zones except the Atrium Ground zone, where the expected result is found.

In light of these results, the role of the fan is clearly ambiguous. It is not possible to ascertain what, if any, effect the fan operation has on the zone temperatures. If tight clusters of prediction differences were found, then one could conclude that the fan had a clear impact on temperatures. However, the wide spread of predictions and the unexpected median impact of running the fan leads one both to question the validity of the model making the predictions and to question whether the fan had a strong impact on the building's thermal response. In recognition of what has been observed, it would not be prudent to attempt to use the building model for predictive control.

7.8 Summary

In this chapter, the modeling framework developed in this thesis was applied to another mixed-mode building: Houghton Hall. Once again, it was found that it was not necessary or beneficial to accommodate any system nonlinearities via the PHDRT algorithm. Improvements in pure simulation performance were demonstrated for models incorporating simple building modes and feature optimization relative to simple models with no modes, or to models using modes, but with no feature optimization. Recall that feature optimization played a minor role at Broadmoor; its role for this building was critical.

It was hoped that a useful model could be developed despite the lack of information regarding the operation of the louvers in the building stacks and the operation of the windows. The model developed did exhibit moderately good agreement with the measured zone temperatures (rms prediction errors of 0.5 to 0.8 in pure simulations). However, one example was provided where an apparent unusual window opening led to significant prediction errors. It is proposed that significant prediction improvement could be realized with a complete data set.

An investigation was made into the impact of the fan on the thermal behavior of the building. The results were inconclusive. To yield a more definitive conclusion, it would be beneficial to augment the data set to include periods when the fan was off at times of high outdoor temperatures and periods when the fan was on at low outdoor temperatures. Of foremost importance in creating models reliable enough to answer the important questions related to fan use in this building is availability of data on the activity of all building apertures⁷². Without such models, model-based optimal control is impossible.

⁷² For systems that are not automated, this information can be obtained by using video cameras or closure sensors.

8 Building Dynamic Modes, Time Constants and Time Shifts

8.1 Introduction

The origin of this chapter lies in some of the questions raised during the analysis of the Houghton Hall data. These include: How important is the thermal mass in the building? Is it engaged in the cooling process? How does the operation of the fan affect the dynamics of the building? To begin to answer these questions, it was necessary to return to some fundamental concepts in dynamics that appear to be overlooked (at least not emphasized) by much of the building design community. The results presented in the chapter should not be taken as conclusive, but rather suggestive of some crucial underlying processes that should be considered in building design and analysis.

What is the importance of the role of thermal mass in a building? In a building, and indeed in all purely thermal systems, thermal mass is the sole energy storage element. There is no analog to the momentum found in mechanical or the inductance found in electrical systems. The only way to introduce dynamic effects is by working with the building's thermal mass, or capacitance. It is the thermal mass of a building that permits the phenomenon of night cooling, whereby the building is cooled during the nighttime hours in preparation for the coming day. If there were no thermal mass in the building, such a strategy would be useless. The extent of the thermal mass, *and* the extent of its excitation are both critical factors in determining the success of a night cooling strategy.

The chapter begins with a discussion of dynamic modes and their relevance in buildings. The remainder of the chapter is devoted to the extraction of time constants from data to ensure that physically reasonable values are obtained.

8.2 Discussion of Dynamic Modes⁷³

As in mechanical and electrical systems, the dynamics of a linear thermal system may be understood in terms of its dynamic modes. To predict the response of a linear system, it is not necessary to do so from the perspective of a modal analysis, but doing so generally provides greater insight than alternative methods.

Consider the system:

$$\begin{aligned} \dot{x} &= Ax + Bu \\ y &= Cx + Du \end{aligned} \quad [8-1]$$

The variables x , u and y are all vectors. A transformation of variables, $Mz = x$, yields:

⁷³ The fundamentals of this discussion are taken from class lecture notes from MIT course 2.151, taught by Professor David Hardt in the Fall semesters of 2000 and 2002. A preliminary book manuscript in preparation at the time by Professor Kamal Youcef-Toumi also served as a reference.

$$\begin{aligned}\dot{z} &= M^{-1}AMz + M^{-1}u \\ y &= CMz + Du\end{aligned}\quad [8-2]$$

When the columns of M , named v , are the eigenvectors of A , then z represents the modal variables of the system. The matrix

$$\Lambda = M^{-1}AM \quad [8-3]$$

is diagonal (provided roots are real and distinct) and the diagonal elements are the eigenvalues of the system. In a thermal system, these eigenvalues are all real and negative, corresponding exponential decays for an initial-condition response. Note that the state variables, x , correspond to a linear combination of the modal variables, so the initial-condition response of a state variable may be comprised of multiple exponential decays.

Consider the initial condition response, where $u = 0$.

$$\begin{aligned}z(t) &= e^{\Lambda t}z(0), \text{ or} \\ x(t) &= Me^{\Lambda t}M^{-1}x(0)\end{aligned}\quad [8-4]$$

The rows of M^{-1} are the left eigenvectors of A . Name the rows w_i^T . Then, for an n^{th} -order system:

$$x(t) = e^{\lambda_1 t}v_1w_1^T x(0) + e^{\lambda_2 t}v_2w_2^T x(0) + \dots + e^{\lambda_n t}v_nw_n^T x(0) \quad [8-5]$$

From Equation [8-5], one can see that the initial-condition response of x is comprised of multiple exponential terms, each weighted by scalars $w_i^T x(0)$ and multiplied by vectors v_i .

The complete response of an n^{th} -order system with m control inputs is given by:

$$x(t) = \sum_{i=1}^n e^{\lambda_i t}v_i[w_i^T x(0)] + \sum_{i=1}^n \sum_{k=1}^m e^{\lambda_i(t-\tau)}v_i[w_i^T b_k]u_k(\tau)d\tau \quad [8-6]$$

Note that the bracketed terms are scalars. The variable b_k represents the k^{th} column of B .

8.2.1 Example of Modal Analysis

Consider the simple room shown below in Figure 8-1, exposed on all six sides to ambient temperature and solar radiation⁷⁴.

⁷⁴ In reality, of course, all six sides could not possibly be exposed to identical ambient conditions, but, for illustration purposes, this model is sufficient.




Figure 8-1 Simple room. Left: Dimensions of room. Right: Envelope cross section, showing the temperature nodes, the inside and outside convection coefficients, solar absorptivity, α , and the thickness of the envelope. Mechanical cooling, Q_{cool} is provided directly to the inside air temperature node. Ventilation using outside air also directly affects the inside air temperature. Half of the envelope thickness is assigned to each mass temperature node. Concrete properties were used for the envelope, and $\alpha = 0.26$.

The dynamic equations for this example may be written:

$$\begin{aligned}
 \frac{d}{dt} \begin{bmatrix} T_{in} \\ T_{mass1} \\ T_{mass2} \end{bmatrix} &= \begin{bmatrix} \frac{-h_{in}A - \dot{m}c_p}{C_{air}} & \frac{h_{in}A}{C_{air}} & 0 \\ \frac{h_{in}A}{C_{mass}/2} & \frac{-h_{in}A - 1/R}{C_{mass}/2} & \frac{1/R}{C_{mass}/2} \\ 0 & \frac{1/R}{C_{mass}/2} & \frac{-h_{out}A - 1/R}{C_{mass}/2} \end{bmatrix} \begin{bmatrix} T_{in} \\ T_{mass1} \\ T_{mass2} \end{bmatrix} + \\
 &+ \begin{bmatrix} -1 & 0 & \frac{\dot{m}c_p}{C_{air}} \\ C_{air} & 0 & 0 \\ 0 & 0 & \alpha \end{bmatrix} \begin{bmatrix} Q_{cool} \\ I_{sun} \\ T_{out} \end{bmatrix} \quad [8-7] \\
 y &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} T_{in} \\ T_{mass1} \\ T_{mass2} \end{bmatrix}
 \end{aligned}$$

The variable, R , is defined to be $kA/thickness$, where k is the thermal conductivity of the concrete. A is the surface area of the envelope. C_{air} and C_{mass} are the heat capacities of the inside air and the entire envelope, respectively. The mass flow rate of ventilation air is given by \dot{m} , and the specific heat by c_p . Note that the system's eigenvalues and eigenvectors are functions of the mass flow rate of ventilation air. This is a key feature to note for systems using air exchange to modify temperatures (be it via a ducted, mechanical air-conditioning system or via natural ventilation).

In the example given, the system time constants are given for the cases where the ventilation rate is 0, 10, 20 and 40 air changes per hour. (Time constants are equal to $(-1)/\text{eigenvalues}$.)

Air Change Rate (ACH)	Tau 1	Eigenvector 1	Tau 2	Eigenvector 2	Tau 3	Eigenvector 3
0	3.1	0.006	518.4	43.3	84.1	-5.5
		0.0		43.0		-5.3
		0.0		13.2		17.3
10	2.1	8.5	419.1	24.6	83.8	-3.6
		0.0		37.2		-5.3
		0.0		12.0		16.4
20	1.5	12.7	383.4	17.2	83.6	-0.2
		0.0		34.9		-0.3
		0.0		11.6		1.0
40	1.0	16.9	353.9	10.8	83.4	-2.7
		0.0		33.2		-5.3
		0.0		11.3		16.0

Table 8-1 Time constants (Taus) and associated eigenvectors (weighted by $w_i^T x(0)$) for different ventilation rates. (Time constants are displayed in minutes.) The weighting originates from Equation [8-5], with $x(0) = [25;25;25]$. (TimeConstants.m)

The weighted eigenvectors may be interpreted as follows (see Equation [8-5]). Consider an initial-condition response for the case with zero ventilation. In that case, the inside temperature is seen to be composed primarily of modes 2 and 3 (note the 43.3 and -5.5 in the first entry of eigenvectors 2 and 3) and, to a much lesser extent, mode 1 (0.006 in the first entry of eigenvector 1). Note that as the ventilation rate increases, the relative contribution of mode 1 to the inside temperature increases.

Note, also, that the time constants associated with the three modes (especially the first two) decrease as the ventilation rate increases. This decrease is the result of the \dot{m} contribution to the A matrix.

To provide a more visual portrait of the impact of the ventilation rate, consider the following three figures, in which the three temperatures are shown for an initial-condition response simulation, as well as the three contributing modes. It is clearly evident in the figures that the state-variable trajectories do not follow pure exponential decays, but rather are combinations of the exponential decays associated with the three modes.




Figure 8-2 Initial-condition responses for the case with no ventilation. Left: State variable trajectories are shown. Right: Pure exponential decays for the system's three modes ($z_1(t)$, $z_2(t)$ and $z_3(t)$) are shown. (TimeConstants.m)




Figure 8-3 Initial-condition responses for the case with ventilation at 20 ACH. Left: State variable trajectories are shown. Right: Pure exponential decays for the system's three modes ($z_1(t)$, $z_2(t)$ and $z_3(t)$) are shown. (TimeConstants.m)




Figure 8-4 Initial-condition responses for the case with ventilation at 20 ACH. Left: State variable trajectories are shown. Right: Pure exponential decays for the system’s three modes ($z_1(t)$, $z_2(t)$ and $z_3(t)$) are shown. (TimeConstants.m)

Up to this point, the relationship between the ventilation rate and the system time constants has been emphasized and illustrated in several different ways. The focus turns now to the case where the system is driven by the various inputs. The important concept of the discussion that follows is that three elements determine the forced portion of the thermal response of the system: the weighting factor $[w_i^T b_k]$ shown in Equation [8-6], the eigenvectors v_i , and the control inputs themselves, u_k ⁷⁵.

The key element of the three is the weighting factor, $[w_i^T b_k]$. If this factor is zero, mode i cannot be influenced by control input k . If this situation is true for all k , then the mode is considered uncontrollable. In other words, no matter what inputs are given to the system, the response of a particular mode cannot be changed. For a (contrived) thermal example, consider the room above, but with large, perfectly insulated water storage tank in the middle. The time constant associated with the introduced state variable T_{water} (which is also now the fourth mode variable, z_4) is infinite. Calculation of $[w_i^T b_k]$ for $i = 4$ and $k = \{1 \dots 3\}$ yields zero in all cases. No control input can change the temperature of the water in the tank.

This previous example is not practically important—perfect insulation is not an option. However, if the problem description is slightly altered so that the tank has high levels of insulation, then a new mode is introduced to the system with a very long, finite time constant. The important point here is that unless that mode is actively engaged by the control inputs (*i.e.*,

⁷⁵ For the purposes of the discussion, both traditional control inputs and exogenous inputs are considered control inputs.

at least one of the weighting factors $[w_i^T b_k]$ for $i = 4$ and $k = \{1 \dots 3\}$ is of the same order of magnitude as for the other modes), then the tank will play a minimal role in the overall dynamics of the room.

To return to one of the original questions posed at the start of the chapter, unless the modes associated with the building's thermal mass are more than weakly controllable, the role of the mass will be negligible. Practically, what this implies for the case of night cooling a building is this: the incoming cool outside air must come into direct contact with the thermal mass to excite the associated modes. If the flow path of the incoming air doesn't coincide with the location of the thermal mass, then it won't be cooled. A building may have modes with very large time constants, but unless they are appreciably excited, their presence is irrelevant.

Before leaving this discussion of dynamic modes, it is helpful to put the topic in the context of the preceding chapters of the thesis. It was shown, particularly in the case of Broadmoor, that the dynamics of the system were most accurately captured when it was recognized that the building exhibited a set of individual operational, or *control* modes. Each *control* mode was described by its own model, with, conceivably, distinct time constants associated with its associated *dynamic* modes. At its essence, the process of determining an optimal night-cooling control strategy was one of selecting the lowest-energy method of "charging" the thermal mass fastest. The control modes that most increased air velocities and flow rates "charged" the mass the fastest for two reasons: these methods generated the greatest temperature difference between air and mass temperatures, and the greater flow rates also reduced the surface resistance of the mass, thereby reducing the associated time constants.

During the daytime at Broadmoor, the principal aim was to preserve the cool temperatures for as long as possible. When the ingress of warm outdoor air was limited by shutting apertures during the day, the result was to increase the thermal resistance between the inside and outside temperatures. Since, for a first-order system, the system time constant is RC (thermal resistance multiplied by the heat capacity), large thermal resistances lead to a slower loss of the cool temperatures attained through the night-cooling strategy. For higher-order systems, the same qualitative effect is achieved by increasing the thermal resistance between inside and outside temperatures.

It should be recognized that the variable flow rates due to natural ventilation necessarily lead to variable time "constants". Changing flow velocity over a surface changes the convection heat transfer coefficient and consequently the time constant. This nonlinearity has been explored by Yam *et al.* (2003). In a system in which MV produces dominant flows, time constants become a function of which apertures are open.

8.3 Extraction of Time Constants from Data

8.3.1 Selection of Model Type

To identify the time constants associated with the various dynamic modes of the building, the appropriate model structure must be selected. Recall that the model developed in Chapter 5, Equation [5-1], had the form:

$$\begin{aligned}
T_1(t) &= a_{11}T_1(t-1) + a_{12}T_1(t-2) + \dots \\
&\quad + a_{21}T_2(t-1) + a_{22}T_2(t-2) + \dots \\
&\quad \vdots \\
&\quad + a_{n1}T_n(t-1) + a_{n2}T_n(t-2) + \dots \\
&\quad + b_{11}I_{sun,horiz}(t-1) + b_{12}I_{sun,horiz}(t-2) + \dots \\
&\quad + b_{21}I_{sun,south}(t-1) + b_{22}I_{sun,south}(t-2) + \dots \quad [8-8] \\
&\quad \vdots \\
&\quad + c_{11}Fan(t-1)(T_{out}(t-1) - T_1(t-1)) + \dots \\
T_2(t) &= \text{similar} \\
&\quad \vdots \\
T_n(t) &= \text{similar}
\end{aligned}$$

The first subscripts refer to the zone variable (1 to n zones) and the second to the coefficient number (1 to p lag terms). Note the assumption that an input at time t cannot affect the output at time t, but rather only at time t+1 and later.

Since the concept of dynamic modes is inherently linear, the nonlinear features with c coefficients must not be included in the model. One might be inclined to use least squares to identify all the a and b parameters for this model (for various numbers of time lags) and identify the time constants from the characteristic equations associated with each zone:

Take the z transform of difference equation:

$$T_1(z)(1 - a_{11}z^{-1} - a_{12}z^{-2}) = T_2(z)(-a_{21}z^{-1} - a_{22}z^{-2}) + \dots \quad [8-9]$$

Identify the characteristic equation:

$$(1 - a_{11}z^{-1} - a_{12}z^{-2}) = 0 \quad [8-10]$$

Factor:

$$(1 - \alpha_{11}z^{-1})(1 - \alpha_{12}z^{-1}) \quad [8-11]$$

However, this approach works only if the zones are independent. Any cross terms linking T_1 to T_2 , etc., affect the roots of the entire system's characteristic equation so that α_{11} and α_{12} are not directly related to the time constants of the system's modes.

Two equivalent (and legitimate) approaches may be undertaken to establish the time constants of the system modes.

The Equations [8-8] can be rewritten in matrix form:

$$\begin{bmatrix} \underline{T}(t) \\ \underline{T}(t-1) \\ \underline{T}(t-2) \end{bmatrix} = \underbrace{\begin{bmatrix} A_{11} & A_{12} & A_{13} \\ I_n & 0 & 0 \\ 0 & I_n & 0 \end{bmatrix}}_A \begin{bmatrix} \underline{T}(t-1) \\ \underline{T}(t-2) \\ \underline{T}(t-3) \end{bmatrix} + \underbrace{\begin{bmatrix} B_{11} & B_{12} \\ 0 & 0 \\ 0 & 0 \end{bmatrix}}_B \begin{bmatrix} \underline{u}(t-1) \\ \underline{u}(t-2) \end{bmatrix} \quad [8-12]$$

where $\underline{T}(t)$ is the vector of all n zone temperatures at time t , $\underline{u}(t)$ is the vector of all system inputs at time t , A_{ij} is the matrix of coefficients relating zone temperatures at time $t-j$ to the temperatures at time t , and B_{jk} is the matrix of coefficients relating inputs at time $t-k$ to the temperatures at time t . The dimensions of the A matrix are $(p*n)*(p*n)$, where, as before, p is the number of lag terms used, and n is the number of zones in the building.

The roots of the characteristic equation of A can be manipulated to determine the time constants for all system modes.

Rather than creating the rather unwieldy A matrix above, the same results may be obtained by creating a single difference equation of order $(p*n)$ for each zone temperature, excluding all other zone temperatures:

$$\begin{aligned} T_1(t) &= a_1 T_1(t-1) + a_2 T_1(t-2) + \dots + a_{(p*n)} T_1(t-(p*n)) + \\ &\quad + b_{11} u_1(t-1) + b_{12} u_1(t-2) + \dots + b_{1(p*n)} u_1(t-(p*n)) + \\ &\quad \vdots \\ &\quad + b_{q1} u_1(t-1) + b_{q2} u_1(t-2) + \dots + b_{q(p*n)} u_q(t-(p*n)) \end{aligned} \quad [8-13]$$

where q is the total number of inputs to the system, including exogenous inputs. The coefficients may be determined from experimental data using standard least squares. The transfer function relating u_1 to T_1 may be written (changing the coefficients b_{11} , b_{12} , b_{13} , *etc.* to b_1 , b_2 , b_3 , *etc.*):

$$T_1(z) = \frac{b_1 z^{-1} + b_2 z^{-2} + \dots + b_{p*n} z^{-p*n}}{1 - a_1 z^{-1} - a_2 z^{-2} - \dots - a_{p*n} z^{-p*n}} u_1(z) \quad [8-14]$$

or

$$T_1(z) = \frac{b_1 z^{-1} \left(1 + \frac{b_2}{b_1} z^{-1} + \dots + \frac{b_{p*n}}{b_1} z^{-p*n+1} \right)}{1 - a_1 z^{-1} - a_2 z^{-2} - \dots - a_{p*n} z^{-p*n}} u_1(z) \quad [8-15]$$

This transfer function is then factored:

$$T_1(z) = \frac{\beta_1 z^{-1} (1 - \beta_2 z^{-1}) (1 - \beta_3 z^{-1}) \dots (1 - \beta_{p*n} z^{-1})}{(1 - \alpha_1 z^{-1}) (1 - \alpha_2 z^{-1}) \dots (1 - \alpha_{p*n} z^{-1})} \quad [8-16]$$

with $\beta_i = b_i$ serving as the common gain constant, K .

Extracting the characteristic equation from Equation [8-16] yields:

$$(1 - \alpha_1 z^{-1})(1 - \alpha_2 z^{-1}) \cdots (1 - \alpha_{p^*n} z^{-1}) = 0 \quad [8-17]$$

the roots of which are the full system poles in the z plane. Note that it will be possible to obtain n sets of roots, one for each zone temperature. Theoretically, they are identical sets, but practically, they will differ due to experimental noise and to the extent the zone temperatures communicate with each other (and *vice versa*).

Standard least squares or linear system identification machinery (*e.g.*, Ljung 1991; Wunsch 1996) may be used to determine the α 's. However, we may bring more knowledge to bear to help create a more physically plausible model. This topic is discussed in the next section.

8.3.2 Constraints Imposed by Dynamic Thermal Systems

Due to the fact that there exists no thermal analog of inertance or momentum (Karnopp *et al.* 1990), an open-loop thermal system exhibits no overshoot or oscillations when excited by an input such as an impulse or step function. These systems are inherently stable, and all poles and zeros fall on the real axis of the left-half s-plane. If time delays exist in the system, zeros are introduced on the real axis of the right half s-plane. This information may be used to our advantage to create more realistic models. We will do so by introducing constraints to the general least-squares problem introduced in the last section.

8.3.2.1 Constraints on Poles

Before presenting the constraints, it is useful to recall that the s-plane maps to the z-plane in the following way:

$$z = e^{sT}$$

$$s = \frac{\ln(z)}{T} \quad [8-18]$$

From the above equations, it is clear that any points on the negative real axis of the s-plane will map to points in the interval [0,1) in the z-plane. Points on the positive real axis of the s-plane map to the interval (1,∞].

Accordingly, the thermal system poles, α_i , which are real and negative in the s-plane, must all be constrained to lie on the interval $\alpha_i \in [0,1)$ in the z-plane. In the z domain, these poles may be understood as follows. The geometric sequence of points

$$\{\alpha^n\}, \text{ for } n = \{0,1,2,\dots,\infty\}$$

$$= \frac{1}{(1 - \alpha z^{-1})} \quad [8-19]$$

provided $|\alpha| < 1$. (A unit sample period was assumed.) Shown below are two plots of this sequence (the impulse response of the transfer function given in [8-19]), given $\alpha = 0.5$ and -0.5 . When $\alpha = 0.5$, the pole is at $z = 0.5$; when $\alpha = -0.5$, the pole is at $z = -0.5$.

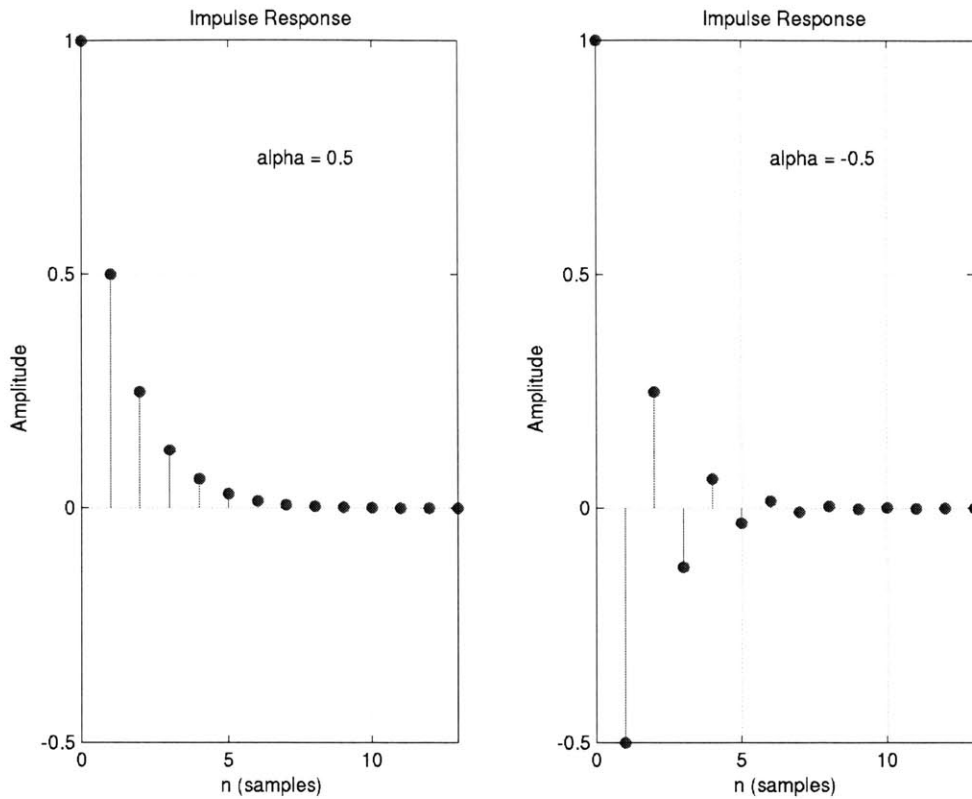


Figure 8-5 Plot of impulse response of first-order system with $\alpha = 0.5$ and -0.5 . (ImpulseResponsePlots.m)

Clearly, the plot with $\alpha = -0.5$ is not physically reasonable for the thermal system. This illustration confirms that poles must be constrained to remain on the positive real axis in the z -plane.

The geometric series shown in the first plot of Figure 8-5 may be written in the form of an exponential decay:

$$\alpha^n = e^{-\frac{nT}{\tau}}, \text{ or} \quad [8-20]$$

$$\tau = -\frac{T}{\ln(\alpha)}$$

where T represents the sample period. This equation yields a positive, real time constant, τ , for $0 \leq \alpha < 1$.

8.3.2.2 Constraints on Zeros

As mentioned at the outset of Section 8.3.2, the zeros of a thermal system must also be real. In the s-plane, they are all negative except in the case of time delays. As for poles, these zeros map to positions in the z-plane on the positive real axis (with any positive s-plane zeros mapping to $z \in (1, \infty]$). It will be shown that it is actually possible to have zeros anywhere on the negative real axis in the z-plane.

Given a continuous transfer function in the s domain,

$$H(s) = \frac{1}{(s+a)(s+b)} \quad [8-21]$$

there are several ways to generate a discrete transfer function in the z domain (Oppenheim *et al.* 1999).

The bilinear transform, or Tustin's approximation, approximates:

$$s = \frac{1}{T} \ln(z) \approx \frac{2}{T} \frac{(z-1)}{(z+1)} \quad [8-22]$$

ignoring higher-order terms. When this expression is substituted for s in the continuous transfer function, zeros appear at $z = -1$. Other methods such as zero-order hold (step-invariant method) or first-order hold (ramp-invariant method) lead to the introduction of zeros at other locations on the negative real z-plane axis. Given this situation, zeros at any location on the real z-plane axis are acceptable when generating a system model.

To provide a sense for the dynamic impact of the zeros, the time constants can be calculated by analogy to those for the poles (recall the definition of β in Equation [8-16]):

$$\beta^{-n} = (1 - \beta z^{-1}) = e^{-\frac{nT}{\tau}}, \text{ therefore} \quad [8-23]$$
$$\tau = \frac{T}{\ln(|\beta|)}$$

where the absolute value is introduced to accommodate negative zeros.

8.3.3 Algorithm for Obtaining a Physically Plausible Model

8.3.3.1 The Algorithm

The constraints developed in Section 8.3.2 were used to modify a standard least squares parameter identification routine⁷⁶. An outline of the algorithm is provided in this section.

⁷⁶ See TestConstrainedOptimized2.m.

Constraint summary:

poles: $0 \leq \alpha \leq \alpha_{max}$, where $\alpha_{max} = \exp(-T/\tau_{max})$

zeros: $-\infty \leq \beta \leq \infty$

τ_{max} is defined to be a maximum realistic time constant for the given problem. As long as $\tau_{max} < \infty$, the algorithm will work.

1. Compute the a and b parameters (Equations [8-14] and [8-15]) using a least-squares algorithm (as before, singular value decomposition was used to perform this task).
2. Factor the numerator and denominator polynomials to obtain poles and zeros. If all the poles and zeros satisfy the constraints, proceed to Step 6; otherwise proceed to Step 3.
3. If any of the poles (α 's) are negative or complex, replace them with randomly selected poles satisfying the constraints. Likewise, if any of the zeros (β 's other than β_1) are complex, then set them to zero. (Alternatively, these could be randomized.)
4. The full set of α 's and β 's becomes the initial decision variable vector for the constrained optimization routine in the MATLAB Optimization Toolbox, `fmincon.m`. The constraints listed above are given to the routine.
5. The objective function to be minimized is the ||one-step-ahead prediction error on the training set||². To evaluate the prediction error, the decision variables (α 's and β 's) must first be converted to polynomial coefficients (a's and b's). These coefficients are then multiplied by the information matrix to yield predictions for the temperatures at the next time step. To aid the optimization routine, the gradient is calculated at each step.
6. Algorithm terminates.

Several comments about the algorithm:

- The algorithm is repeated for a range of p's (number of lag terms used). The model yielding the lowest root mean square prediction error on a test set is selected.
- The algorithm is repeated n times to create a model for each zone's temperature.
- The routine `fmincon.m` locates local minima—results vary depending on the initial decision vector. The robustness of the algorithm could be improved using genetic algorithms or simulated annealing.
- It may be possible to obtain models showing a higher order if all n models were established simultaneously, *i.e.*, solving for the parameters in the A and B matrices of Equation [8-12]. However, it is not clear how to impose the nonlinear constraint that the roots of the characteristic polynomial of the A matrix be real and positive (in the z-plane).
- It may be possible to determine the appropriate order more quickly by examining the results of Step 1 more carefully. The presence of complex zeros or poles, or negative poles may indicate that those poles/zeros should be eliminated, thereby reducing the system order.

8.3.3.2 Demonstration of Algorithm

To demonstrate the performance of the algorithm just described, the following system with time constants of 6 and 10 was generated⁷⁷:

$$H(z) = \frac{\exp(-1/3)z^{-1}}{\left(1 - \exp(-1/6)z^{-1}\right)\left(1 - \exp(-1/10)z^{-1}\right)} \quad [8-24]$$

The system, with sample time $T=1$, was driven by the input:

$$u(t) = \sin(nT/7) + \sin(nT/12) \quad [8-25]$$

Noise ($\sigma^2 = 0.01$) was added to the measured signal.

The results of the algorithm are shown in Figure 8-6.




Figure 8-6 Demonstration of the constrained system identification algorithm. Top: Train and test errors versus the number of lag terms used. Middle: Calculated time constants versus the number of lag terms used. Bottom: Time constants for system zeros versus the number of lag terms used.

⁷⁷ RunConstrainedOptimizationKnown.m

Note that the actual time constants were well approximated in all models shown with time lags exceeding 1. Extra time constants of approximately zero were found for systems of order 3 and higher. Theoretically, these time constants are identically zero. In the upper plot, the error on the training and test sets are shown. The minimum test error shown is for the third-order system, and is only $2e-5$ smaller than that for the second. A test such as Akaike's information criterion (Ljung and Glad 1994) could be used to select the second-order system automatically.

It should be noted that the system to be identified in this case was a pure discrete system. When data from a continuous system are used, the results are less promising. The continuous system:

$$H(s) = \frac{1/60}{(s + 1/6)(s + 1/10)} \quad [8-26]$$

excited by the same input yields:




Figure 8-7 Demonstration of the constrained system identification algorithm given a continuous system. Top: Train and test errors versus the number of lag terms used. Middle: Calculated time constants versus the number of lag terms used. Bottom: Time constants for system zeros versus the number of lag terms used.

The time constants identified are incorrect. For this second-order system, they are calculated to be 16.9 and 0.03. The origin of the discrepancy is not understood. When no noise is present, the time constants are accurately identified. The situation does not change if sample period is reduced or if the input is changed to a random signal. The same behavior was observed for higher order systems with (inaccurate) and without (accurate) added noise. The results are not an artifact of the constrained algorithm—they were observed when standard, unconstrained least squares was used as well.

Recall that experimental measurements recorded at time t were actually the time-averaged values of those measurements from $t - 1$ to t . These particular conditions under which building data were acquired were replicated using SIMULINK to determine if this miscalculation of time constants occurs under those conditions.⁷⁸ No change was found when the continuous system data were gathered as they were during experimental data acquisition.

This curious behavior should be more carefully examined.

8.3.3.3 A Caveat

One issue to keep in mind is that small errors in the determination of the roots of the characteristic equation can lead to a large errors in the associated time constants due to the exponential relation between the two (see Equation [8-20]). Roots of $0.9 \pm .05$ lead to time constants of $9.5 + 10/-3.3$. This error increase is not practically important, since when time constants are used to establish system phase lags, the root is used directly. In this particular example with $\tau = 9.5 + 10/-3.3$, the time lag exhibited by a first order system driven by a sinusoidal input with frequency $2\pi/24$ ranges from 4.4 hours to 5.8 hours⁷⁹.

8.4 Houghton Hall Time Constants

In an effort to obtain an understanding of how the thermal mass in Houghton Hall affected the dynamic response of the building, models were created with the aim of extracting the building time constants. This was done using data gathered in the summer months, when the building was cooled via natural ventilation, sometimes boosted with exhaust fans. The expectation was that, since airflow reduces the effective thermal resistance of the building, the time constants would be shorter in the summer than in the winter. Several wintertime datasets were examined to test this hypothesis.

8.4.1 Summertime

A linear model (with no sub-modes) of each building zone was made, with the following measurements as inputs:

- T_{out}
- Solar irradiation on five building faces (roof and four sides)
- Wind speed
- Electrical gains on three floors
- Constant term (not needed if data means are removed)

⁷⁸ RunExperimentalDAQ.m

⁷⁹ TimeDelayDemo.m

Data from mid-August through early October 2003 were used. The results of model formation for the ground-floor atrium are presented here as an example. Plots similar to those shown in Figure 8-6 are given below.




Figure 8-8 Results of constrained optimization for the ground-floor atrium. Top: Train and test errors versus the number of lag terms used. Bottom: Calculated time constants versus the number of lag terms used. (AG_order_selection.m (and tau_record.mat))

As shown in Figure 8-8, the rms error on the test set decreases as the system order is increased. A minimum is not found for the range of lag terms used. However, the dominant time constant remains steady in the vicinity of 6 hours. All other time constants are 0.01 hours or shorter. Given that the sample period is 0.25 hour, it is impossible to detect such high-frequency behavior, so those small time constants must be neglected. Although the rms test error decreases for orders greater than 2, it was decided to select a model of second order to estimate time constants. It should be noted that the behavior observed here is consistent with that described in Section 8.3.3.2 for continuous systems—there is a dominant time constant accompanied by several others of negligible size. (This is a cautionary note that the time constants calculated in this section may be distorted. If the same pattern occurs in this section as was observed in Section 8.3.3.2, the time constants calculated here are too large.)

Time constants were calculated for all zones as above using two time lags. The results are shown in Figure 8-9. The mean time constant for the entire building was 6.8 hours.




Figure 8-9 Time constants by building zones. (RunAllConstrainedOptim.m)

The results shown in Figure 8-9 are consistent with the following intuitive observations:

- The three atria have decreasing time constants as the floor level increases. This pattern may be due to the relatively smaller thermal mass in the upper atria and the reduced thermal resistance due to greater airflow.
- The three atria have smaller time constants than the office areas due to their relative lack of thermal mass.
- The first-floor office areas may have the greatest time constants due to their central location within the building (thermal mass above in the ceiling and, to a lesser degree, the floor).

The previous calculations were performed without regard to the various modes of operation of the building. The same type of calculations were performed below, but using subsets of the data according to whether the exhaust fans were running and whether the building was occupied.

Figure 8-10 shows the results of the calculations, assuming that two time lags were still adequate. Note that there were several instances where the second calculated time constant was

approximately 0.1 hour and one where it was nearly 0.3 hour. Since the sample period is 0.25 hour, a time constant of 0.1 hour would be nearly undetectable and probably should be ignored. The time constant of 0.3 hour might be legitimate. It occurs in tandem with a particularly short first time constant of 0.7 hour.




Figure 8-10 Time constants by building zones and by mode of operation. Top: First time constants. Bottom: Second time constants. (RunAllConstrainedOptim_modes.m)

When the building was unoccupied, all occupant-controlled windows were shut, leading to *potentially* lower airflows than occurring during occupied periods⁸⁰. This observation would lead to greater time constants during unoccupied periods (second and fourth bars) than during occupied periods (first and third bars). The expectation is borne out for the periods when the fan is off. When the fan is on, the second-floor zones and North 1 display greater time constants when unoccupied versus occupied (the expected result), but the other zones display the opposite pattern.

Since the operation of the fan may increase airflow through the building (see discussion in Section 7.7), the time constants for periods with the fan on would be expected to be lower than those found when the fans are off. It can be seen that the Fan ON/Fan OFF changes did follow the expected pattern during unoccupied periods (with the one exception of Atrium 1). However, during occupied periods, the opposite pattern is generally observed.

⁸⁰ Since the occupants were observed to not always open windows during the day.

As was the case in Figure 8-9, the atrium time constants tend to be lower than those in the office areas. It is not clear why the time constants in Figure 8-10 are uniformly smaller than those in Figure 8-9.

8.4.2 Wintertime

In order to assess the impact of airflow through the building on its dynamic response, two wintertime periods in December 2003 and January 2004 were selected for modeling and analysis. The days selected were on weekends and holidays, so there was no heat supplied from the boilers (due to implementation of a thermostat setback strategy). The inputs to the model were then restricted to:

- Solar irradiation on five building faces (roof and four sides)
- T_{out}
- Constant term

Electrical data were not supplied, but rather were partially accounted for with a unity “input” that served as the intercept for the regression. The data available at the time were for the North office on the first floor. The modeling process was not entirely successful—the model developed for the December period gave poor results when tested with the January period, and *vice versa*. It may be that the constant terms used for one model were entirely inappropriate for the other due to changes in boundary conditions (such as internal loads, wind speed, outdoor temperature) from one data set to the next. Outdoor temperatures displayed very little periodicity. They are shown below in Figure 8-11.

Despite these less-than-ideal conditions, time constants were obtained from the two data sets: 12.2 hours and 18.0 hours for the December and January data, respectively. The summertime time constant for the North 1 office was approximately 11 hours (when the building modes of operation were ignored—see Figure 8-9). Given the modeling difficulties encountered with the wintertime data, it appears safe to conclude that closed-building time constants tend to be at least slightly longer than open-building time constants.

To improve the wintertime time constant calculation, additional overnight data (when the boilers were off) could be used to boost the size of the data set.




Figure 8-11 Wintertime outside and inside temperature data. (PlotWintertimeData.m)

8.4.3 First-order Models

According to Equation [8-13], one would expect the difference equations describing the various zones of Houghton Hall to be at least seventh order—one order per zone variable. Given several time lags per variable, the order of the system is expected to grow rapidly. Why then are first-order difference equations obtained?

There are several possible scenarios at least partially consistent with the observation.

First, consider the case of the South office on the ground floor, and assume that all air enters the office windows, passes through to the ground-floor atrium, then up and out the stack vents. In this scenario, the airflow dictates the interaction between zones. Since there is no backflow of air in this scenario, it is plausible that the difference equation describing the South ground-floor office zone could be first order since the office zone is entirely independent of the atrium zone (and others)⁸¹. Of course, were this scenario observed, then the difference equation for the ground-floor atrium would be at least second order. We might expect to discern a pattern of

⁸¹ This argument assumes the flow rate is independent of the temperatures in the other zones.

typical airflow in the building by the order of the governing difference equations in the different zones.

Second, consider the difficulties encountered above in Section 8.3.3.2 when the underlying system to be identified was continuous rather than discrete. In that case, dominant first time constants were found irrespective of whether standard least squares or the constrained least squares algorithm was used. Secondary and higher-order time constants were generally negligible. Unfortunately, the dominant time constant found did not reflect any of the time constants of the underlying system. If the pattern observed in the Houghton Hall data is analogous to that seen in Section 8.3.3.2, then the calculated time constants are too large. This question will be taken up again in Section 8.5.

8.5 Discussion of Time Constants and Time Shifts

8.5.1 Calculating Time Shifts with the System Transfer Function

The time constants and, more directly, system poles and zeros provide information regarding the building's response to a variety of inputs. The initial-condition response (all inputs are zero) is determined solely by the time constants in the system characteristic equation. The response of the system to ordinary inputs cannot be interpreted solely by examining system poles (unless, of course, there are no zeros).

The most important inputs to a naturally ventilated building, such as the outside temperature and solar radiation, follow a fundamental frequency of oscillation of 1 day^{-1} . Internal gains also follow such a pattern. Therefore, it is reasonable to ask how the building responds to sinusoidal inputs of frequency $2\pi/24 \text{ rad/hour}$.

The effect of the system transfer function $H(z)$ is to attenuate the input signal and to impose a phase shift on it. Of great interest in the naturally ventilated building is the phase shift—by how many hours can the effects of the peak outside temperature be delayed? If it is possible to delay those effects until after occupied hours, the building has an improved chance of remaining comfortable.

8.5.2 Frequency Response for Houghton Hall

Recall that 6.8 hours was the average time constant found for the seven temperature zones. Given a simple first-order system of time constant 6.8 hours, the associated transfer function is:

$$H(s) = \frac{1/6.8}{s + 1/6.8} \quad [8-27]$$

whose frequency response is shown in the Bode plot in Figure 8-12. The phase lag is shown for an input whose frequency is $2\pi/24 \text{ rad/hour}$. This phase lag of 61° corresponds to a time delay of $(24 \text{ hours}) \cdot (61^\circ) / (360^\circ) = 4 \text{ hours}$. Note that this time delay is what would be expected for a system with a single pole.




Figure 8-12 Bode plot for $H(s)$ as given in Equation [8-27]. (FreqResponse.m)

Consider the plots of the indoor and outdoor temperatures shown for Houghton Hall in Figure 8-13, below. Examination of the data reveals that this estimated time delay is consistent with observed time delays between peak outdoor temperatures and peak indoor temperatures. Note that the estimated times of the extrema of the outdoor temperatures were established by using overall curve shape rather than by locating the true extrema.

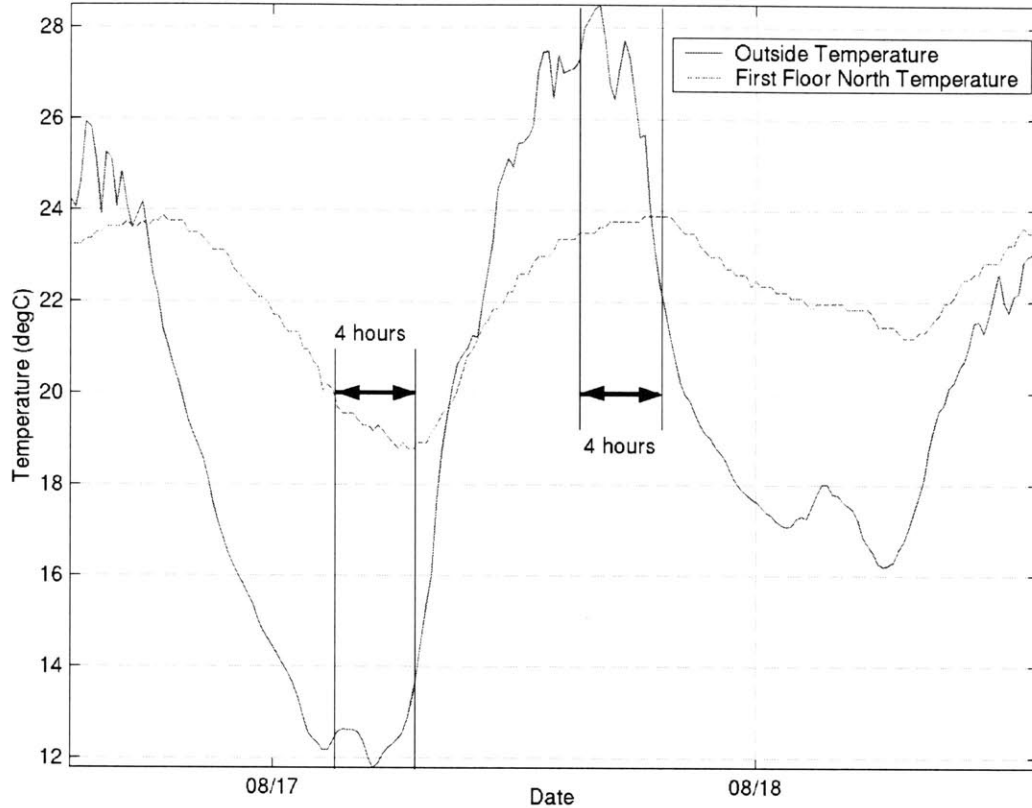


Figure 8-13 Demonstration of time shift of indoor temperature extrema relative to outdoor temperature extrema. (GenerateTimeShiftFig.m)

The time delays may be calculated directly from the discrete transfer function linking outside and inside temperatures. Such an approach allows the incorporation of the effect of zeros on the time shift. Zeros exert a positive influence on the phase angle, so time shifts for systems with zeros are expected to be smaller than those with poles alone. Bode plots were generated as above⁸² and the phase angle at an input of $2\pi/24$ rad/hour was determined. This phase lag was converted to a time lag. The results for the seven building zones are shown below in Figure 8-14. The average time shift was 3.8 hours, slightly shorter than the average 4 hour shift reported above.

The pattern of shifts shown in Figure 8-14 exhibits some similarity to the pattern of time constants in Figure 8-9—the Atrium 1 and second-floor shifts are shorter than those for other

⁸² This time a discrete transfer function was used. It is possible to convert $H(z)$ to $H(s)$ by using the relation $s = \ln(z)/T$ and find the phase lags using $H(s)$, but this approach is not straightforward when negative real zeros exist. Also, results when positive real zeros > 1 exist, the time shifts become strongly negative (*i.e.*, there is a phase lead). This behavior is explored in `TimeShifts_continuous.m`.

zones, and the North and South 1 zones demonstrate the longest time shifts. However, the overall range of time shifts is compressed relative to that of the time constants




Figure 8-14 Calculated time shifts of inside temperature extrema relative to outside temperature extrema for the seven building zones. (TimeShifts.m)

8.6 Broadmoor Time Constants

A similar set of calculations was performed for Broadmoor to obtain information about the time constants of the various zones. The following measurements were used as inputs:

- T_{out}
- Solar irradiation on five building faces (roof and four sides)
- Wind speed
- Electrical gains in the Basement and first floor
- Constant term

The following table was generated using the first 35,000 data points of the data set, terminating before the Sunspace door was opened. The presence of control modes in the building was ignored for these initial calculations.

Number of Lags	1		2		
	constrained		unconstrained		
Zone	Tau_1	Tau_1	Tau_2	Tau_1	Tau_2
Assembly	8.1	8.3	0.0005	9.4	.002+.004i
Mass	11.7	11.5	0.0006	11.6	.002+.004i
Attic 1	4.3	4.3	0.0005	4.4	.003+.003i
Attic 2	2.3	2.3	0.0005	2.3	.002+.004i
Sunspace	0.6	0.7	0.007	0.7	0.007
Basement	23.5	18.8	0.008	18.8	0.008

Table 8-2 Time constants calculated for the six building zones at Broadmoor (expressed in hours). The first-order system time constants are listed under the column for 1 lag term. The constrained and unconstrained algorithm yielded identical results. The second-order system (2 lags) time constants are shown for the constrained and unconstrained cases. (BroadmoorTauCalc.m)

Several aspects of Table 8-2 are worthy of note.

First, the difference between the first-order time constant and the dominant time constant of the second-order system is small, with the exception of the Basement zone. The secondary time constants for the second-order system were negligible in comparison with the primary time constants. Recall that the sample time for the Broadmoor experiments was one minute, or 0.017 hours. If a dynamic mode exists with a time constant roughly one-quarter or less than the sample period, it will have decayed in less time than the duration of the sample period and therefore will be invisible. Consequently, the secondary time constants for the first four zones can be ignored. The secondary time constants for the Sunspace and Basement are larger than the arbitrary threshold of (sample time)/4, so it is conceivable from these results that some fast dynamics do exist in those zones. It is difficult to rationalize, however, why there would be a fast mode in the Basement and not in the Assembly zone. Due to the dominance of the first time constant, however, the systems can be regarded as essentially first order.

Second, it is evident that when the least-squares routine is not constrained, the resulting time constants can be complex, as they are for the first four zones. For the case of the Assembly zone, constraining all time constants to be real had the effect of reducing the time constant by more than an hour.

In the previous discussion, the presence of control modes was ignored. Since the building's thermal responses varied according to the control mode in effect, it is interesting to examine the time constants associated with the individual control modes at Broadmoor. Shown in the following table are the time constants associated with the various modes of the Assembly and Mass zones. The data used in the calculations were the first 54580 points of the data set, with the exception of the data from the period when the Sunspace door was opened.

Mode	Assembly		Mass	
	Tau_1	Time Shift	Tau_1	Time Shift
1	0.4	0.4	6.2	3.9
4	11.3	4.8	22.3	5.4
5	2.6	2.3	14.4	5
8	6.0	3.8	11.0	4.7
9	2.7	2.4	7.3	4.2

Table 8-3 First-order system time constants for the different control modes of the Assembly and Mass zones (shown in hours). The time shifts associated with excitation at a frequency of $2\pi/24$ rad/hr are shown (in hours) as well. Mode definitions are given in Table 5-2. (BroadmoorTauCalc.m)

The time constants shown in Table 8-3 demonstrate some qualitative behavior consistent with observations of the system performance. For example, consider the time constants for the Mass zone. The system is in Mode 1 when the fan is on. It is reasonable that this time constant is the fastest. The slowest is for Mode 4, which represents the building fully closed. This is also as expected. Modes 5, 8 and 9 are for cross ventilation, building open with Attic 2 windows open, and building open with Attic 2 windows shut, respectively. The pattern of decreasing time constants for these three modes is consistent with observations that building cooling rates increased from Mode 5 to 8 to 9. Essentially, the time constants give an indication of the flow rates, or equivalently, the cooling capacity of the different modes⁸³.

Unfortunately, the pattern breaks down for the Assembly zone. Again, the Mode 1 time constant is fastest, and the Mode 4 time constant is slowest, as one would expect. The time constants for the remaining modes do not follow a pattern that is easily rationalized. One should recall, however, that the models used to estimate time constants are not the same models that led to such accurate predictions of the temperature trajectories. To use those models for time constant calculations, one must assemble the A matrix in Equation [8-12], calculate and factor the characteristic equation and then extract the time constants. As of this writing, no method exists of ensuring that the time constants be real.

The unexpected pattern exhibited by some of the time constants in Table 8-3 also serves as a reminder that the time constants are functions of airflow rates, which are functions of temperature differences. Other than for the fan mode, a given mode is not characterized by a fixed flow rate. Therefore, the time constant calculated for that mode will reflect only the time-averaged conditions of the data set used to generate the model.

8.7 Summary and Possible Future Extension of Concepts

The questions triggering the work begun in this chapter have not been definitively answered. It was possible to demonstrate that the more massive elements of the two buildings examined displayed longer time constants. For example, the Basement at Broadmoor had the longest time constant, followed by the Mass temperature. At Houghton Hall, the Atria had the shortest time constants, which is consistent with their smaller mass relative to the office spaces. As far as the

⁸³ When $T_{\text{outside}} < T_{\text{Assembly}}$.

importance of thermal mass was concerned, it was shown that the peak of a sinusoidal input was shifted approximately four hours at Houghton Hall. The time shifts at Broadmoor were calculated for different control modes, but examples in the data set could not be found showing a five-hour shift for Mode 4, for example. A shorter time shift was observed. Thermal mass clearly plays a role in the buildings—it is engaged in the cooling process. Otherwise night cooling would be fruitless. However, it was found at Broadmoor that there was no benefit of long-horizon control since the thermal mass in the building was insufficient to carry one day's "charge" to the next. The role of the fan and its impact on the time constants at Broadmoor are striking. The Assembly time constant is a fraction of an hour when the fan is in operation, but is at least several hours when the fan is off. Such a clear result cannot be found at Houghton Hall. The effect of the fans remains elusive.

In the process of creating a framework for answering the above questions, a method was developed to calculate time constants from experimental data in such a way that they were ensured to be physically plausible. Time constants were constrained to be real and positive. In addition, the role of the system zeros was highlighted in calculating the time shifts of system inputs.

While the definition of time constants for a naturally ventilated building (or any building with variable airflow rates) is problematic due to the inherent non-linearity of the system, the concepts of dynamic modes and controllability are worth consideration during the design process. For instance, even with an infinite time constant, one is limited to a six-hour time shift for a first-order system (since a first-order system can supply at most a 90° phase lag). If it were desirable to obtain greater shifts, one could design a system to have higher order. Essential to the design of such a system would be the determination of how to rapidly charge the mass, but then allow a controlled discharge over time. Reflection on the design problem quickly reveals that this is a mixed-mode system—one whose control could be approached in a manner similar to that presented in Chapter 6. Great opportunities remain to probe some of these complex issues related to the dynamics of naturally ventilated buildings.

9 Summary, Conclusions and Future Research

9.1 Summary

The research described in this thesis was motivated by two factors: the recognition that mixed-mode cooling carries the potential for significant energy savings in many locations worldwide, and that very little progress had been made on the challenging problem of controlling such systems in a way that maximizes the energy-savings benefits of mixed-mode cooling while preserving comfortable conditions for building occupants. The primary contributions of the thesis directly addressed the second factor. Specifically,

1. A flexible system identification framework was developed to accommodate the unique features of mixed-mode buildings.
2. The framework was demonstrated to be effective and yield accurate predictions of the thermal behavior of a mixed-mode building more complex than any modeled (and validated) in the literature. The prediction accuracy met or exceeded that reported for the models of the simpler buildings studied.
3. An approach to optimal model-predictive control was developed that took advantage of an efficient method for encoding model information to simplify and accelerate the optimization calculations.

The primary topics covered in the thesis are summarized in the following paragraphs.

In Chapter 2, a series of simulations was performed to assess the potential impact of using mixed-mode cooling in a variety of locations in the United States. While the strategy was found to be poorly suited to hot and humid climates, it was found that cooling energy savings resulting from the use of mixed-mode cooling rather than standard air conditioning could be expected to range from 5% to greater than 20% in other locations. Such estimates were for a building not optimized for mixed-mode cooling. Improvements in building design, such as the incorporation of shading, increased thermal mass and the use of natural rather than mechanical ventilation, increased the savings potential dramatically. For example, in Boston, MA, the annual cooling energy savings were estimated to be approximately 10% when purely mechanical mixed-mode cooling was used. The savings were estimated at 30% if natural ventilation was incorporated into the mixed-mode cooling strategy.

The challenges of modeling a thermal system influenced by natural ventilation were discussed in Chapter 3. It was recognized that many different airflow patterns, or regimes, could result from the constellation of elements affecting natural ventilation: the outside temperature, the inside zone temperatures, the configuration and state of all of the apertures in the building, and the wind speed and direction. Several issues were asserted to preclude the use of first-principles models to predict airflow: the uncertainty of the model structure due to the wide variety of influencing elements just mentioned, and the large set of unknown parameters required by such models. Furthermore, even if a careful first-principles model were developed (and calibrated) for one building or building zone, the effort would need to be repeated for each new building configuration encountered. A flexible type of regression, called PHDRT, was presented as an effective means of handling the nonlinearities expected in systems influenced by natural

ventilation without requiring *a priori* assumptions of model structure. This method was found to model accurately the nonlinear thermal impact of airflow in a simple simulated building.

The multi-zone, mixed-mode building used for the bulk of the research in this thesis was introduced in Chapter 4. The building was the Saltonstall Nature Center at the Massachusetts Audubon Society's Broadmoor Wildlife Sanctuary in Natick, MA. Experimental details were provided, as was a description of the extensive data acquisition and control system developed for this research. Simple comparisons between the operation of the building as recorded in an earlier study and its operation during this research provided evidence that several building retrofits completed as part of the research (in addition to improved control strategies) were indeed beneficial to comfort in the building.

Chapters 5 and 6 contain the most important elements of the thesis research. Chapter 5 was devoted to the development of a modeling framework appropriate for mixed-mode buildings, and Chapter 6 was devoted to using the models developed in Chapter 5 to optimize temperature control within the building.

The most surprising element of Chapter 5 (which might not have been discovered had experiments not been performed on a real building) was that the approach developed in Chapter 3 for handling the nonlinearities associated with natural ventilation was unnecessary. No benefit was found relative to the performance achieved using a linear model. Furthermore, it was found that nonlinear "black-box" models, such as neural networks or kernel recursive least squares (KRLS) did not improve model predictions. In fact, cases were found where instabilities and erratic predictions were introduced by such models. In light of these discoveries, an attempt was made to construct the most effective linear model possible. The modeling framework consisted of two elements: 1) the division of the building's operation into control modes, each representing a distinct group of control settings, and 2) the optimization of the selection of features (or inputs) used to construct the models.

The control modes were determined manually, based on knowledge of how the building was operated. For example, one mode described the building's behavior when all apertures were shut. Another mode described the behavior when apertures were set to encourage cross ventilation. A procedure to automate the process of defining the key building modes was proposed. Dedicated models were created to learn the operation of the building in each mode. Therefore, the building "model" actually consisted of a set of models, each appropriate for a given control mode.

Many elements influence the thermal behavior of a building. Due to the difficulty of assessing the relevance of these elements *a priori*, an automated method for selecting model features (or inputs) was developed and demonstrated. Features were judged strictly on how their inclusion in a model impacted the accuracy of model predictions. The algorithm led to some interesting findings, such as that wind did not play an important role in the thermal behavior at Broadmoor. For the case of Broadmoor, the performance benefit obtained from automated feature selection was less important than that obtained from the incorporation of building modes into the modeling approach.

Model predictions in multi-day (and multi-week) pure simulations differed from measured temperatures by an rms error of 0.3 – 0.4 °C for most zones. This modeling accuracy exceeded that reported in the literature for much simpler buildings.

With confidence that an accurate model had been developed, a model-based approach to optimal control was pursued in Chapter 6. No model-based control strategy has been presented in the literature for temperature control in an actual occupied naturally ventilated or mixed-mode building. The primary goal of the chapter was to develop a method of optimizing the control settings over a 24-hour horizon in order to maintain daytime temperatures in a comfortable region and to minimize fan energy consumption. Borrowing a concept from others' work, the 24 hours were divided into larger segments to reduce the computational load. Two or three nighttime periods were used in addition to two daytime periods.

The initial approach taken was to determine the optimal control schedule using genetic algorithms and simulated annealing. Such an approach was effective, but not efficient. Insight about the nature of diurnal temperature swings coupled with model-derived knowledge permitted the entire optimization procedure to be reduced to the selection of a single code number out of an array of numbers ranked by their cooling capacity. The essence of the ranking was the recognition that the most aggressive cooling approach should be undertaken at the time of greatest inside-outside temperature difference. The proposed efficient encoding of the decision variables may be undertaken for any building where daytime outside temperatures are too warm for cooling (as was the case for Broadmoor). The optimal control method was also extended to apply to buildings in other locations. Algorithms permitting an extension of the control horizon to 48 or 72 hours were presented. It was shown that such methods did not benefit the cooling at Broadmoor due to insufficient thermal storage.

The success of the modeling framework outlined in Chapter 5 provided encouragement to apply the technique to another mixed-mode building, situated in the U.K. This work was described in Chapter 7. It was shown that the use of modes combined with optimized feature selection provided significant improvements in model accuracy. Unfortunately, no data were available regarding the operation of the apertures in the building, so it was not possible to construct a model suitable for predictive control.

Among the questions raised during the work on the U.K. building and the extended horizon optimal control for Broadmoor related to the importance of the building's thermal storage capacity, or thermal mass. The main contribution of Chapter 8 was the development of a procedure for extracting thermal time constants from experimental data in such a way that the time constants are constrained to be physically plausible (real and positive). The remainder of the chapter was devoted to an exploration of the possibilities of incorporating the notion of dynamic modes into the design and control of buildings. Finally, a discussion of the relevance of time constants in a building with variable airflow was presented.

9.2 Conclusions

The initial study on the feasibility and effectiveness of mixed-mode cooling in the U.S. provides encouraging support to the possibility of bringing greater efficiency to the built environment in the U.S. These studies must be reinforced by well-documented performance of real buildings to achieve greater acceptance in the building design and construction community. It is hoped that the tools presented in this thesis will serve to enhance the chances of success of these mixed-mode buildings, leading to their more widespread development.

It is fully acknowledged that the findings related to the modeling of the Saltonstall Nature Center at the Broadmoor Wildlife Sanctuary were unexpected. The potency of linear models in this complex building was not foreseen. However, it was the hypothesis that one could not determine *a priori* the structure of the model needed to capture the thermal behavior of the building that led to the adoption of data-driven models. Ironically, the assumption that it was a *nonlinear* structure that had to be discovered through data analysis proved incorrect. The acceptance of the data as opposed to reliance on “correct analytical predictions” permitted the creation of very flexible and very accurate linear models.

The ultimate simplicity of the framework for optimizing control strategies was also unexpected. Accounting for some elementary aspects of the exogenous inputs, such as the roughly monotonic decrease of temperatures at night and the monotonic increase during occupied hours, transformed a randomized search orchestrated by genetic algorithms and simulated annealing into an efficient binary search problem. The fundamental concepts highlighted by the work on optimal control of mixed-mode cooling may be very useful for any building relying on ambient conditions to provide, or contribute to a comfortable indoor environment.

An important question to consider at this stage is: how extensible are the modeling and control frameworks to other buildings? Also, what modifications in the approach may be necessary for another building? As long as the first stage of the modeling process is begun with a test of the nonlinearity of the system, the modeling framework elucidated here is expected to perform well in other buildings. Had data from a strongly nonlinear building been available, nonlinear models could have been developed for those particular operational modes requiring them. The simplest, and—from the positive results found in the flow-regime change modeling experiments—the most effective tool to begin with would be the PHDRT algorithm. However, any other tool, such as KRLS, could be substituted in its place. The key is the focus on the essential operational modes of the building and the selection of the correct features (or model inputs)—both topics discussed at length in this thesis.

Practically, to obtain a strong model, sufficient measurements are required. Certainly all zone temperatures should be monitored, in addition to the temperatures of the building fabric. Outdoor conditions, including solar radiation, wind speed and direction, temperature (and later, humidity) should all be monitored. Only after the fact can it be determined that wind speed measurements are not beneficial, for example. Additional measurements with potential modeling import are electrical loads, which may be monitored as described in the thesis, and occupancy levels. Finally, of utmost importance is information regarding all adjustments to the building, be they aperture changes, fan-speed changes, or the use of air conditioning. Without

access to this information, one may not be able to form an accurate model useful for control (witness the case at Houghton Hall). Again, it may be possible to establish after the fact that some of the control settings may be neglected, but it is dangerous to do so from the outset.

Transferring the modeling framework to another building may also necessitate the change of sample time and/or the maximum number of lag terms used to form the model. In the interest of maintaining access to all information that could be relevant (an essential practice for data-driven modeling), it would be beneficial to maintain a small sample time and combine measurements at a later time if found useful.

Given an accurate model, the proposed framework for optimizing control strategies is expected to be robust. For buildings with access to cool temperatures at nighttime only, the exact strategy implemented in the thesis should be effective. Regarding the use of optimal control strategies over a multi-day horizon, it must be determined on a case-by-case basis whether any benefit can be derived. It is not known now how to determine *a priori* whether a building will have enough thermal storage capacity to make this technique worthwhile. Certainly, simulations during the design phase could provide some insight.

For buildings with access to cool outside temperatures during the occupied period (and for those with air conditioning), the number of control options increases. Guidance was given in the text that draws on the same fundamental insights in order to optimize control strategies in these situations.

It is the author's perspective that the concept of time constants in buildings experiencing variable flow rates needs careful re-examination. Since flow rates change, the time constants are not constant. Of greater importance and relevance to the problem of obtaining high levels of building performance is the notion of dynamic modes and the effectiveness with which they are excited. Given a clever design, a building control system may be able to manipulate the charging and discharging of the building's thermal mass to ultimately increase the effective cooling capacity of the building.

9.3 Future Research

Several of the ideals outlined in the introduction for control of a mixed-mode building have yet to be realized. One elusive ideal is the automation of building control strategy development before adequate data have been gathered to train building models. A promising research topic would be the exploration of how to incorporate (then phase out) a design-stage building model into an early-time control strategy. Another remaining ideal is an occupant-responsive control strategy, wherein the comfort preferences (*e.g.*, set points) of the occupants are learned over time and are used to inform the optimal control schedules.

Several other projects and areas of inquiry follow naturally from this research. A few are listed below:

- A mixed-mode building other than Broadmoor, such as Houghton Hall, should be fully equipped with monitoring equipment to provide a fair assessment of the validity of the modeling framework on another building.
- Once this first project has been completed, a full implementation of the modeling framework and optimal control strategy on a real, fully instrumented building should be pursued. As part of this implementation, weather and load prediction modeling approaches would need to be selected from the literature or developed.
- The focus on single-zone temperature control in a multi-zone building could be broadened to allow multi-zone temperature control.
- The techniques presented in this thesis could be extended to humidity modeling and control.
- The creative use of thermal mass for greater flexibility in thermal control remains an intriguing topic.
- The use of optimal control strategy development using other tools such as IP and DP has yet to be extended to mixed-mode buildings and may prove to be beneficial.
- The issue regarding the poor identification of continuous dynamic systems remains unresolved.
- The proposed method of optimizing mode selection could be refined.
- A careful study of the model behavior at times of mode switching could be performed to establish a reliable method for treating these transitional periods.
- A careful airflow measurement and analysis study could be performed on an actual mixed-mode building such as Broadmoor to establish the presence and importance of any nonlinearities in the flows.

10 References

- http://www.massaudubon.org/Nature_Connection/Sanctuaries/Broadmoor/index.php.
- <http://www.erh.noaa.gov/box/fcsts/BOSAFMBOX.html>.
- Abe, S. and M.-S. Lan (1995). "A Method for Fuzzy Rules Extraction Directly from Numerical Data and Its Application to Pattern Classification." IEEE Transactions on Fuzzy Systems **3**(1): 18-28.
- Aggerholm, S. (2001). Control Strategies for Hybrid Ventilation. Second International One-day Forum on Hybrid Ventilation, Delft: 33-41.
- Aggerholm, S. (2002). Hybrid Ventilation and Control Strategies in the Annex 35 Case Studies. IEA Annex 35 Technical Report. IEA.
- Ahmed, O., J. W. Mitchell and S. A. Klein (1996). "Application of General Regression Neural Network (GRNN) in HVAC Process Identification and Control." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **102**(1): 1147-1156.
- Ahmed, O., J. W. Mitchell and S. A. Klein (1998a). "Feedforward-Feedback Controller Using General Regression Neural Network (GRNN) for Laboratory HVAC System: Part I--Pressure Control." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **104**(2): 613-625.
- Ahmed, O., J. W. Mitchell and S. A. Klein (1998b). "Feedforward-Feedback Controller Using General Regression Neural Network (GRNN) for Laboratory HVAC System: Part II--Temperature Control--Cooling." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **104**(2): 626-634.
- Ahmed, O., J. W. Mitchell and S. A. Klein (1998c). "Feedforward-Feedback Controller Using General Regression Neural Network (GRNN) for Laboratory HVAC System: Part III--Temperature Control--Heating." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **104**(2): 635-642.
- Ahn, B. C. and J. W. Mitchell (2001). "Optimal Control Development for Chilled Water Plants Using a Quadratic Representation." Energy and Buildings **33**: 371-378.
- Aird, T. J. and J. R. Rice (1977). "Systematic Search in High Dimensional Sets." SIAM Journal on Numerical Analysis **14**: 296-312.
- Allard, F., Ed. (1998). Natural Ventilation in Buildings: A Design Handbook, James X. James.
- Allard, F. and Y. Utsumi (1992). "Airflow Through Large Openings." Energy and Buildings **18**: 133-145.
- Allocca, C. (2001). Single-Sided Natural Ventilation Design Analysis and General Guidelines. M.S., Mechanical Engineering, MIT, Cambridge.
- Andersen, K. T. (1995). "Theoretical Considerations on Natural Ventilation by Thermal Buoyancy." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **101**(2): 1103-1117.
- Arens, A. D. (2000). Evaluation of Displacement Ventilation for Use in High-ceiling Facilities. Masters Thesis, Mechanical Engineering, MIT, Cambridge.
- Armstrong, P. R. (2004). Model Identification with Application to Building Control and Fault Detection. Ph.D. Thesis, Architecture, MIT, Cambridge.
- ASHRAE (1997). ASHRAE Handbook: Fundamentals. Atlanta, GA, American Society of Heating, Refrigerating and Air-Conditioning Engineers.
- Awbi, H. (2003). Ventilation of Buildings. New York, Spon Press.
- Axley, J. W. and S. J. Emmerich (2002). A Method to Assess the Suitability of a Climate for Natural Ventilation of Commercial Buildings. Indoor Air **2002**: 854-859.
- Babuska, R. and H. B. Verbruggen (1996). "An Overview of Fuzzy Modeling for Control." Control Engineering Practice **4**(11): 1593-1606.
- Bertsekas, D. P. (2000). Dynamic Programming and Optimal Control. Belmont, MA, Athena Scientific.

- Blondeau, P., M. Sperandio and F. Allard (1997). "Night Ventilation for Building Cooling in Summer." Solar Energy **61**(5): 327-335.
- Bloomfield, D. P. and D. J. Fisk (1977). "The Optimisation of Intermittent Heating." Building and Environment **12**: 43-55.
- Brager, G. S. and R. de Dear (2000). A Standard for Natural Ventilation. ASHRAE Journal: 21-28.
- Braun, J. E. (1988). Methodologies for Design and Control of Central Cooling Plants. Ph.D. Thesis, University of Wisconsin-Madison, Madison.
- Braun, J. E. (1990). "Reducing Energy Costs and Peak Electrical Demand Through Optimal Control of Building Thermal Storage." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **96**(2): 38-50.
- Braun, J. E., K. W. Montgomery and N. Chaturvedi (2001). "Evaluating the Performance of Building Thermal Mass Control Strategies." International Journal of HVAC&R Research **7**(4): 403-428.
- BRE (1994). Natural Ventilation in Non-domestic Buildings. BRE Digest No 399. Building Research Establishment.
- BRE (1998). Night Ventilation for Cooling Office Buildings. Information Paper IP4/98. Building Research Establishment, Ltd.
- BRE (1999). Natural Ventilation for Offices. Building Research Establishment/NatVent.
- BRE (1999a). Mixed-mode Buildings and Systems: an Overview. General Information Report 56. Building Research Establishment, Ltd.
- Bruant, M., G. Guarraccino and P. Michel (2001). "Design and Tuning of a Fuzzy Controller for Indoor Air Quality and Thermal Comfort Management." International Journal of Solar Energy **21**: 81-109.
- BS (1991). Code of Practice for Ventilation Principles and Designing for Natural Ventilation. BS 5925. British Standards.
- BSE (1996). "Future Buildings." Building Services Journal: 22-24.
- Buckles, B. P., F. E. Petry, D. Prabhu, R. George and R. Srikanth (1994). Fuzzy Clustering with Genetic Search. First IEEE Conference on Evolutionary Computation: 46-50.
- Bunn, R. (1997). "Green Demo: Building Analysis - Environmental Office of the Future." Building Services Journal: 18-23.
- Chan, H. L. and A. B. Rad (2000). "Real-time Flow Control Using Neural Networks." ISA Transactions **39**: 93-101.
- Chang, J.-C. R. (2002). Case Studies of Naturally Ventilated Commercial Buildings in the United States. M.S., Mechanical Engineering, MIT, Cambridge, MA.
- Chaturvedi, N. and J. E. Braun (2001). "An Inverse Gray-Box Model for Transient Building Load Prediction." International Journal of HVAC&R Research **8**(1): 73-99.
- Chaudhuri, P., M.-C. Huang, W.-Y. Loh and R. Yao (1994). "Piecewise-Polynomial Regression Trees." Statistica Sinica **4**: 143-167.
- Chen, T. Y. (1997). A Methodology for Thermal Analysis and Predictive Control of Building Envelope Heating Systems. Ph.D. Thesis, Concordia University, Montreal.
- Chen, T. Y. (2001). "Real-Time Predictive Supervisory Operation of Building Thermal Systems with Thermal Mass." Energy and Buildings **33**: 141-150.
- Chen, T. Y. and A. K. Athienitis (1996). "Ambient Temperature and Solar Radiation Prediction for Predictive Control of HVAC Systems and a Methodology for Optimal Building Heating Dynamic Operation." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **102**(1): 26-36.
- Chen, Z. D. and Y. Li (2002). "Buoyancy-driven Displacement Natural Ventilation in a Single-zone Building with Three-level Openings." Building and Environment **37**(3): 295-303.
- Cheng, C.-S. and K. C. Li (1995). "A Study of the Method of Principal Hessian Direction for Analysis of Data from Designed Experiments." Statistica Sinica **5**: 617-639.

- Chiu, S. (1996). Method and Software for Extracting Fuzzy Classification Rules by Subtractive Clustering. Biennial Conference of the North American Fuzzy Information Processing Society: 461-465.
- Chiu, S. L. (1994). "Fuzzy Model Identification based on Cluster Estimation." Journal of Intelligent and Fuzzy Systems **2**: 267-278.
- Chiu, S. L. (1994a). A Cluster Estimation Method with Extension to Fuzzy Model Identification. Third IEEE Conference on Fuzzy Systems: 1240-1245.
- Chiu, S. L. (1996). "Selecting Input Variables for Fuzzy Models." Journal of Intelligent and Fuzzy Systems **4**: 243-256.
- CIBSE (1997). Natural Ventilation in Non-domestic Buildings. CIBSE.
- CIBSE (2000). CIBSE Guide Volume H: Building Control Systems.
- Clarke, D. W., C. Mohtadi and P. S. Tuffs (1987a). "Generalized Predictive Control--Part I. The Basic Algorithm." Automatica **23**(2): 137-148.
- Clarke, D. W., C. Mohtadi and P. S. Tuffs (1987b). "Generalized Predictive Control--Part II. Extensions and Interpretations." Automatica **23**(2): 149-160.
- Coley, D. A. and J. M. Penman (1992). "Second Order System Identification in the Thermal Response of Real Buildings. Paper II: Recursive Formulation for On-Line Building Energy Management and Control." Building and Environment **27**(3): 269-277.
- Collobert, R. and S. Bongio (2001). "SVM Torch: Support Vector Machines for Large-Scale Regression Problems." Journal of Machine Learning Research **1**: 143-160.
- Cui, X. and K. G. Shin (1993). "Direct Control and Coordination Using Neural Networks." IEEE Transactions on Systems, Man, and Cybernetics **23**(3): 686-697.
- Curtiss, P. S. (1992). Artificial Neural Networks for Use in Building Systems Control and Energy Management. Ph.D., University of Colorado at Boulder, Boulder, CO.
- Curtiss, P. S. (1996). "Experimental Results from a Network-Assisted PID Controller." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **102**(1): 1157-1168.
- Curtiss, P. S., M. J. Brandemuehl and J. F. Kreider (1994). "Energy Management in Central HVAC Plants Using Neural Networks." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **100**(1): 476-493.
- Curtiss, P. S., J. F. Kreider and M. J. Brandemuehl (1993a). "Adaptive Control of HVAC Processes Using Predictive Neural Networks." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **99**(1): 496-504.
- Curtiss, P. S., J. F. Kreider and M. J. Brandemuehl (1993b). Artificial Neural Networks Proof of Concept for Local and Global Control of Commercial Building HVAC Systems. Joint Solar Engineering Conference, ASME: 429-443.
- Curtiss, P. S., G. Shavit and J. F. Kreider (1996). "Neural Networks Applied to Buildings--A Tutorial and Case Studies in Prediction and Adaptive Control." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **102**(1): 1141-1146.
- Daly, A. (2002). "Operable Windows and HVAC Systems." HPAC Engineering(12): 22-30.
- Delsante, A. and T. A. Vik (2001). Hybrid Ventilation: State-of-the-Art Review. IEA Annex 35. IEA.
- Dexter, A. L. and P. Haves (1989). "A Robust Self-Tuning Predictive Controller for HVAC Applications." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **95**(2): 431-438.
- Dodier, R. H. and G. P. Henze (2004). "Statistical Analysis of Neural Networks as Applied to Building Energy Prediction." Journal of Solar Energy Engineering **126**(1): 592-600.
- DOE <http://www.eia.doe.gov/emeu/aer/txt/ptb0201a.html>.
- Dols, W. S. and G. N. Walton (2002). CONTAMW 2.0 User Manual. NISTIR 6921. Building and Fire Research Laboratory, National Institute of Standards and Technology.

- Dounis, A. I., M. Bruant, G. Guarraccino, P. Michel and M. Santamouris (1996). "Indoor Air-Quality Control by a Fuzzy-Reasoning Machine in Naturally Ventilated Buildings." Applied Energy **54**(1): 11-28.
- Dounis, A. I., M. Santamouris, C. C. Lefas and D. E. Manolakis (1994). "Thermal-Comfort Degradation by a Visual Comfort Fuzzy-Reasoning Machine Under Natural Ventilation." Applied Energy **48**: 115-130.
- Duffie, J. A. and W. A. Beckman (1991). Solar Engineering of Thermal Processes. New York, John Wiley & Sons, Inc.
- Economou, C. G. and M. Morari (1986). "Internal Model Control. 5. Extension to Nonlinear Systems." Ind. Eng. Chem. Process Des. Dev. **25**: 403-411.
- EDSL (1996). <http://ourworld.compuserve.com/homepages/edsl/designin.htm>.
- EIA (1999). 1999 Commercial Buildings Energy Consumption Survey: Building Characteristics Tables. Table B5, Census Region and Division, Floorspace, 1999. Energy Information Administration.
- Engel, Y., S. Mannor and R. Meir (2002). Sparse Online Greedy Support Vector Regression. 13th European Conference on Machine Learning, Helsinki: 84-96.
- Engel, Y., S. Mannor and R. Meir (2004). "The Kernel Recursive Least Squares Algorithm." IEEE Trans. Signal Processing: In press.
- Engel, Y., S. Mannor and R. Meir (Unpublished). "Kernel Recursive Least Squares Regression: Online Sparsification, Kernel PCA and Loss Bounds."
- Erbs, D. G., S. A. Klein and J. A. Duffie (1982). "Estimation of the Diffuse Radiation Fraction for Hourly, Daily, and Monthly-Average Global Radiation." Solar Energy **28**: 293.
- ESRU (1997). The ESP-r System for Building Energy Simulation, User Guide Version 9. Energy System Research Unit, University of Strathclyde.
- Etheridge, D. and M. Sandberg (1996). Building Ventilation: Theory and Measurement. Chichester, John Wiley and Sons.
- Ferguson, M. N. and J.-L. Scartezini (1992). "Evaluation of an Optimal Stochastic Controller in a Full-scale Experiment." Energy and Buildings **18**: 1-10.
- Ferreira, P. M., E. A. Faria and A. E. Ruano (2002). "Neural Network Models in Greenhouse Air Temperature Prediction." Neurocomputing **43**(1-4): 51-75.
- Ferreira, P. M. and A. E. Ruano (2002). Choice of RBF Model Structure for Predicting Greenhouse Inside Air Temperature. IFAC 15th Triennial World Congress: 91-96.
- Ferreira, P. M., A. E. Ruano and C. M. Fonseca (2003). Genetic Assisted Selection of RBF Model Structures for Greenhouse Inside Air Temperature Prediction. 2003 IEEE Conference on Control Applications: 576-581.
- Feustel, H. E. and J. Dieris (1992). "A Survey of Airflow Models for Multizone Structures." Energy and Buildings **18**: 79-100.
- Feuston, B. P. and J. H. Thurtell (1994). "Generalized Nonlinear Regression with Ensemble of Neural Nets: The Great Energy Predictor Shootout." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **100**(2): 1075-1080.
- Filliben, J. J. and K. C. Li (1997). "A Systematic Approach to the Analysis of Complex Interaction Patterns in Two-Level Factorial Designs." Technometrics **39**(3): 286-297.
- Fletcher, J. (1996). Night Cooling Control Strategies. TA 14/96. Building Services Research and Information Association.
- Fonseca, C. M. and P. J. Fleming (1998a). "Multiobjective Optimization and Multiple Constraint Handling with Evolutionary Algorithms--Part I: a Unified Formulation." IEEE Transactions on Systems, Man, and Cybernetics--Part A: Systems and Humans **28**(1): 26-37.
- Fonseca, C. M. and P. J. Fleming (1998b). "Multiobjective Optimization and Multiple Constraint Handling with Evolutionary Algorithms--Part II: Application Example." IEEE Transactions on Systems, Man, and Cybernetics--Part A: Systems and Humans **28**(1): 38-47.
- Geros, V., M. Santamouris and G. Guarraccino (1997). Applying Night Ventilation Techniques in Office Buildings. 18th AIVC Conference, "Ventilation and Cooling", Athens, Greece: 467-475.

- Ghiaus, C. and F. Allard (2001). "Fuzzy Modeling and Controlling of a Fan-Coil." International Journal of Solar Energy **21**: 131-145.
- Gill, P. E., W. Murray and M. G. Wright (1981). Practical Optimization. London, Academic Press.
- Gunn, S. R. (1998). Support Vector Machines for Classification and Regression. University of Southampton.
- Heiselberg, P., Ed. (2002). Principles of Hybrid Ventilation, IEA Annex 35.
- Heiselberg, P. and P. O. Tjeflaat (1999). Design Procedure for Hybrid Ventilation. HybVent Forum '99, Sydney, IEA Annex 35.
- Henze, G. P. (2003). "An Overview of Optimal Control for Central Cooling Plants with Ice Thermal Energy Storage." Journal of Solar Energy Engineering **125**(3): 302-309.
- Henze, G. P., R. H. Dodier and M. Krarti (1997). "Development of a Predictive Optimal Controller for Thermal Energy Storage Systems." International Journal of HVAC&R Research, American Society of Heating, Refrigeration, and Air-Conditioning Engineers, Atlanta, Georgia **3**(3): 233-264.
- Henze, G. P., C. Felsmann, D. E. Kalz and G. Knabe (2004a). "Impact of Forecasting Accuracy on Predictive Optimal Control of Active and Passive Building Thermal Storage Inventory." International Journal of HVAC&R Research **10**(2): 153-178.
- Henze, G. P., C. Felsmann and G. Knabe (2004b). "Evaluation of Optimal Control for Active and Passive Building Thermal Storage." International Journal of Thermal Sciences **43**: 173-183.
- Henze, G. P. and R. E. Hindman (2002). "Control of Air-Cooled Chiller Condensers using Clustering Neural Networks." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **108**(2).
- Herrlin, M. K. and F. Allard (1992). "Solution Methods for the Air Balance in Multizone Buildings." Energy and Buildings **18**: 159-170.
- Iijima, M., K. Takagi, R. Takeuchi and T. Matsumoto (1994). "A Piecewise-Linear Regression on the ASHRAE Time-Series Data." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **100**(2): 1088-1095.
- Incropera, F. P. and D. P. DeWitt (1996). Introduction to Heat Transfer. New York, John Wiley & Sons.
- Irving, S. J., P. J. Conccannon and H. S. Dhargalkar (1995). Sizing and Location of Passive Ventilation Openings. ETSU Report S/N5/00142. Energy Technology Support Unit.
- Jaakkola, T. S. (2001). Personal Correspondence.
- Jang, J.-S. R., C.-T. Sun and E. Mizutani (1997). Neuro-Fuzzy and Soft Computing: A Computational Approach to Learning and Machine Intelligence. Upper Saddle River, NJ, Prentice Hall.
- Jeannette, E., K. Assawamartbunlue, P. S. Curtiss and J. F. Kreider (1998). "Experimental Results of a Predictive Neural Network HVAC Controller." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **104**(2): 192-197.
- Jeong, Y. and F. Haghghat (2002). Modeling of a Hybrid Ventilated Building "Grong School". Hybrid Ventilation 2002 Fourth International Forum, Montreal: 198-207.
- Jordan, M. I. (1994). "Hierarchical Mixtures of Experts and the EM Algorithm." Neural Computation **6**: 181-214.
- Kalogirou, S. A., M. M. Eftekhari and D. J. Pinnock (2001). "Artificial Neural Networks for Predicting Air Flow in a Naturally Ventilated Test Room." Building Services Engineering Research and Technology **22**(2): 83-93.
- Kammerud, R., E. Ceballos, B. Curtis, W. Place and B. Andersson (1984). "Ventilation Cooling of Residential Buildings." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **90**(1b): 226-252.
- Karnopp, D. C., D. I. Margolis and R. C. Rosenberg (1990). System Dynamics: A Unified Approach. New York, John Wiley & Sons, Inc.
- Kasahara, M., T. Matsuba, Y. Hashimoto, I. Murasawa, A. Kimbara, K. Kamimura and S. Kurosu (1998). "Optimal Preview Control for HVAC System." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **104**(1a): 502-513.

- Kawashima, M. (1994). "Artificial Neural Network Backpropagation Model with Three-Phase Annealing Developed for the Building Energy Predictor Shootout." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **100**(2): 1096-1103.
- Kawashima, M., C. E. Dorgan and J. W. Mitchell (1995). "Hourly Thermal Load Prediction for the Next 24 Hours by ARIMA, EWMA, LR, and an Artificial Neural Network." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **101**(1): 186-200.
- Kawashima, M., C. E. Dorgan and J. W. Mitchell (1996). "Optimizing System Control with Load Prediction by Neural Networks for an Ice-Storage System." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **102**(1): 1169-1178.
- Keeney, K. and J. Braun (1996). "A Simplified Method for Determining Optimal Cooling Control Strategies for Thermal Storage in Building Mass." International Journal of HVAC&R Research **2**(1): 59-78.
- Kindangen, J. I. (1996). "Artificial Neural Networks and Naturally Ventilated Buildings: A Method of Predicting Window Size and Location with Subsequent Effect on Interior Air Motion Using Neural Networks." Building Research and Information **24**(4): 203-208.
- Kintner-Meyer, M. and A. F. Emery (1995). "Optimal control of an HVAC system using cold storage and building thermal capacitance." Energy and Buildings **23**(1): 19-31.
- Klein, S. A. (1990). TRNSYS: A Transient System Simulation Program. Solar Energy Laboratory, University of Wisconsin, Madison.
- Kolokotroni, M. (2001). "Night Ventilation Cooling of Office Buildings: Parametric Analyses of Conceptual Energy Impacts." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **107**(1): 479-489.
- Kolokotroni, M., A. Tindale and S. J. Irving (1997). NiteCool: Office Night Ventilation Pre-Design Tool. 18th AIVC Conference, Athens, Greece: 213-224.
- Kramer, M. A. (1991). "Nonlinear Principal Components Analysis Using Autoassociative Neural Networks." AIChE J. **37**: 233-243.
- Krarti, M. (2003). "An Overview of Artificial Intelligence-Based Methods for Building Energy Systems." Journal of Solar Energy Engineering **125**(3): 331-342.
- Krauss, G., J. I. Kindangen and P. Depecker (1997). "Using artificial neural networks to predict interior velocity coefficients." Building and Environment **32**(4): 295-303.
- Kreider, J. F. and J. S. Haberl (1994). "Predicting Hourly Building Energy Use: The Great Energy Predictor Shootout--Overview and Discussion of Results." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **100**(2): 1104-1118.
- Kreider, J. F., S. Wang, D. Anderson and J. Dow (1992). "Expert Systems, Neural Networks and Artificial Intelligence Applications in Commercial Building HVAC Operations." Automation in Construction **1**(3): 225-238.
- Kummert, M. (2001). Contribution to the Application of Modern Control Techniques to Solar Buildings: Simulation-Based Approach and Experimental Validation. Doctoral, Sciences de l'Environnement, Fondation Universitaire Luxembourgeoise,
- Kummert, M., P. Andre and J. Nicolas (2000). "Optimal Heating Control in a Passive Solar Commercial Building." Solar Energy **69**(Suppl.)(1-6): 103-116.
- Lee, K.-h. and J. E. Braun (2004). Development and Application of an Inverse Building Model for Demand Response in Small Commercial Buildings. SimBuild 2004, IBPSA-USA National Conference, Boulder, CO.
- Leung, H. and Y. Li (2001). Analytical Solutions for Buoyancy-Driven Natural Ventilation in Buildings with Three Openings. The 4th International Conference on Indoor Air Quality, Ventilation and Energy Conservation in Buildings (IAQVEC 2001), Changsha, China.

- Levermore, G. J. (2000). Building Energy Management Systems, E&FN SPON.
- Li, H. and J. E. Braun (2002). On-Line Models for Use in Automated Fault Detection and Diagnosis for HVAC&R Equipment. ACEEE Conference on Energy Efficiency in Buildings, Monterey, CA.
- Li, K. C. (1992). "On Principal Hessian Directions for Data Visualization and Dimension Reduction: Another Application of Stein's Lemma." Journal of the American Statistical Association **87**(420): 1025-1038.
- Li, K. C., H.-h. Lue and C.-H. Chen (2000). "Interactive Tree-Structured Regression via Principal Hessian Directions." Journal of the American Statistical Association **95**(450): 547-560.
- Li, Y. (2000). "Buoyancy-driven Natural Ventilation in a Thermally Stratified One-Zone Building." Building and Environment **35**: 207-214.
- Li, Y. (2002). Analysis of Natural Ventilation--A Summary of Existing Analytical Solutions. Technical Report TR12. IEA Annex 35.
- Li, Y. and A. Delsante (2001). "Natural Ventilation Induced by Combined Wind and Thermal Forces." Building and Environment **36**(1): 59-71.
- Li, Y., A. Delsante and J. Symons (2000). "Prediction of Natural Ventilation in Buildings with Large Openings." Building and Environment **35**: 191-206.
- Liem, S. H. and A. H. C. van Paassen (1997). Hardware and Controls for Natural Ventilation Cooling. 18th AIVC Conference on Ventilation and Cooling, Athens: 60-68.
- Linden, P. F., G. F. Lane-Serff and D. A. Smeed (1990). "Emptying Filling Boxes: The Fluid Mechanics of Natural Ventilation." Journal of Fluid Mechanics **212**: 309-335.
- Linkens, D. A. and M.-Y. Chen (1998). Hierarchical Fuzzy Clustering Based on Self-Organising Networks. The 1998 IEEE International Conference on Fuzzy Systems: 1406-1410.
- Linker, R., P. O. Gutman and I. Seginer (1999). "Robust Controllers for Simultaneous Control of Temperature and CO₂ Concentration in Greenhouses." Control Engineering Practice **7**: 851-862.
- Ljung, L. (1991). System Identification Toolbox User's Guide, The MathWorks.
- Ljung, L. and T. Glad (1994). Modeling of Dynamic Systems. Englewood Cliffs, NJ, PTR Prentice Hall.
- Loh, W.-Y. (2002a). GUIDE User Manual.
- Loh, W.-Y. (2002b). "Regression Trees With Unbiased Variable Selection and Interaction Detection." Statistica Sinica **12**: 361-386.
- Lu, H. and G. G. Yen (2001). Multiobjective Optimization Design via Genetic Algorithm. 2001 IEEE International Conference on Control Applications, Mexico City: 1190-1195.
- Lue, H.-h. (1994). Principal-Hessian-Direction-Based Regression Trees. Ph.D., Mathematics, UCLA, Los Angeles.
- Lute, P. and D. v. Paassen (1995). Optimal Indoor Temperature Control Using a Predictor. IEEE Control Systems Magazine. **15**: 4-10.
- Lute, P. J. (1992). The Use of Predictions in Temperature Control in Buildings: A Passive Climate Application. Ph.D. Thesis, Faculty of Mechanical Engineering and Marine Technology, Delft University of Technology, Delft.
- MacKay, D. J. C. (1994). "Bayesian Nonlinear Modeling for the Prediction Competition." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **100**(2): 1053-1062.
- Mackey, M. and L. Glass (1977). "Oscillation and Chaos in Physiological Control Systems." Science **197**(287-289).
- Madsen, H. and J. Holst (1995). "Estimation of Continuous-Time Models for the Heat Dynamics of a Building." Energy and Buildings **22**: 67-79.
- Martin, A. J. (1995). Control of Natural Ventilation. Building Services Research and Information Association.
- Massie, D. D., P. S. Curtiss and J. F. Kreider (1998). "Predicting Central Plant HVAC Equipment Performance Using Neural Networks--Laboratory System Test Results." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **104**(1a): 221-228.

- McQuiston, F. C. and J. D. Parker (1994). Heating, Ventilating, and Air Conditioning Analysis and Design, John Wiley and Sons, Inc.
- Meteotest http://www.meteotest.ch/en/mn_home?w=ber.
- Moody, J. and C. Darken (1989). "Fast Learning in Networks of Locally-Tuned Processing Units." Neural Computation **1**: 281-294.
- Morel, N., M. Bauer, M. El-Khoury and J. Krauss (2001). "Neurobat, a Predictive and Adaptive Heating Control System Using Artificial Neural Networks." International Journal of Solar Energy **21**: 161-201.
- Morris, F. B., J. E. Braun and S. J. Treado (1994). "Experimental and Simulated Performance of Optimal Control of Building Thermal Storage." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **100**(1): 402-414.
- Mozer, M. C., L. Vidmar and R. H. Dodier (1996). The Neurothermostat: Predictive Optimal Control of Residential Heating Systems. Advances in Neural Information Processing Systems 1996, MIT Press: 953-959.
- Mueller, K.-R., A. Smola, G. Raetsch, B. Schoelkopf, J. Kohlmorgen and V. Vapnik (1997). Predicting Time Series with Support Vector Machines. International Conference on Artificial Neural Networks: 999-1004.
- Mueller, K.-R., A. Smola, G. Raetsch, B. Schoelkopf, J. Kohlmorgen and V. Vapnik (1999). Using Support Vector Machines for Time Series Prediction. Advances in Kernel Methods--Support Vector Learning. B. Schoelkopf, C. J. C. Burges and A. J. Smola. Cambridge, MIT Press: 243-253.
- Nameroff, S. (1989). QuickBASIC: The Complete Reference. Berkeley, CA, McGraw-Hill.
- NatVent (1999). Overcoming Barriers to Natural Ventilation (CD-ROM).
- NREL http://rredc.nrel.gov/solar/old_data/nsrdb/tmy2/.
- Ohlsson, M. B. O., C. O. Peterson, H. Pi, T. S. Roegnvaldsson and B. P. W. Soederberg (1994). "Predicting System Loads with Artificial Neural Networks--Methods and Results from "The Great Energy Predictor Shootout"." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **100**(2): 1063-1074.
- Oja, E. (1991). Data Compression, Feature Extraction and Autoassociation in Feed-Forward Neural Networks. Artificial Neural Networks. Amsterdam, Elsevier: 737-745.
- Oppenheim, A. V., R. W. Schaffer and J. R. Buck (1999). Discrete-Time Signal Processing. Upper Saddle River, NJ, Prentice Hall.
- Orgill, J. F. and K. G. Hollands (1997). "Correlation Equation for Hourly Diffuse Radiation on a Horizontal Surface." Solar Energy **19**: 357-359.
- Pareto, V. (1906). Manuale di Economia Politica.
- Portney, L. G. and M. P. Watkins (1993). Foundations of Clinical Research: Applications to Practice. Norwalk, CT, Appleton & Lange.
- Rao, S. (1984). Optimization. New Delhi, Wiley Eastern Limited.
- Rink, R. E., V. G. Gourishankar and M. Zaheer-Uddin (1988). "Optimal Control of Heat-Pump/Heat-Storage Systems with Time-of-Day Energy Price Incentive." Journal of Optimization Theory and Applications **58**(1): 93-108.
- Santamouris, M., A. Argiriou, D. Asimakopoulos, N. Klitsikas and A. I. Dounis (1995). "Heat and Mass Transfer Through Large Openings by Natural Convection." Energy and Buildings **23**: 1-8.
- Sebald, A. V. and J. Schlenzig (1994). "Minimax Design of Neural Net Controllers for Highly Uncertain Plants." IEEE Transactions on Neural Networks **5**(1): 73-82.
- Seem, J. E. (1998). "A New Pattern Recognition Adaptive Controller with Application to HVAC Systems." Automatica **34**(8): 969-982.
- Seginer, I. (1997). "Some Artificial Neural Network Applications to Greenhouse Environmental Control." Computers and Electronics in Agriculture **18**: 167-186.
- Seginer, I., T. Boulard and B. J. Bailey (1994). "Neural Network Models of the Greenhouse Climate." Journal of Agricultural Engineering Research **59**(3): 203-216.

- Seginer, I., Y. Hwang, T. Boulard and J. W. Jones (1996). "Mimicking an Expert Greenhouse Grower with a Neural-Net Policy." Transactions of the ASAE **39**(1): 299-306.
- Sheen, D. and S. R. T. Kumara (1993). Adaptive Neural Network Control Schemes for Unknown Nonlinear Dynamic Systems. Intelligent Engineering Systems Through Artificial Neural Networks. C. H. Dagli, L. I. Burke, B. R. Fernandez and J. Ghosh. **3**: 535-540.
- Skrjanc, I., B. Zupancic, B. Furlan and A. Krainer (2001). "Theoretical and Experimental FUZZY Modelling of Building Thermal Dynamic Response." Building and Environment **36**: 1023-1038.
- Smola, A. J. and B. Schoelkopf (1998). A Tutorial on Support Vector Regression.
- Soeterboek, R. (1992). Predictive Control. New York, Prentice Hall.
- Sousa, J. M., R. Babuska and H. B. Verbruggen (1997a). Internal Model Control with a Fuzzy Model: Application to an Air-Conditioning System. Sixth IEEE International Conference on Fuzzy Systems: 207-212.
- Sousa, J. M., R. Babuska and H. B. Verbruggen (1997b). "Fuzzy Predictive Control Applied to an Air-Conditioning System." Control Engineering Practice **5**(10): 1395-1406.
- Specht, D. F. (1991). "A General Regression Neural Network." IEEE Transactions on Neural Networks **2**(6): 568-576.
- Spindler, H. (1998). Residential Building Energy Analysis: Development and Uncertainty Analysis of a Simplified Model. Masters Thesis, Building Technology, MIT, Cambridge.
- Spindler, H., L. Glicksman and L. Norford (2002). The Potential for Natural and Hybrid Cooling Strategies to Reduce Cooling Energy Consumption in the United States. RoomVent 2002, Copenhagen: 517-520.
- Stein, C. (1981). "Estimation of the Mean of a Multivariate Normal Distribution." The Annals of Statistics **9**: 1135-1151.
- Stevenson, W. J. (1994). "Using Artificial Neural Nets to Predict Building Energy Parameters." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **100**(2): 1081-1087.
- Swami, M. V. and S. Chandra (1988). "Correlations for Pressure Distribution on Buildings and Calculation of Natural-ventilation Airflow." ASHRAE Transactions, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta, Georgia **94**(1): 243-266.
- Teeter, J. and M.-Y. Chow (1998). "Application of Functional Link Neural Network to HVAC Thermal Dynamic System Identification." IEEE Transactions on Industrial Electronics **45**(1): 170-176.
- Tindale, A. (1993). "Third Order Lumped-parameter Simulation Method." BSER&T **14**(3): 87-97.
- Townsend, M. A., D. B. Charchas and A. Abdelmessih (1986). "Optimal Control of a General Environmental Space." Transactions of the ASME, Journal of Dynamic Systems, Measurement, and Control **108**: 330-339.
- van der Aa, A. (2002). Hybrid Ventilation Waterland School Building, The Netherlands--First Results of the Monitoring Phase. Hybrid Ventilation 2002 Fourth International Forum, Montreal: 79-86.
- van Paassen, A. H. C., S. H. Liem and B. P. Groeninger (1998). Control of Night Cooling with Natural Ventilation: Sensitivity Analysis of Control Strategies and Vent Openings. 19th Annual AIVC Conference, Oslo.
- van Paassen, A. H. C. and P. J. Lute (1993). Energy Saving through Controlled Ventilation Windows. 3rd European Conference on Architecture, Florence, Italy: 208-211.
- Vapnik, V. (1995). The Nature of Statistical Learning, Springer-Verlag.
- Walker, C., S. Manchanda, H. Spindler and L. Norford (2004). Building Performance: Analysis of Naturally Ventilated UK Office Building. RoomVent 2004, Coimbra.
- Weigend, A. S. and N. A. Gershenfeld, Eds. (1994). Time Series Prediction, Addison Wesley.
- Wright, J. A. and H. Loosemore (2001). The Multi-criterion Optimization of Building Thermal Design and Control. Seventh International IBPSA Conference, Rio de Janeiro: 873-880.
- Wunsch, C. (1996). The Ocean Circulation Inverse Problem, Cambridge University Press.
- Yam, J., Y. Li and Z. Zheng (2003). "Nonlinear Coupling Between Thermal Mass and Natural Ventilation in Buildings." International Journal of Heat and Mass Transfer **46**: 1251-1264.

- Yoshida, H. and H. Yamaguti (2001). Optimal Operation of a HVAC System with a Thermal Storage Water Tank. Seventh International IBPSA Conference, Rio de Janeiro: 1249-1256.
- Zaheer-Uddin, M. (1992). "Optimal Control of a Single Zone Environmental Space." Building and Environment **27**(1): 93-103.
- Zaheer-Uddin, M., V. G. Gourishankar and R. E. Rink (1988). "Dynamic Suboptimal Control of a Heat Pump/Heat Storage System." Optimal Control Applications and Methods **9**: 341-355.

Appendix 1 Electrical Logger Program Files

The two sections of this appendix contain the program files uploaded to the C180E (Section A1.1) and the K20 (Section A1.2).

A1.1 C180E (P1630.001)

// Free form comments up to 3k in size OK

RECORDER_TYPE

DESCRIP: C180
 MODEL: K20-6
 SERIAL: 1630

RECORDER_INFO

PSID: 1
 MIN: 1
 PSDESC:
 RINGS: 2
 CUTOFF: 0
 OPTIONA: 0
 OPTIONB: 0

K20_CT_TABLE

CH	DESCRIP	AMPS	VH	VL	VMULT	VLT	AMP	DLT	PW
0	Primary_A	100.0	A1	N1	1.0	ON	ON	OFF	0
1	Primary_B	100.0	B1	N1	1.0	ON	ON	OFF	1
2	Panel_Outlet	20.00	B1	N1	1.0	OFF	ON	OFF	2
3	Parking_Outside	20.00	B1	N1	1.0	OFF	ON	OFF	3
4	Networks	20.00	A1	N1	1.0	OFF	ON	OFF	4
5	Secondary_A	100.0	A1	N1	1.0	OFF	ON	OFF	5
6	Secondary_B	100.0	B1	N1	1.0	OFF	ON	OFF	6
7	UpstairsMain_A	100.0	A1	N1	1.0	OFF	ON	OFF	7
8	UpstairsMain_B	100.0	B1	N1	1.0	OFF	ON	OFF	8
9	Fans_Clivus	20.00	B1	N1	1.0	OFF	ON	OFF	9
10	HotWater	40.00	A1	B1	1.0	OFF	ON	OFF	10
11	Upstairs_Light	20.00	B1	N1	1.0	OFF	ON	OFF	11
12	Lts/Plug_Storage	20.00	B1	N1	1.0	OFF	ON	OFF	12
13	Lts_Storage	20.00	A1	N1	1.0	OFF	ON	OFF	13
14	Bath_Recept	20.00	A1	N1	1.0	OFF	ON	OFF	14
15	Attic	20.00	A1	N1	1.0	OFF	ON	OFF	15

K20_PW_TABLE

PW	DESCRIP	KW	KWH	KVA	KVH
0	Primary_A	ON	OFF	OFF	OFF
1	Primary_B	ON	OFF	OFF	OFF
2	Panel_Outlet	ON	OFF	OFF	OFF
3	Parking_Outside	ON	OFF	OFF	OFF
4	Networks	ON	OFF	OFF	OFF
5	Secondary_A	ON	OFF	OFF	OFF
6	Secondary_B	ON	OFF	OFF	OFF

7	UpstairsMain_A	ON	OFF	OFF	OFF
8	UpstairsMain_B	ON	OFF	OFF	OFF
9	Fan_Clivus	ON	OFF	OFF	OFF
10	HotWater	ON	OFF	OFF	OFF
11	Upstairs_Light	ON	OFF	OFF	OFF
12	Lts/Plug_Storage	ON	OFF	OFF	OFF
13	Lts_Storage	ON	OFF	OFF	OFF
14	Bath_Recept	ON	OFF	OFF	OFF
15	Attic	ON	OFF	OFF	OFF

K20_AN_TABLE

CH	DESCRIP	TYPE	SCALE	OFFSET	UNITS	SNAP	TSR	TYP	LCH	INV
0		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
1		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
2		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
3		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
4		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
5		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
6		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
7		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
8		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
9		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
10		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
11		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
12		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
13		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
14		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
15		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF

K20_DI_TABLE

CH	DESCRIP	SCALE	OFFSET	UNITS	METHOD	TSR
0		1.0	0.0		COUNT	OFF
1		1.0	0.0		COUNT	OFF
2		1.0	0.0		COUNT	OFF
3		1.0	0.0		COUNT	OFF
4		1.0	0.0		COUNT	OFF
5		1.0	0.0		COUNT	OFF
6		1.0	0.0		COUNT	OFF
7		1.0	0.0		COUNT	OFF
8		1.0	0.0		COUNT	OFF
9		1.0	0.0		COUNT	OFF
10		1.0	0.0		COUNT	OFF
11		1.0	0.0		COUNT	OFF
12		1.0	0.0		COUNT	OFF
13		1.0	0.0		COUNT	OFF
14		1.0	0.0		COUNT	OFF
15		1.0	0.0		COUNT	OFF

A1.2 K20 (P10046.001)

// Free form comments up to 3k in size OK

RECORDER_TYPE

DESCRIP: C180
 MODEL: K20-6
 SERIAL: 1630

RECORDER_INFO

PSID: 1
 MIN: 1
 PSDESC:
 RINGS: 2
 CUTOFF: 0
 OPTIONA: 0
 OPTIONB: 0

K20_CT_TABLE

CH	DESCRIP	AMPS	VH	VL	VMULT	VLT	AMP	DLT	PW
0	Primary_A	100.0	A1	N1	1.0	ON	ON	OFF	0
1	Primary_B	100.0	B1	N1	1.0	ON	ON	OFF	1
2	Panel_Outlet	20.00	B1	N1	1.0	OFF	ON	OFF	2
3	Parking_Outside	20.00	B1	N1	1.0	OFF	ON	OFF	3
4	Networks	20.00	A1	N1	1.0	OFF	ON	OFF	4
5	Secondary_A	100.0	A1	N1	1.0	OFF	ON	OFF	5
6	Secondary_B	100.0	B1	N1	1.0	OFF	ON	OFF	6
7	UpstairsMain_A	100.0	A1	N1	1.0	OFF	ON	OFF	7
8	UpstairsMain_B	100.0	B1	N1	1.0	OFF	ON	OFF	8
9	Fans_Clivus	20.00	B1	N1	1.0	OFF	ON	OFF	9
10	HotWater	40.00	A1	B1	1.0	OFF	ON	OFF	10
11	Upstairs_Light	20.00	B1	N1	1.0	OFF	ON	OFF	11
12	Lts/Plug_Storage	20.00	B1	N1	1.0	OFF	ON	OFF	12
13	Lts_Storage	20.00	A1	N1	1.0	OFF	ON	OFF	13
14	Bath_Recept	20.00	A1	N1	1.0	OFF	ON	OFF	14
15	Attic	20.00	A1	N1	1.0	OFF	ON	OFF	15

K20_PW_TABLE

PW	DESCRIP	KW	KWH	KVA	KVH
0	Primary_A	ON	OFF	OFF	OFF
1	Primary_B	ON	OFF	OFF	OFF
2	Panel_Outlet	ON	OFF	OFF	OFF
3	Parking_Outside	ON	OFF	OFF	OFF
4	Networks	ON	OFF	OFF	OFF
5	Secondary_A	ON	OFF	OFF	OFF
6	Secondary_B	ON	OFF	OFF	OFF
7	UpstairsMain_A	ON	OFF	OFF	OFF
8	UpstairsMain_B	ON	OFF	OFF	OFF
9	Fan_Clivus	ON	OFF	OFF	OFF
10	HotWater	ON	OFF	OFF	OFF
11	Upstairs_Light	ON	OFF	OFF	OFF

12	Lts/Plug_Storage	ON	OFF	OFF	OFF
13	Lts_Storage	ON	OFF	OFF	OFF
14	Bath_Recept	ON	OFF	OFF	OFF
15	Attic	ON	OFF	OFF	OFF

K20_AN_TABLE

CH	DESCRIP	TYPE	SCALE	OFFSET	UNITS	SNAP	TSR	TYP	LCH	INV
0		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
1		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
2		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
3		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
4		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
5		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
6		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
7		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
8		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
9		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
10		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
11		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
12		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
13		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
14		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF
15		UNUSED	1.0	0.0		OFF	OFF	NA	0	OFF

K20_DI_TABLE

CH	DESCRIP	SCALE	OFFSET	UNITS	METHOD	TSR
0		1.0	0.0		COUNT	OFF
1		1.0	0.0		COUNT	OFF
2		1.0	0.0		COUNT	OFF
3		1.0	0.0		COUNT	OFF
4		1.0	0.0		COUNT	OFF
5		1.0	0.0		COUNT	OFF
6		1.0	0.0		COUNT	OFF
7		1.0	0.0		COUNT	OFF
8		1.0	0.0		COUNT	OFF
9		1.0	0.0		COUNT	OFF
10		1.0	0.0		COUNT	OFF
11		1.0	0.0		COUNT	OFF
12		1.0	0.0		COUNT	OFF
13		1.0	0.0		COUNT	OFF
14		1.0	0.0		COUNT	OFF
15		1.0	0.0		COUNT	OFF

Appendix 2 Campbell Scientific, Inc. Data Logger Usage Details

The two sections of this appendix include wiring details (Section A2.1) and programs (Section A2.2) for the 21X data logger.

A2.1 Wiring

The following table catalogues the wiring connections to the 21X. Locations of ground attachments are generally omitted for clarity. Refer to the photograph of the 21X logger in Figure 4-35 for a view of the ports used.

21X Port	Device	Wire color	Notes
<i>Analog Voltage Input</i>			
1H	Assembly Thermocouple (1)	blue	red to ground
1L	Mass Thermocouple (2)	blue	red to ground
2H	Louvered Door E Thermocouple (3)	blue	red to ground
2L	Louvered Door W Thermocouple (4)	blue	red to ground
3H	Attic Opening Thermocouple (5)	blue	red to ground
3L	Attic 1 Thermocouple (6)	blue	red to ground
4H	Attic 2 Thermocouple (7)	blue	red to ground
4L	Sunspace Thermocouple (8)	blue	red to ground
5H	Basement Thermocouple (9)	blue	red to ground
5L	Outside Thermocouple (10)	blue	red to ground
6H	North Pyranometer		other to ground
6L	East Pyranometer		other to ground
7H	South Pyranometer		other to ground
7L	West Pyranometer		other to ground
8H	Horizontal Pyranometer		other to ground
8L	Potentiometer voltage (wind vane)	green	red to ground
<i>Analog Voltage Output</i>			
CAO1	Voltage to fan controller card		
CAO2	Voltage to wind vane	yellow	
<i>Pulse Input</i>			
Pulse Input 1	Anemometer	black	
<i>Ground</i>	House ground connected to ground port adjacent to +12V input		

Table 1-1 Wiring Table for the Campbell Scientific, Inc. 21X data logger. For the pyranometer connections, one wire was connected to the 21X port indicated, and the other was connected to ground.

A2.2 Code Listing of Program Loaded onto 21X

Provided in this section are the commented and more readable file (BRDMR.CSI) as well as the actual file (BRDMR.DLD) uploaded to the logger. The files are two representations of the identical set of instructions.

A2.2.1 BRDMR.CSI

```
{21X}
;
*Table 1 Program
01: 1      Execution Interval (seconds)

; mod 60 seconds in 1/10 of second intervals
1: Time (P18)
1: 00      Option
2: 600     Mod/By
3: 15      Loc [ time   ]

2: Internal Temperature (P17)
1: 1       Loc [ Tref   ]

;first time around, initialize param 3
3: If Flag/Input (P91)
1: 21      Do if Flag 1 is Low
2: 30      Then Do

4: Z=F (P30)
1: 0.0     F
2: 3       Z Loc [ CtrlVolt ]

5: Z=F (P30)
1: 2500    F
2: 4       Z Loc [ WindPot ]

6: If Flag/Input (P91)
1: 21      Do if Flag 1 is Low
2: 11      Set Flag 1 High ;

7: End (P95)

8: Analog Out (P21)
1: 1       CAO Channel
2: 3       mV Loc [ CtrlVolt ]

;potentiometer voltage supply for wind dir
9: Analog Out (P21)
1: 2       CAO Channel
2: 4       mV Loc [ WindPot ]

;T_common
10: Thermocouple Temp (SE) (P13)
1: 1       Reps
2: 1       5 mV Slow Range
3: 1       SE Channel
4: 1       Type T (Copper-Constantan)
5: 1       Ref Temp (Deg. C) Loc [ Tref   ]
6: 2       Loc [ T_in   ]
7: 1.0     Mult
8: 0.0     Offset

;Tmass
11: Thermocouple Temp (SE) (P13)
1: 1       Reps
2: 1       5 mV Slow Range
3: 2       SE Channel
4: 1       Type T (Copper-Constantan)
5: 1       Ref Temp (Deg. C) Loc [ Tref   ]
6: 16      Loc [ T_mass ]
7: 1.0     Mult
8: 0.0     Offset

;T_inletE
12: Thermocouple Temp (SE) (P13)
1: 1       Reps
2: 1       5 mV Slow Range
3: 3       SE Channel
4: 1       Type T (Copper-Constantan)
5: 1       Ref Temp (Deg. C) Loc [ Tref   ]
6: 17      Loc [ T_inletE ]
7: 1.0     Mult
8: 0.0     Offset
```

```

;T_inletW
13: Thermocouple Temp (SE) (P13)
  1: 1    Reps
  2: 1    5 mV Slow Range
  3: 4    SE Channel
  4: 1    Type T (Copper-Constantan)
  5: 1    Ref Temp (Deg. C) Loc [ Tref    ]
  6: 18   Loc [ T_inletW ]
  7: 1.0  Mult
  8: 0.0  Offset

```

```

;T_inletAttic
14: Thermocouple Temp (SE) (P13)
  1: 1    Reps
  2: 1    5 mV Slow Range
  3: 5    SE Channel
  4: 1    Type T (Copper-Constantan)
  5: 1    Ref Temp (Deg. C) Loc [ Tref    ]
  6: 19   Loc [ T_inletAt ]
  7: 1.0  Mult
  8: 0.0  Offset

```

```

;T_attic1
15: Thermocouple Temp (SE) (P13)
  1: 1    Reps
  2: 1    5 mV Slow Range
  3: 6    SE Channel
  4: 1    Type T (Copper-Constantan)
  5: 1    Ref Temp (Deg. C) Loc [ Tref    ]
  6: 20   Loc [ T_att1   ]
  7: 1.0  Mult
  8: 0.0  Offset

```

```

;T_attic2
16: Thermocouple Temp (SE) (P13)
  1: 1    Reps
  2: 1    5 mV Slow Range
  3: 7    SE Channel
  4: 1    Type T (Copper-Constantan)
  5: 1    Ref Temp (Deg. C) Loc [ Tref    ]
  6: 21   Loc [ T_att2   ]
  7: 1.0  Mult
  8: 0.0  Offset

```

```

;T_sunspace
17: Thermocouple Temp (SE) (P13)

```

```

  1: 1    Reps
  2: 1    5 mV Slow Range
  3: 8    SE Channel
  4: 1    Type T (Copper-Constantan)
  5: 1    Ref Temp (Deg. C) Loc [ Tref    ]
  6: 22   Loc [ T_sunsp  ]
  7: 1.0  Mult
  8: 0.0  Offset

```

```

;T_bsmt
18: Thermocouple Temp (SE) (P13)
  1: 1    Reps
  2: 1    5 mV Slow Range
  3: 9    SE Channel
  4: 1    Type T (Copper-Constantan)
  5: 1    Ref Temp (Deg. C) Loc [ Tref    ]
  6: 23   Loc [ T_bsmt   ]
  7: 1.0  Mult
  8: 0.0  Offset

```

```

;Tout
19: Thermocouple Temp (SE) (P13)
  1: 1    Reps
  2: 1    5 mV Slow Range
  3: 10   SE Channel
  4: 1    Type T (Copper-Constantan)
  5: 1    Ref Temp (Deg. C) Loc [ Tref    ]
  6: 24   Loc [ T_out    ]
  7: 1.0  Mult
  8: 0.0  Offset

```

```

;10: Thermocouple Temp (DIFF) (P14)
; 1: 1    Reps
; 2: 1    5 mV Slow Range
; 3: 1    DIFF Channel
; 4: 1    Type T (Copper-Constantan)
; 5: 1    Ref Temp (Deg. C) Loc [ Tref    ]
; 6: 2    Loc [ T_in     ]
; 7: 1.0  Mult
; 8: 0.0  Offset

```

```

;windspeed measurement
; multiplier is 1/59.66
20: Pulse (P3)
  1: 1    Reps
  2: 1    Pulse Input Channel
  3: 2    Switch Closure, All Counts
  4: 5    Loc [ windspd  ]

```

```

5: .01676 Mult
6: 0.0 Offset

;wind direction measurement
21: Volt (SE) (P1)
1: 1 Reps
2: 15 5000 mV Fast Range
3: 16 SE Channel
4: 6 Loc [ winddir ]
5: .144 Mult
6: 0.0 Offset

;calculate the sin and cosine of the wind
direction
22: Z=SIN(X) (P48)
1: 6 X Loc [ winddir ]
2: 7 Z Loc [ sinwindr ]

;loc 80 holds this temporary value
23: Z=X+F (P34)
1: 6 X Loc [ winddir ]
2: 90 F
3: 25 Z Loc [ temp1 ]

;this actually is calculating COS(angle-loc 6)
24: Z=SIN(X) (P48)
1: 25 X Loc [ temp1 ]
2: 8 Z Loc [ coswindr ]

;solar measurements

;S/N 5829
25: Volt (SE) (P1)
1: 1 Reps
2: 12 15 mV Fast Range
3: 11 SE Channel
4: 9 Loc [ solar1 ]
5: -81.486 Mult
6: 0.0 Offset

;S/N 5822
26: Volt (SE) (P1)
1: 1 Reps
2: 12 15 mV Fast Range
3: 12 SE Channel
4: 10 Loc [ solar2 ]
5: -94.09 Mult
6: 0.0 Offset

```

```

;S/N 5801
27: Volt (SE) (P1)
1: 1 Reps
2: 12 15 mV Fast Range
3: 13 SE Channel
4: 11 Loc [ solar3 ]
5: -111.09 Mult
6: 0.0 Offset

```

```

;S/N 4626/5626
28: Volt (SE) (P1)
1: 1 Reps
2: 12 15 mV Fast Range
3: 14 SE Channel
4: 12 Loc [ solar4 ]
5: -95.355 Mult
6: 0.0 Offset

```

```

;S/N 4951
29: Volt (SE) (P1)
1: 1 Reps
2: 12 15 mV Fast Range
3: 15 SE Channel
4: 13 Loc [ solar5 ]
5: -73.179 Mult
6: 0.0 Offset

```

=====OUTPUT ROUTINES=====

```

; is the number of tenths of a second into this
minute<5?
;(this allows some "slop" in timing, rather than
using
something like:

```

```

;22: If time is (P92)
; 1: 0 Minutes into a
; 2: 1 Minute Interval
; 3: 10 Set Output Flag High

```

```

30: If (X<=>F) (P89)
1: 15 X Loc [ time ]
2: 4 <
3: 5 F
4: 10 Set Output Flag High

```

```

31: Set Active Storage Area (P80)
1: 1 Final Storage
2: 101 Array ID

```

32: Real Time (P77)
1: 1221 Year,Day,Hour/Minute,Seconds
(midnight = 2400)

33: Average (P71)
1: 1 Reps
2: 2 Loc [T_in]

34: Average (P71)
1: 1 Reps
2: 16 Loc [T_mass]

35: Average (P71)
1: 1 Reps
2: 17 Loc [T_inletE]

36: Average (P71)
1: 1 Reps
2: 18 Loc [T_inletW]

37: Average (P71)
1: 1 Reps
2: 19 Loc [T_inletAt]

38: Average (P71)
1: 1 Reps
2: 20 Loc [T_att1]

39: Average (P71)
1: 1 Reps
2: 21 Loc [T_att2]

40: Average (P71)
1: 1 Reps
2: 22 Loc [T_sunsp]

41: Average (P71)
1: 1 Reps
2: 23 Loc [T_bsmt]

42: Average (P71)
1: 1 Reps
2: 24 Loc [T_out]

;SOLAR MEASUREMENTS

43: Average (P71)
1: 1 Reps
2: 9 Loc [solar1]
44: Average (P71)

1: 1 Reps
2: 10 Loc [solar2]
45: Average (P71)

1: 1 Reps
2: 11 Loc [solar3]
46: Average (P71)

1: 1 Reps
2: 12 Loc [solar4]
47: Average (P71)

1: 1 Reps
2: 13 Loc [solar5]

;wind direction

48: Average (P71)
1: 1 Reps
2: 7 Loc [sinwindr]

49: Average (P71)
1: 1 Reps
2: 8 Loc [coswinddr]

;output wind speed

50: Totalize (P72)
1: 1 Reps
2: 5 Loc [windspd]

;output control voltage

51: Average (P71)
1: 1 Reps
2: 3 Loc [CtrlVolt]

*Table 2 Program
02: 0.0000 Execution Interval (seconds)

*Table 3 Subroutines

End Program

-Input Locations-

1 Tref 1 1 1
2 T_in 1 1 2
3 CtrlVolt 1 2 1
4 WindPot 1 1 1
5 windspd 1 1 1
6 winddir 1 3 1
7 sinwindr 1 1 1
8 coswinddr 1 1 1
9 solar1 1 1 1
10 solar2 1 1 1
11 solar3 1 1 1

12 solar4 1 1 1
13 solar5 1 1 1
14 _____ 1 0 0
15 time 1 1 1
16 T_mass 1 1 1
17 T_inletE 1 1 1
18 T_inletW 1 1 1
19 T_inletAt 1 1 1
20 T_att1 1 1 1
21 T_att2 1 1 1
22 T_sunsp 1 1 1

23 T_bsmt 1 1 1
24 T_out 1 1 1
25 temp1 1 1 1
26 _____ 0 0 0
27 _____ 0 0 0
28 _____ 0 0 0
-Program Security-
0
0000
0000

A2.2.2 BRDMR.DLD

```

};21X
;BRDMR.DLD
;$
;:Tref :T_in :CtrlVolt:WindPot :windspd
;:winddir :sinwindr :coswindr:solar1 :solar2
;:solar3 :solar4 :solar5 :_____time
;:T_mass :T_inletE:T_inletW:T_inletAt:T_att1
;:T_att2 :T_sunsp :T_bsmt :T_out :temp1
;:_____:_____:_____
;$

```

MODE 1	10:P13	8:0.0	6:22	6:0.0
SCAN RATE	1:1		7:1.0	
1	2:1	14:P13	8:0.0	22:P48
	3:1	1:1		1:6
1:P18	4:1	2:1	18:P13	2:7
1:00	5:1	3:5	1:1	
2:600	6:2	4:1	2:1	23:P34
3:15	7:1.0	5:1	3:9	1:6
	8:0.0	6:19	4:1	2:90
2:P17		7:1.0	5:1	3:25
1:1	11:P13	8:0.0	6:23	
	1:1		7:1.0	24:P48
3:P91	2:1	15:P13	8:0.0	1:25
1:21	3:2	1:1		2:8
2:30	4:1	2:1	19:P13	
	5:1	3:6	1:1	25:P1
4:P30	6:16	4:1	2:1	1:1
1:0.0	7:1.0	5:1	3:10	2:12
2:3	8:0.0	6:20	4:1	3:11
		7:1.0	5:1	4:9
5:P30	12:P13	8:0.0	6:24	5:-81.486
1:2500	1:1		7:1.0	6:0.0
2:4	2:1	16:P13	8:0.0	
	3:3	1:1		26:P1
6:P91	4:1	2:1	20:P3	1:1
1:21	5:1	3:7	1:1	2:12
2:11	6:17	4:1	2:1	3:12
	7:1.0	5:1	3:2	4:10
7:P95	8:0.0	6:21	4:5	5:-94.09
		7:1.0	5:.01676	6:0.0
8:P21	13:P13	8:0.0	6:0.0	
1:1	1:1			27:P1
2:3	2:1	17:P13		1:1
	3:4	1:1	21:P1	2:12
9:P21	4:1	2:1	1:1	3:13
1:2	5:1	3:8	2:15	4:11
2:4	6:18	4:1	3:16	5:-111.09
	7:1.0	5:1	4:6	6:0.0
			5:.144	

28:P1	2:101	38:P71	1:1	2:5
1:1	32:P77	1:1	2:10	51:P71
2:12	1:1221	2:20	45:P71	1:1
3:14	33:P71	39:P71	1:1	2:3
4:12	1:1	1:1	2:11	MODE 2
5:-95.355	2:2	2:21	46:P71	SCAN RATE
6:0.0	34:P71	40:P71	1:1	0.0000
29:P1	1:1	1:1	2:12	MODE 3
1:1	2:16	2:22	47:P71	MODE 10
2:12	35:P71	41:P71	1:1	1:28
3:15	1:1	1:1	2:13	2:64
4:13	2:17	2:23	48:P71	MODE 12
5:-73.179	36:P71	42:P71	1:1	1:0
6:0.0	1:1	1:1	2:7	2:0000
30:P89	2:18	2:24	49:P71	
1:15	37:P71	43:P71	1:1	
2:4	1:1	1:1	2:8	
3:5	2:19	2:9	50:P72	
4:10	31:P80	44:P71	1:1	
1:1	1:1			

Appendix 3 Data Retrieval from Loggers

The first two sections of this appendix contain QuickBASIC code used to retrieve data from the Campbell data logger and write it to a file. The code in the first section is used the first time the Campbell is polled; the code in the second section is used for subsequent pollings. The third section contains code used to run `kdump.exe`⁸⁴ for retrieving data from the C180E logger. The fourth section contains the parameter file required by `kdump.exe`. The fifth and sixth sections contain analogous code and parameter files for the K20 logger.

A3.1 FIRSTDL.D.BAS

```
DECLARE SUB sendCommand (comport%, filenumber%, cmd$, reps%, retVal$)
DECLARE SUB readA (A$, AR%, AF%, AL%)
DECLARE SUB FullDump (comport%, filenumber%, AF%, AR%, DL%)
CLEAR
CLS
CLOSE

COM(1) ON
ON COM(1) GOSUB ReadData      'as soon as something appears in read
ready% = 0                   'buffer, then read it in...

=====
comport% = FREEFILE
OPEN "com1:9600,N,8,1,ASC,DS0,CS0,CD0,RS" FOR RANDOM AS #comport%

filenumber% = FREEFILE
OPEN "CAMPDT.DAT" FOR OUTPUT AS filenumber%

statusfile% = FREEFILE
OPEN "STATUS.DAT" FOR OUTPUT AS statusfile%
=====
DO
    PRINT #comport%, CHR$(13);
    PLAY "mf T255 p14"
LOOP UNTIL ready% = 1
COM(1) OFF
=====

CALL sendCommand(comport%, filenumber%, "A", 0, retVal$)
'PRINT retVal$

CALL readA(retVal$, AR%, AF%, AL%)

CALL FullDump(comport%, filenumber%, AF%, AR%, DL%)
'PRINT DL%

CALL sendCommand(comport%, filenumber%, TIMES$ + "C" + CHR$(13), 0, retVal$)
```

⁸⁴ Fishbaugher and Associates, Oakland, OR.

```
WRITE #statusfile%, 1
WRITE #statusfile%, DL%
WRITE #statusfile%, 0
WRITE #statusfile%, 0
```

```
CLOSE #comport%
CLOSE #filenumber%
CLOSE #statusfile%
```

```
END
```

```
ReadData:
```

```
INPUT #comport%, inputstring$
PRINT inputstring$
ready% = 1
RETURN
```

```
'+++++
```

```
SUB FullDump (comport%, filenumber%, AF%, AR%, DL%)
```

```
'AF is the num of full entries in storage
```

```
'AR is the num position of the pointer in the storage area
```

```
avgRecLength! = 6!   'NOTE THESE 2 MUST BE ENTERED BY USER!!
storageSize% = 14284 'Get this # by pressing *AAA on the Campbell keypad
                    'avgRecLength is the number of entries the Campbell
                    'makes in the final storage area, including labels
                    'and times
```

```
IF AF% = storageSize% THEN
```

```
PRINT "Storage is full"
```

```
numD% = INT(storageSize% / avgRecLength!) - 6
```

```
IF (numD% MOD 2) = 0 THEN
```

```
    posOfOldestData% = AR% - (numD% + 4) * avgRecLength!
```

```
ELSE
```

```
    posOfOldestData% = AR% - (numD% + 3) * avgRecLength!
```

```
    numD% = numD% + 1
```

```
END IF
```

```
IF posOfOldestData% < 0 THEN
```

```
    posOfOldestData% = posOfOldestData% + storageSize%
```

```
END IF
```

```
memLoc$ = STR$(posOfOldestData%)
```

```
CALL sendCommand(comport%, filenumber%, memLoc$ + "G", 0, retVal$)
```

```

ELSE
    PRINT "Storage not full yet"
    numD% = AF% / avgRecLength!
    CALL sendCommand(comport%, filenumber%, "IG", 0, retVal$)
END IF

```

```

cmd$ = STR$(numD%) + "D"
PRINT "Reading all data"

```

```

CALL sendCommand(comport%, filenumber%, cmd$, numD%, retVal$)
PRINT "Data have been read"

```

```

'PRINT retVal$
lPosition% = INSTR(retVal$, "L+")

```

```

DL% = VAL(MID$(retVal$, lPosition% + 2, 6))

```

```

END SUB

```

```

'+++++

```

```

SUB readA (A$, AR%, AF%, AL%)

```

```

lenA% = LEN(A$)
AR% = VAL(MID$(A$, 3, 6))
AF% = VAL(MID$(A$, 12, 6))
position% = INSTR(A$, "L+")
AL% = VAL(MID$(A$, position% + 2, 6))

```

```

END SUB

```

```

'+++++

```

```

SUB sendCommand (comport%, filenumber%, cmd$, reps%, retVal$)

```

```

lenCmd% = LEN(cmd$)
location% = INSTR(cmd$, " ")
IF location% = 1 THEN
    cmd$ = RIGHT$(cmd$, lenCmd% - 1)
END IF

```

```

PRINT #comport%, cmd$

```

```

ctr% = 0
done% = 0

```

```

DO
    ctr% = ctr% + 1

```

```

INPUT #comport%, A$
correct% = INSTR(A$, MID$(cmd$, 1, 1))

IF correct% > 0 THEN
    done% = 1
ELSE
    done% = 0
END IF

IF ctr% >= 100 THEN
    PRINT "stuck"
    INPUT temp$
END IF

LOOP UNTIL done% = 1

IF reps% > 0 THEN
FOR i% = 1 TO reps%
    INPUT #comport%, A$
    ALen% = LEN(A$)
    IF ALen% < 4 THEN
        PRINT "stuck"
    END IF

    RITE #filenumber%, A$
    IF i% MOD 200 = 0 THEN
        PRINT i% / reps%
    END IF
NEXT
END IF

INPUT #comport%, A$
ALen% = LEN(A$)
retVal$ = A$

END SUB

```

A3.2 UPDATE.BAS

```
DECLARE SUB sendCommand (comport%, filename%, cmd$, reps%, retVal$)
DECLARE SUB readA (A$, AR%, AF%, AL%)
DECLARE SUB FullDump (comport%, filename%, AF%, AR%, DL%)
CLEAR
CLS
CLOSE

COM(1) ON
ON COM(1) GOSUB ReadData      'as soon as something appears in read
ready% = 0                    'buffer, then read it in...

=====
comport% = FREEFILE
OPEN "com1:9600,N,8,1,ASC,DS0,CS0,CD0,RS" FOR RANDOM AS #comport%

filename% = FREEFILE
OPEN "CAMPDT.DAT" FOR OUTPUT AS filename%

statusfile% = FREEFILE
OPEN "STATUS.DAT" FOR OUTPUT AS statusfile%
=====
DO
    PRINT #comport%, CHR$(13);
    PLAY "mf T255 p14"
LOOP UNTIL ready% = 1
COM(1) OFF
=====

CALL sendCommand(comport%, filename%, "A", 0, retVal$)
'PRINT retVal$

CALL readA(retVal$, AR%, AF%, AL%)

CALL FullDump(comport%, filename%, AF%, AR%, DL%)
'PRINT DL%

CALL sendCommand(comport%, filename%, TIMES$ + "C" + CHR$(13), 0, retVal$

WRITE #statusfile%, 1
WRITE #statusfile%, DL%
WRITE #statusfile%, 0
WRITE #statusfile%, 0

CLOSE #comport%
CLOSE #filename%
CLOSE #statusfile%
```

END

ReadData:

```
INPUT #comport%, inputstring$
PRINT inputstring$
ready% = 1
RETURN
```

'+++++

SUB FullDump (comport%, filenumber%, AF%, AR%, DL%)

'AF is the num of full entries in storage

'AR is the num position of the pointer in the storage area

avgRecLength! = 6! 'NOTE THESE 2 MUST BE ENTERED BY USER!!

storageSize% = 14284 'Get this # by pressing *AAA on the Campbell keypad

'avgRecLength is the number of entries the Campbell

'makes in the final storage area, including labels

'and times

IF AF% = storageSize% THEN

PRINT "Storage is full"

numD% = INT(storageSize% / avgRecLength!) - 6

IF (numD% MOD 2) = 0 THEN

posOfOldestData% = AR% - (numD% + 4) * avgRecLength!

ELSE

posOfOldestData% = AR% - (numD% + 3) * avgRecLength!

numD% = numD% + 1

END IF

IF posOfOldestData% < 0 THEN

posOfOldestData% = posOfOldestData% + storageSize%

END IF

memLoc\$ = STR\$(posOfOldestData%)

CALL sendCommand(comport%, filenumber%, memLoc\$ + "G", 0, retVal\$)

ELSE

PRINT "Storage not full yet"

numD% = AF% / avgRecLength!

CALL sendCommand(comport%, filenumber%, "1G", 0, retVal\$)

END IF

cmd\$ = STR\$(numD%) + "D"

PRINT "Reading all data"

```
CALL sendCommand(comport%, filenumber%, cmd$, numD%, retVal$)
PRINT "Data have been read"
```

```
'PRINT retVal$
lPosition% = INSTR(retVal$, "L+")
```

```
DL% = VAL(MID$(retVal$, lPosition% + 2, 6))
```

```
END SUB
```

```
'+++++
```

```
SUB readA (A$, AR%, AF%, AL%)
```

```
lenA% = LEN(A$)
AR% = VAL(MID$(A$, 3, 6))
AF% = VAL(MID$(A$, 12, 6))
position% = INSTR(A$, "L+")
AL% = VAL(MID$(A$, position% + 2, 6))
```

```
END SUB
```

```
'+++++
```

```
SUB sendCommand (comport%, filenumber%, cmd$, reps%, retVal$)
```

```
lenCmd% = LEN(cmd$)
location% = INSTR(cmd$, " ")
IF location% = 1 THEN
    cmd$ = RIGHT$(cmd$, lenCmd% - 1)
END IF
```

```
PRINT #comport%, cmd$
```

```
ctr% = 0
done% = 0
```

```
DO
```

```
    ctr% = ctr% + 1
```

```
    INPUT #comport%, A$
    correct% = INSTR(A$, MID$(cmd$, 1, 1))
```

```
    IF correct% > 0 THEN
```

```
        done% = 1
```

```
    ELSE
```

```
        done% = 0
```

```
    END IF
```

```

IF ctr% >= 100 THEN
    PRINT "stuck"
    INPUT temp$
END IF

LOOP UNTIL done% = 1

IF reps% > 0 THEN
FOR i% = 1 TO reps%
    INPUT #comport%, A$
    ALen% = LEN(A$)
    IF ALen% < 4 THEN
        PRINT "stuck"
    END IF

    WRITE #filenumber%, A$
    IF i% MOD 200 = 0 THEN
        PRINT i% / reps%
    END IF
NEXT
END IF

INPUT #comport%, A$
ALen% = LEN(A$)
retVal$ = A$

END SUB

```


A3.3 C180.BAS

```
CLEAR
CLS

inp$ = ""

DO
  DO
    'wait to line up with minute
    keyhit$ = INKEY$
    IF keyhit$ <> "" THEN
      inp$ = UCASE$(keyhit$)
    END IF

    LOOP UNTIL VAL(MID$(TIME$, 7, 2)) = 5 OR inp$ = "Q"

    start$ = TIME$

    filenumber% = FREEFILE
    ' KILL "stopC.dat"
    ' PRINT "deleted stop file"
    OPEN "startC.dat" FOR OUTPUT AS filenumber%
    PRINT "created startC"
    CLOSE #filenumber%

    IF inp$ = "V" THEN
      SHELL "kdump sc1630.dat -v"
    ELSE
      SHELL "kdump sc1630.dat"
    END IF

    filenumber% = FREEFILE
    OPEN "stopC.dat" FOR OUTPUT AS filenumber%
    PRINT "created stopC"
    CLOSE #filenumber%

    KILL "startC.dat"
    PRINT "deleted start file"

    PRINT start$ + " --> " + TIME$

  LOOP UNTIL inp$ = "Q"

  PRINT "Press any key to quit"
  DO
    keyhit$ = INKEY$
  LOOP UNTIL keyhit$ <> ""

END
```

A3.4 SC1630.DAT

The parameter file referenced below (p1630.001) is the same one included in A1.1.

```
COMPORT: 3
  BAUD: 2400
  AUTO: 0
MSETUP:
  PHONE: DIRECT
TIMEOUT: 2.0
  CALLS: 5
  SYNC: 2.0
  START: 640
  STOP: -1
  LAST: 639
LSTAMP: 10/07/2003 01:43:00
  DTYPE: 2
  TSRNDX: 1
SKIPCRC: 0
PARFILE: \ktools\param\p1630.001
  KEY: 222F6246
  D1FILE: T1630.TXT
  D2FILE: DUMPC180.TXT
LOGFILE: LOGC180.TXT
ERRFILE: ERRC180.TXT
```

A3.5 K20.BAS

```
CLEAR
CLS
```

```
inp$ = ""
```

```
DO
```

```
  DO
```

```
    'wait to line up with minute
    keyhit$ = INKEY$
    IF keyhit$ <> "" THEN
      inp$ = UCASE$(keyhit$)
    END IF
```

```
  LOOP UNTIL VAL(MID$(TIME$, 7, 2)) = 5 OR inp$ = "Q"
```

```
  start$ = TIME$
```

```
  filenumber% = FREEFILE
```

```
'  KILL "stopK.dat"
'  PRINT "deleted stop file"
  OPEN "startK.dat" FOR OUTPUT AS filenumber%
  PRINT "created startK"
  CLOSE #filenumber%
```

```

IF inp$ = "V" THEN
    SHELL "kdumpk s10046.dat -v"
ELSE
    SHELL "kdumpk s10046.dat"
END IF

filenumber% = FREEFILE
OPEN "stopK.dat" FOR OUTPUT AS filenumber%
PRINT "created stopK"
CLOSE #filenumber%

KILL "startK.dat"
PRINT "deleted start file"

PRINT start$ + " --> " + TIMES$

LOOP UNTIL inp$ = "Q"

PRINT "Press any key to quit"
DO
    keyhit$ = INKEY$
LOOP UNTIL keyhit$ <> ""
END

```

A3.6 S10046.DAT

The parameter file referenced below (p10046.001) is the same one included in A1.2.

```

COMPORT: 4
    BAUD: 2400
    AUTO: 0
MSETUP:
    PHONE: DIRECT
TIMEOUT: 2.0
    CALLS: 5
    SYNC: 2.0
    START: 1281
    STOP: -1
    LAST: 1280
LSTAMP: 10/07/2003 05:36:00
    DTYPE: 2
    TSRNDX: 1
SKIPCRC: 0
PARFILE: \ktools\param\p10046.001
    KEY: 558009A2
    D1FILE: T10046.TXT
    D2FILE: DUMPK20.TXT
LOGFILE: LOGK20.TXT
ERRFILE: ERRK20.TXT

```


Appendix 4 Description of Manual Monitoring and Control Spreadsheet

The following spreadsheet was used to enter information into the computer regarding aperture openings, desired fan settings and day types. A detailed explanation is provided below. The table below is reproduced from Table 4-6 for convenience.

minute starting at:	Egrill:Wgrill:Conf/AssDoor:Bsmtdoor:Att1Win:Att2Win WinS,WinW,WinAss,AtticSlider,FanSlider	Fan	Setpoint	Override (fan/setpt)	daytype	confirm
0:00	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1
0:01	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1
0:02	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1
0:03	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1
0:04	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1
0:05	"000001.000,000,000,000,000"	0.0	25.0	1.0	10	1

Table 10-1 Excerpt of one day's manual monitoring and control spreadsheet.

The aperture data were entered in code form for compactness. In the second column of the table, the digits before the decimal point were dedicated to apertures with discrete opening amounts: 0/1 (2, in the case of the Conference an Assembly doors, which were lumped together). The corresponding aperture numbers before the decimal are, in order, #1, #2, #3, #12, #9⁸⁵, #10. Apertures that could be fractionally open were assigned three digits, separated by commas. The sequence of digits was divided by 100 to determine the opening fraction. The aperture numbers for the entries after the decimal point are: #4, #5, #11, #6, #7.

The fan setting in mV (0-5000) was entered in the third column. A column for the desired setpoint was provided, but was not used during the experiments. The fifth column indicated whether the fan and setpoint settings provided on the spreadsheet should override any automatically generated settings within MATLAB. This was always true.

The "daytype" column was used to indicate how the space was being used. The first digit was used to specify if the day was a regular weekday (1), a weekday when camp was in session (2), or a holiday or weekend day (3). The second digit was used to denote hours when a presentation or special event was being held in the assembly room (0 for no event, 1 for an event).

The final column was used to alert MATLAB when the data in a given line had been verified and was ready to be loaded into the workspace. This feature permitted notes taken by staff members regarding when various apertures were opened to be entered into the computer at a later time.

This setup was clearly rather cumbersome, and would be greatly improved through the use of automatic sensors on all apertures.

⁸⁵ At the time the spreadsheets were developed, there was a possibility of using the windows in Attic 1 as controlled apertures. That was not done, so this entry position was reassigned to the Sunspace door.