

Title	Modelling and predictive control techniques for building heating systems
Authors	O'Dwyer, Edward
Publication date	2016
Original Citation	O'Dwyer, E. 2016. Modelling and predictive control techniques for building heating systems. PhD Thesis, University College Cork.
Type of publication	Doctoral thesis
Rights	© 2016, Edward O'Dwyer. - http://creativecommons.org/licenses/by-nc-nd/3.0/
Download date	2024-04-20 08:45:58
Item downloaded from	https://hdl.handle.net/10468/3472

Modelling and Predictive Control Techniques for Building Heating Systems

Edward O'Dwyer



NATIONAL UNIVERSITY OF IRELAND, CORK

SCHOOL OF ENGINEERING

**Thesis submitted for the degree of
Doctor of Philosophy**

1st April 2016

Head of School: Professor Nabeel A. Riza

Supervisor: Dr Gordon Lightbody

Research supported by The Irish Research Council and United Technologies Research
Centre Ireland

Contents

List of Figures	v
List of Tables	xi
Acknowledgements	xiii
Abstract	xiv
1 Introduction	1
1.1 Motivation	1
1.2 Research Objectives	4
1.3 Thesis Outline	5
2 Technologies used in Building Energy Systems	9
2.1 Introduction	9
2.2 Control of Building Energy Systems	10
2.2.1 Passive and Active Energy Efficiency	10
2.2.2 Traditional Approaches for Thermal Control of Buildings	11
2.2.3 Intelligent Model-Free Approaches for the Control of Building Energy Systems	14
2.2.3.1 Defining a Desired Thermal Performance	14
2.2.3.2 Obtaining a Desired Performance	15
2.3 Optimal Model-based Control Formulations	16
2.3.1 Solving Constrained Optimisation Problems	17
2.3.2 Model Predictive Control	21
2.3.2.1 General Formulation	21
2.3.2.2 Applying MPC to Building Energy Systems	23
2.4 Conclusion	24
3 Representing a Building's Thermodynamics in a Validated Simulation Platform	31
3.1 Introduction	31
3.2 Modelling the Thermal Properties of Buildings	32
3.3 Representative RC-Network for Analysis of Modelling Methodologies	34
3.4 Identification of Model Parameters from Measured Data	37
3.4.1 Outlining the Identification Problem - Complexity and Identifiability Considerations	37
3.4.1.1 Excitation of the System	37
3.4.1.2 Complexity of the Parameter Selection Process	38
3.4.2 Metaheuristic Search Algorithms for Complex Optimisation Problems	39
3.5 Chemotaxis Search for Identifying Model Parameters	42
3.5.1 The Chemotaxis Algorithm	42
3.5.2 Evaluation of solutions	44
3.6 Platform Parameter Identification in the Presence of Unmeasured Disturbances	47
3.6.1 Generation of disturbance profiles	49
3.6.2 Chemotaxis in the Presence of Unmeasured Disturbances	51

3.7	Methods of Disturbance Estimation for Improved System-Identification	52
3.7.1	Kalman Filtered Disturbance Estimation	53
3.7.2	Spatial Filtering using Principal Component Analysis	55
3.7.3	Incorporating Disturbance Estimation in the Identification Process	58
3.7.4	Chemotaxis Algorithm with Disturbance Estimation	60
3.8	Identification Algorithm Based on Particle Swarm Optimisation and Disturbance Estimation	61
3.8.1	PSO - Implementation with test network	62
3.8.2	Quantum-behaved Particle Swarm Optimisation	63
3.8.2.1	Q-PSO - Implementation with test network	64
3.8.2.2	Q-PSO Decentralised Search	66
3.9	Case Study - Modelling the Nimbus Centre from Process Data	67
3.9.1	Processing Raw Data for Use in Identification Strategy	70
3.9.1.1	Heating System Sensor Data	70
3.9.1.2	Zone and External Temperature Sensor Data	72
3.9.2	Physics-based RC-Network Representing Initial Estimate of Nimbus Centre	72
3.9.3	Training of Physics-based Network Using Measured Data	72
3.10	Modelling of the Nimbus Centre Heating System	73
3.10.1	Boiler & Header Thermodynamic Models	74
3.10.2	Modelling the PI-Controlled Mixing Valves	75
3.10.3	Modelling the Zone-level Thermostatic Valves	76
3.10.4	System Pumps and Flow Rate Considerations	77
3.11	Conclusion	78
4	Development of Low-Order Building Zone Models Suitable for Predictive Control Applications	84
4.1	Introduction	84
4.2	Model Suitability for Predictive Control Formulations	86
4.2.1	Decomposing the Full-Scale Building Model	86
4.2.2	Evaluating Performance of the Zone Models	87
4.3	Physics-based Techniques for Identifying Reduced-Order Models	87
4.3.1	Model Reduction - Hankel Singular Value Decomposition	88
4.4	Deriving Zone Models from Measured Data	92
4.4.1	System-Identification Methodology	93
4.4.2	Identifiability - Limitations Due to Excitation	95
4.4.3	Models Derived Using Undisturbed Data from the Test Network	96
4.5	Including Unmeasured Disturbances in the Identification Data	100
4.5.1	Identification without Disturbance Estimation	100
4.5.2	Inclusion of the Spatio-Temporally Filtered Disturbance Estimates in the Identification Process	102
4.5.3	Models Derived Using Disturbed Test Network Data with Disturbance Estimation	104
4.6	Deriving Zone Models Using Real Data - The Nimbus Centre	107
4.6.1	Nimbus Centre Inputs, Outputs and Disturbances	107

4.6.2	Zone Models Derived from Measured Nimbus Centre Data . .	109
4.7	Conclusion	112
5	Scalable Model Predictive Control Strategies for Building Heating Systems	118
5.1	Introduction	118
5.2	Defining the Cost - Balancing Energy Consumption with Occupant Comfort	119
5.2.1	Performance Criteria to be Optimised	119
5.2.2	Standard Formulations - Quadratic Cost Functions	120
5.2.3	Implementation of Standard Quadratic Objective in the Simulation Platform	121
5.3	Scale and Scalability	129
5.3.1	Handling the Growing Complexity in Larger Buildings	130
5.3.1.1	Distributing the Optimisation Problem	130
5.3.1.2	A Hierarchical Approach - Reducing the Problem . .	131
5.3.2	Tuning Considerations - Setting a Desired Level of Performance	133
5.4	A Prioritised-Objectives MPC Approach for Building Heating Systems	137
5.4.1	Multiple Objectives and Lexicographic MPC Formulation . .	137
5.4.2	Development of the Cost Functions	140
5.4.3	Supply and Demand - Nonlinear Heating System Components	142
5.4.3.1	Heating System Considerations and Nonlinearities .	143
5.4.3.2	Separating the Supply and Demand Problems	146
5.4.4	Formulation of the Prioritised-Objective MPC Strategy	147
5.4.4.1	Demand-Side Constraints	147
5.4.4.2	Supply-Side Constraints	150
5.4.5	Performance of Prioritised Objective Strategy in Simulation Platform	150
5.4.6	Load Flexibility - Choosing the Acceptable Level of Set-point Deviation	152
5.5	Conclusion	154
6	Inaccurate Predictions and System Faults in the MPC Formulation	160
6.1	Introduction	160
6.2	Analysing the Impact of Inaccurate Predictions	161
6.2.1	Effect of Zone Model Inaccuracies on Optimisation	162
6.2.2	Forecasting Disturbances in the Formulation	165
6.2.2.1	Inaccuracies in Measured Disturbance Forecast . .	166
6.2.2.2	Prediction of the Unmeasured Disturbances for Improved Control Performance	168
6.3	Faults and System Changes - Detection and Identification	172
6.3.1	Fault Detection and Identification using Measured Data	173
6.3.2	Detection and Identification of Sensor Faults	173
6.3.3	Detection and Identification of Actuator Faults	180
6.4	Reconfiguration of the MPC Formulation when a Fault Occurs	185
6.4.1	Updating the Model and the Optimisation Constraints	186
6.4.2	Reformulating the Objective Function in Faulted Conditions .	187

6.5	Conclusion	190
7	Conclusions and Future Work	196
7.1	Conclusions	196
7.2	Future Work	200
A	Boiler Efficiency Curves	A1
B	Quantum-behaved Particle Swarm Optimisation	B1

List of Figures

2.1	Active energy efficiency, using wireless sensors to determine electrical and thermal load requirements	12
2.2	Weather-compensated mixing valve in hydronic heating system	13
2.3	Convex (strict), convex and non-convex surface plots	17
2.4	Interior-point algorithm pathway as barrier weight is reduced to 0 . .	20
2.5	Receding Horizon for 3 samples, using a horizon length of 5 samples .	22
3.1	Equivalent 3R2C RC-network model of the thermal connection between a zone and the external environment	33
3.2	Equivalent RC Network model of two zones with one internal wall, six external walls and six windows	34
3.3	RC-network used for testing identification strategies	35
3.4	Convective heat supply profile applied to Zone 1 of test network used for training process	36
3.5	External temperature profile applied to test network used for training process	37
3.6	Chemotaxis search on a nonlinear surface	40
3.7	Particle Swarm Optimisation with 10 agents	42
3.8	Fitness-function reduction as algorithm progresses for different levels of search aggression	44
3.9	Validation data-set output comparison (temperature of Zone 8) of identified networks with target network	46
3.10	Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 1 using target system, initial system and identified systems (for different values of γ_{ch}) .	47
3.11	Frequency response comparison of transfer function between external temperature and output temperature for Zone 1 using target system, initial system and identified systems (for different values of γ_{ch}) . . .	47
3.12	Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 25 using target system, initial system and identified systems (for different values of γ_{ch}) .	48
3.13	Frequency response comparison of transfer function between external temperature and output temperature for Zone 25 using target system, initial system and identified systems (for different values of γ_{ch}) . . .	48
3.14	Zone 2 internal gain profile (occupancy and equipment gains) for 5 day period	49
3.15	Relevant angles for determining intensity of solar heat gain	50
3.16	Direct solar radiation profile for each orientation over 5 day period . .	51
3.17	Fitness-function reduction as algorithm progresses - unmeasured disturbances included	52
3.18	Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 1 using target system, initial system and identified systems (with disturbances included)	53

3.19	Frequency response comparison of transfer function between external temperature and output temperature for Zone 1 using target system, initial system and chemotaxis-identified systems (with disturbances included)	53
3.20	Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 25 using target system, initial system and chemotaxis-identified systems (with disturbances included)	54
3.21	Frequency response comparison of transfer function between external temperature and output temperature for Zone 25 using target system, initial system and chemotaxis-identified systems (with disturbances included)	54
3.22	Set of 10 correlated variables comprising common and individual constituent parts	56
3.23	Set of 10 variables after filtering process (PCA)	56
3.24	Disturbance Estimation - Zones are grouped by geographical similarity. Zone models and disturbance estimates are determined iteratively	59
3.25	Spectrum of variance described by first 8 principal components for central zone disturbances	60
3.26	Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 1 using target system, initial system and decentralised Q-PSO-identified systems (with disturbances and disturbance estimation included)	67
3.27	Frequency response comparison of transfer function between external temperature and output temperature for Zone 1 using target system, initial system and decentralised Q-PSO-identified systems (with disturbances and disturbance estimation included)	67
3.28	Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 25 using target system, initial system and decentralised Q-PSO-identified systems (with disturbances and disturbance estimation included)	68
3.29	Frequency response comparison of transfer function between external temperature and output temperature for Zone 25 using target system, initial system and decentralised Q-PSO-identified systems (with disturbances and disturbance estimation included)	68
3.30	Validation data-set output (temperature of Zone 17) using identified network and initial estimated network compared to the target network	69
3.31	The Nimbus Centre	69
3.32	Schematic of ground floor section of heating system used in the Nimbus Centre	70
3.33	Supply and demand boiler water temperature measurements taken for a period of three days	70
3.34	Comparison of real data (Nimbus Centre) with simulation platform: Q-PSO adjusted vs. original RC network comparison with real data - Ground Floor Meeting Room	73

3.35	Comparison of real data (Nimbus Centre) with simulation platform: Q-PSO adjusted vs. original RC network comparison with real data - Ground Floor Corridor	74
3.36	Block diagram of header dynamics with heat added by the boiler and heat removed by the building	75
3.37	Block diagram of mixing valve control loop	76
3.38	Block diagram of zone valve associated with i^{th} zone	77
3.39	Change in heat flow to zone as mass flow rate through radiators is increased	78
4.1	15 largest Hankel singular values for transfer function between the heat input and zone temperature output of Zone 1 of test network (plotted on log scale)	90
4.2	Frequency responses of full-order system and reduced-order approximations for transfer function between heat input and internal temperature for Zone 1 of the test network	90
4.3	Frequency responses of full-order system and reduced-order approximations for transfer function between external and internal temperatures for Zone 1 of the test network	91
4.4	Frequency responses of full-order system and reduced-order approximations for transfer function between heat input and internal temperature for Zone 1 of the test network - preserving DC-gain	92
4.5	Frequency responses of full-order system and reduced-order approximations for transfer function between external and internal temperatures for Zone 1 of the test network - preserving DC gain	92
4.6	Frequency responses of full-order system and data-driven low-order transfer functions between heat input and internal temperature for Zone 1 of the test network	97
4.7	Frequency responses of full-order system and data-driven low-order transfer functions between external and internal temperatures for Zone 1 of the test network	98
4.8	Frequency responses of full-order system and data-driven low-order transfer functions between heat input and internal temperature for Zone 1 of the test network using PRBS input for training	98
4.9	Frequency responses of full-order system and data-driven low-order transfer functions between external and internal temperature for Zone 1 of the test network using PRBS input for training	99
4.10	p -step (6 hours) ahead predicted output for $1^{st} - 4^{th}$ order models of Zone 1	100
4.11	Frequency responses of full-order system and data-driven low-order transfer functions between heat input and internal temperature for Zone 1 of the test network derived with disturbances present	101
4.12	Frequency responses of full-order system and data-driven low-order transfer functions between external and internal temperatures for Zone 1 of the test network derived with disturbances present	101

4.13	p -step (6 hours) ahead predicted output for $1^{st} - 4^{th}$ order models of Zone 1 (zone models derived in presence of disturbance)	102
4.14	Iterative process of parameter identification and disturbance estimation	104
4.15	Frequency responses of full-order system and data-driven low-order transfer functions between heat input and internal temperature for Zone 1 of the test network derived with and without disturbance estimation	105
4.16	Frequency responses of full-order system and data-driven low-order transfer functions between external and internal temperatures for Zone 1 of the test network derived with and without disturbance estimation	105
4.17	72-step (6 hours) ahead predicted output for 2^{nd} order models of Zone 1 derived with and without disturbance estimation	106
4.18	72-step (6 hours) ahead predicted output for 2^{nd} order models of Zone 4 derived with and without PCA	107
4.19	Bode plot comparison of simulation platform and identified second-order control model (first-floor office)	110
4.20	p -step (6 hours) ahead predicted output for 2^{nd} order models of first-floor office	110
4.21	Bode plot comparison of simulation platform and identified second-order control model (first floor seminar room)	111
4.22	p -step (6 hours) ahead predicted output for 2^{nd} order models of first-floor seminar room (models derived using poorly excited data)	112
4.23	p -step (6 hours) ahead predicted output for 2^{nd} order models of first-floor seminar room derived using data from high order simulation platform	113
5.1	Zone temperature using MPC with different prediction horizons for the ground-floor meeting room (2 day period)	126
5.2	Input ($T_{fl_{sp1}}$) determined using MPC with different prediction horizons (ground-floor, 2 day period)	126
5.3	Zone temperature using re-tuned MPC with different prediction horizons (ground-floor meeting room - 2 day period)	127
5.4	Input ($T_{fl_{sp1}}$) determined using re-tuned MPC with different prediction horizons (ground-floor - 2 day period)	127
5.5	Accumulated set-point deviation across full building for different prediction horizons and tuning parameters	128
5.6	Accumulated energy consumption of building for different prediction horizons and tuning parameters	128
5.7	Distributed MPC - Problem is divided into manageable subproblems	131
5.8	Control hierarchy of approach using hot and cold representative zones to simplify optimisation	132
5.9	Accumulated set-point deviation across the Nimbus Centre model using hierarchical MPC formulation and a more standard weather compensated formulation	133
5.10	Accumulated energy consumption of the Nimbus Centre model using hierarchical MPC formulation and a more standard weather compensated formulation	133

5.11	Relationship between the steady-state set-point deviation and the tuning weight for ground floor meeting room	136
5.12	6 ground-floor zone temperatures for three day period using a tuning strategy based on the steady-state set-point deviation	137
5.13	Combined Pareto optimal solution for two objectives	138
5.14	Lexicographic optimal solution for two prioritised-objectives	139
5.15	Three-layer heating system control hierarchy	143
5.16	Input and output power for different values of T_h for <i>Viessmann Virocrossal 200</i> series boiler	144
5.17	Example boiler output as a function of boiler input & header temperature for <i>Viessmann Virocrossal 200</i> series boiler	145
5.18	Boiler output as a function of boiler input & header temperature	146
5.19	Three-layer control hierarchy	147
5.20	Energy consumption for the Nimbus Centre simulation platform using prioritised MPC vs. Weather Compensation (10 days, Feb 2015)	151
5.21	Temperature set-point deviation in the Nimbus Centre simulation platform using prioritised MPC vs. Weather Compensation (10 days, Feb 2015)	152
5.22	Energy consumption change as acceptable comfort range is increased using prioritised model predictive control strategy	153
5.23	Simulated zone temperature obtained using weather-compensation and prioritised MPC for 10 day period (first-floor administration office of Nimbus Centre)	153
5.24	Projected energy consumption over period of 11.5 hours for different acceptable comfort ranges	154
6.1	Optimal input determined using zone models of different accuracy	163
6.2	Difference between the inputs calculated with low-accuracy and with high-accuracy models at each point of the horizon	164
6.3	First element of optimal solution to unconstrained QP problem for different horizons with high-accuracy and low-accuracy models	165
6.4	Difference between ground-floor meeting room output determined using low-accuracy and high-accuracy zone models within MPC formulation	166
6.5	External temperature measurement and forecast for 66 hour period, as well as the corrected forecast obtained at time $k = 100$	167
6.6	Difference between the inputs calculated with correct and with incorrect weather forecast information at each point of the horizon	168
6.7	Difference between inputs calculated with perfect prediction of disturbance estimate and prediction using Option 1 (future disturbance = 0)	170
6.8	Difference between inputs calculated with perfect prediction of disturbance estimate and prediction using Option 2 (future disturbance = current estimated disturbance)	170

6.9	Difference between inputs calculated with perfect prediction of disturbance estimate and prediction using Option 3 (future disturbance = EWMA of previous 20 samples)	171
6.10	Difference between inputs calculated with perfect prediction of disturbance estimate and prediction using Option 4 (future disturbance = profile from previous day)	171
6.11	Diagram of sensor fault detection and isolation process using PCA-based techniques with measured data	176
6.12	Q-statistic showing detection of sensor faults in two zones	177
6.13	Q-contribution plot showing the proportion of total data variance contributed by each sensor in the unfaulted scenario	177
6.14	Q-contribution plot showing the proportion of total data variance contributed by each sensor with fault in Sensor 9	178
6.15	Q-contribution plot showing the proportion of total data variance contributed by each sensor with fault in Sensor 3	178
6.16	Q-statistic showing detection of sensor faults in two zones when disturbed data is used	178
6.17	Q-contribution plot determined with disturbances present in the data - unfaulted scenario	179
6.18	Q-contribution plot determined with disturbances present in the data - fault in Sensor 9	179
6.19	Q-contribution plot determined with disturbances present in the data - fault in Sensor 3	180
6.20	k-Nearest Neighbour and k-Nearest Prototype classification methods for classifying new data-point	182
6.21	Diagram of actuator fault detection and identification process using Kalman filtering and machine learning	182
6.22	Training clusters developed for classification of actuator faults - $\ell = 3$	183
6.23	Training clusters - largest deviation from model estimate vs. second largest	183
6.24	Training clusters - largest deviation from model estimate vs. third largest	184
6.25	Training clusters - second largest deviation from model estimate vs. third largest	184
6.26	Controlling 2 zones - Reconfiguration in Fault Scenario	189
6.27	Simulated temperature measurement of ground floor meeting room with and without the presence of a zone-level fault in corridor - Standard quadratic MPC approach	191
6.28	Simulated temperature measurement of ground floor meeting room with and without the presence of a zone-level fault in corridor - Prioritised MPC approach	191
A.1	Boiler efficiency curves (<i>Viesmann Vitocrossal 200</i>) given for different return temperatures	A1
B.1	Probability density functions of hypothetical wavefunction for different diversity parameters	B2

List of Tables

3.1	Model error present before and after chemotaxis adjustment using test network for different algorithm configurations	45
3.2	Model error present before and after chemotaxis adjustment using test network with unmeasured disturbances present in the system	52
3.3	Model error present before and after chemotaxis adjustment using test network with disturbance estimation	61
3.4	Model error present before and after PSO adjustment incorporating disturbance estimation using test network	63
3.5	Model error present before and after Q-PSO adjustment using test network	65
3.6	Model error present before and after decentralised Q-PSO adjustment using test network	66
3.7	Mean daily output error present before and after Q-PSO adjustment of RC-network compared to data measured in Nimbus Centre	74
4.1	Coefficient of determination for p -step ahead predicted outputs of 1^{st} – 4^{th} order models of Zone 1	100
4.2	Coefficient of determination for p -step ahead predicted outputs of 1^{st} – 4^{th} order models of Zone 1 (zone models derived in presence of disturbance)	102
4.3	Coefficient of determination for 72-step ahead predicted outputs of 2^{nd} order models of Zone 1 derived without disturbances present, with disturbances present and with disturbances present and estimated	106
4.4	Coefficient of determination for 72-step ahead predicted outputs of 2^{nd} order models of Zone 4	107
4.5	Coefficient of determination for p -step ahead predicted outputs of 2^{nd} order models of first-floor office derived using data from Nimbus Centre with and without disturbance estimation	111
4.6	Coefficient of determination for p -step ahead predicted outputs of 2^{nd} order models of first-floor seminar room derived using data from Nimbus Centre with and without disturbance estimation	112
6.1	Difference between first element of derived input sequence calculated with perfect prediction of disturbance estimate and those calculated with different disturbance prediction methods	172
6.2	Accuracy of fault identification for unfaulted scenario, zone-valve fault scenario and mixing-valve fault scenario	185
6.3	Possible actions to be taken when the system is affected by a particular fault	188

I, Edward O'Dwyer, certify that this thesis is my own work and I have not obtained a degree in this university or elsewhere on the basis of the work submitted in this thesis. All external references and sources are clearly acknowledged and identified within the contents. I have read and understood the regulations of University College Cork concerning plagiarism.

Edward O'Dwyer

Acknowledgements

First and foremost, I'd like to thank my supervisor Gordon for help in all regards over the years of research. The substantial amount of time offered up was greatly appreciated, while the technical education given was outstanding, not to mention the general advice on how to actually approach a PhD, which seemed to make everything run far more smoothly.

Having completed this work as part of the Irish Research Council Enterprise Partnership Scheme, I would also like to acknowledge the funding that was provided both by the IRC itself and by the industrial partner United Technologies Research Centre Ireland (UTRCI). This leads to my industrial supervisors through the course of the work carried out from UTRC, namely Luciano, Kostas and Marcin. I'm extremely grateful for the guidance and advice that they provided, as well as the significant amount of time that they were always more than willing to offer.

At this point, I have worked alongside too many people in the Electrical Engineering building to write a list of names without leaving people out, but I am genuinely appreciative of everybody's influence. In particular, the ever-unpredictable day-to-day happenings in the lab made the whole project a thoroughly enjoyable experience. Other friends from other walks-of-life should also be mentioned for their support and the necessary distractions that were always provided.

Finally, I would like to thank my family. Between Mam, Dad, Cliodhna and Donal, there is plenty of experience with type of pursuit. True to form, support and encouragement was never in short supply in any circumstance, while advice was always readily available.

Abstract

Model predictive control (MPC) has often been referred to in literature as a potential method for more efficient control of building heating systems. Though a significant performance improvement can be achieved with an MPC strategy, the complexity introduced to the commissioning of the system is often prohibitive. Models are required which can capture the thermodynamic properties of the building with sufficient accuracy for meaningful predictions to be made. Furthermore, a large number of tuning weights may need to be determined to achieve a desired performance. For MPC to become a practicable alternative, these issues must be addressed.

Acknowledging the impact of the external environment as well as the interaction of occupants on the thermal behaviour of the building, in this work, techniques have been developed for deriving building models from data in which large, unmeasured disturbances are present. A spatio-temporal filtering process was introduced to determine estimates of the disturbances from measured data, which were then incorporated with metaheuristic search techniques to derive high-order simulation models, capable of replicating the thermal dynamics of a building. While a high-order simulation model allowed for control strategies to be analysed and compared, low-order models were required for use within the MPC strategy itself. The disturbance estimation techniques were adapted for use with system-identification methods to derive such models.

MPC formulations were then derived to enable a more straightforward commissioning process and implemented in a validated simulation platform. A prioritised-objective strategy was developed which allowed for the tuning parameters typically associated with an MPC cost function to be omitted from the formulation by separation of the conflicting requirements of comfort satisfaction and energy reduction within a lexicographic framework. The improved ability of the formulation to be set-up and reconfigured in faulted conditions was shown.

Chapter 1

Introduction

1.1 Motivation

The building energy sector has been widely recognized as a significant contributor to global energy consumption and as such, the effects of human-influenced climate-change. While the stated figures can vary, the contribution of buildings to global energy consumption is given in (Pérez-Lombard et al. 2008) as lying in a range of 20-40%, while in (Harish & Kumar 2016) the figure is stated to be approximately 40%. More generally, in (IEA 2015), the International Energy Agency state that the services and households sector is responsible for 35% of energy consumption. An inevitable consequence of energy consumption on this scale is a contribution to global CO_2 emissions of approximately 30% (Shaikh et al. 2014), with the contribution to total EU CO_2 emissions given as 36% (EU 2010).

A breakdown of how energy is used in buildings (for commercial and residential buildings in the U.S.) is given in (Harish & Kumar 2016), noting that in commercial buildings, 45% of energy consumption is associated with space heating, while a further 9% is accounted for by space cooling. In (Costa et al. 2013), the energy saving potential in the building domain is given in the range of 5% to 30%. Any improvements in terms of energy efficiency made in this sector could potentially have a significant impact (a figure of 22.6Mtoe is indicated for the EU in (Ahern & Norton 2015)).

The desire to reduce energy consumption, limit the effects of climate-change and improve energy security in the global context of diminishing fossil fuel supplies has lead to an intensification of research in energy efficient buildings. While the need for large-scale improvements is clear and stricter building regulations have encouraged better

insulation and more efficient equipment, it is shown in (Peeters et al. 2008) that typically, modern heating systems are not used efficiently and are not adjusted to meet the needs of changing conditions. The issue is further emphasised in (Liao et al. 2005), in which it is stated that work in the development of energy efficient systems is largely carried out on isolated components of the overall system, without considering the efficiency of the system as a whole.

The widespread proliferation of wireless sensor networks provides cause for optimism. With appropriate use of sensor data, each individual component of the system can be coordinated to achieve a more energy efficient performance. Just as the availability of energy efficient equipment does not guarantee an energy efficient performance however, the availability of a larger quantity of measurement data alone is not sufficient without effective methods for its use. As such, different control techniques have been developed, suitable for the application of building energy. In (Dounis & Caraiscos 2009), intelligent learning-based control approaches are reviewed, which can determine the best control action to take given a particular set of conditions without the need for the underlying relationship between inputs and outputs to be explicitly known. Fuzzy logic control (Paris et al. 2010), neural-networks (Ferreira et al. 2012) and multi-agent based systems (Yang & Wang 2013) have been outlined as potential options. These can be trained to suit the problem at hand using techniques such as reinforcement learning (Dalamagkidis et al. 2007) and evolutionary algorithms (López et al. 2004) implemented using measured data.

The benefits of improved building energy control strategies can be seen from different perspectives. For the user, better thermal regulation in the building can enhance comfort while additionally, control decisions can be taken to reduce the operational cost. Though the initial cost of installation of a more advanced control strategy may be higher, money can be saved over time (with a payback period highly dependent on the strategy used). From an environmental perspective, the general implementation of energy efficient control strategies can allow for buildings to be thermally regulated with less fuel consumption and as such, less greenhouse gas emission. All of these improvements can be achieved by reacting intelligently to the requirements of the building at any given time, avoiding overheating and overcooling while running equipment at efficient operating points.

To enable the control action to be (in some sense) optimally chosen, while taking the future predicted conditions of the building into account, as well as the present measurements, model predictive control (MPC) approaches have been widely considered as a promising alternative (Ma et al. 2012, Liao & Dexter 2010). Using an MPC for-

mulation, an optimal control input trajectory is calculated over a specified prediction horizon at each measurement sample by solving a constrained optimisation problem. The objective of the optimisation is to minimise (or maximise, depending on the perspective taken) some cost function which describes a particular set of performance criteria.

Design of an MPC formulation requires suitable models to be developed which represent the thermodynamic properties of the system, as well as the internal and external processes which affect it. Furthermore, the control strategy itself must be formed in a way that can be set up and adjusted in a straightforward manner. This implies that it should be possible to characterise a desired set of performance criteria, which is tangible to a typical operator or user of the building, within a cost function.

When applying MPC methods to the thermal control of a building, a number of additional factors must be taken into account, including the following:

- A building is comprised of many rooms (or zones), each with a corresponding set of thermal dynamics and comfort objectives which typically must all be considered by a single control strategy.
- In addition to the thermodynamics of the building itself, many external forces can affect the system, which should be taken into account. People inside the building for example, and equipment such as computers can emit significant quantities of heat. Weather conditions can also have a large impact.
- Each occupant may consider the notion of thermal comfort differently and so a large degree of flexibility should be present.
- Given that the full system includes many individual pieces of equipment (such as valves and sensors), faults and failures occurring to different components must be considered. Preferably, in many cases an individual fault should not result in the complete system being compromised.

Though an extensive range of literature is available in which MPC is shown to outperform traditional approaches for building energy control, as yet, examples of its use in practice are rare (the Desigo predictive heating controller by Siemens being an exception). In (Sturzenegger et al. 2015), it is noted that at the present time, the cost savings available by using an MPC-based strategy are not sufficient to make it a financially competitive technology, given the additional commissioning costs involved. Traditional control approaches can appear more attractive as they are easier to set-up without requiring a high level of knowledge of the specific building, or of advanced

control in general. While the efficiency may be inferior, a definitive quantification of the improvements that would be available with a more sophisticated strategy can be a little abstract. Furthermore, the traditional technologies have been in use for many years and are well understood.

To improve the prospects of MPC as a technology in this domain, simplicity is required from two perspectives. First of all, it must be possible to compute a solution to the numerical optimisation at the heart of the MPC formulation on-line and in an efficient manner. Secondly, perhaps a larger obstacle is the need to set up and adjust the strategy with a low level of commissioning effort.

The crux of the matter lies in resolving this required simplicity with the actual complexity of the problem. Large building structures can be thermodynamically intricate, with a wide range of possible sizes and configurations. Developing models which can accurately capture the dominant dynamics, while maintaining the required simplicity, can necessitate an extensive modelling effort by a practitioner with a high level of expertise. Beyond the need for appropriate models, the MPC strategy itself may be comprised of different and competing thermal and energy objectives associated with many zones in the building. This can result in a formidable tuning problem, which would require re-computation if any adjustments were to be implemented. To become a realistic alternative to traditional methods on a large scale, these issues must be addressed.

1.2 Research Objectives

This work seeks to develop MPC strategies suitable for implementation within a hydronic building heating system control strategy. The focus of the design process is to incorporate the performance benefits that are often attributed to MPC with an improved level of scalability and flexibility.

To achieve a greater level of scalability, methods for developing models with low commissioning effort must be considered. More specifically, modelling techniques that can be used without the need for a comprehensive a-priori knowledge of the system are required, while also acknowledging the fact that any data measurements used for modelling are likely to be corrupted by the presence of significant quantities of unmeasured disturbance, caused by weather and building occupants.

The problem of tuning must also be considered. Given the large number of possible zones and objectives present, the requirement of a substantial number of suitable tun-

ing weights for each building in which the strategy is implemented is far from ideal. Formulations in which the cost functions can be characterised in a way that precludes the need for this type of tuning, while maintaining the flexibility needed for adjustments to be made are investigated. An ability to reconfigure the formulation in faulted conditions is also a desirable aspect which should be incorporated.

1.3 Thesis Outline

In Chapter 2, a background of technologies used in the building energy domain is provided. Typical building energy systems are outlined, along with a description of some of the limitations associated with the control techniques most commonly used at present. Methods for exploiting the increased availability of sensor data are introduced, including learning-based techniques which can implicitly handle more complex relationships between user comfort and thermal load requirement than would be possible using traditional, classical control methods. Predictive strategies are outlined which can base current control decisions on future thermal load requirements by using models of the system and forecasts of external and internal conditions. Model predictive control (MPC) is introduced, as well as the ways in which it has been applied to building energy systems.

Chapter 3 outlines a novel set of techniques which can be used to develop detailed mathematical representations of the thermal properties of a building using measured data. By representing the thermal dynamics of a building in a simulation platform, control strategies can be compared more easily, as consistent external and internal conditions can be simulated for different experiments. Parameter identification procedures using different metaheuristic search algorithms are analysed for the adjustment of the parameters of a high-order RC-network-based building model. A novel statistical methodology for estimating unmeasured disturbances by exploiting the multi-zone nature of buildings is introduced and included in the model identification procedures to allow for the biasing impact of disturbances to be reduced. As a case study, a model is derived to represent the thermal dynamics of a two-storey office building (specifically the Nimbus Centre), validated using measured data from the building.

In Chapter 4, building modelling is considered from a different perspective. While detailed simulation models are useful for comparison and analysis of control strategies, a far greater level of simplicity is needed for use within an effective numerical optimisation-based control strategy such as MPC. The separation of the high-order

simulation model into lower-order zone models is considered. Model reduction techniques and data-based techniques are investigated in terms of accuracy, complexity and identifiability. The disturbance estimation techniques of Chapter 3 are then expanded to allow for integration with data-based system identification methods to form a new method for deriving zone models.

With an appropriate set of models derived, as well as a validated simulation platform in which strategies can be tested, novel MPC formulations are developed in Chapter 5 for use in building energy systems. In addition to achieving an improved thermal performance with better energy efficiency, the primary focus of the developed MPC strategies is scalability. Methods for reducing the problem complexity and determining appropriate tuning parameters are derived and simulated using the simulation platform. A new MPC formulation using prioritised-objectives is then designed which can be set up without the need for an extensive tuning procedure. By using a lexicographic framework, the need for tuning parameters is removed. The formulation is set up to be flexible enough to allow a user, uninitiated in advanced control theory, to make adjustments. The developed strategy is considered to be a practical alternative to traditional strategies.

In Chapter 6, the performances of MPC formulations are considered in scenarios in which the dynamics of the building do not evolve in the predicted manner, either due to inaccurate characterisation of the underlying models and thermodynamic processes, or due to faults or system changes occurring unexpectedly. The effect of inaccurate modelling is first considered by comparing scenarios in which high-accuracy predictions of all processes are available with those in which inaccuracies are present. Methods for predicting unmeasured disturbances are then developed and compared. Techniques for fault detection are outlined. A model-free, statistical approach for sensor-fault detection using PCA is developed and simulated as well as a novel model-based method for detection of actuator faults, which combines Kalman filtering with simple machine learning techniques. The reconfigurable properties in faulted conditions of the prioritised-objective formulation (developed in Chapter 5) are compared to those of more standard quadratic MPC approaches. It is shown that the prioritised approach can be beneficial in such circumstances.

References

- Ahern, C. & Norton, B. (2015), 'Energy savings across EU domestic building stock by optimizing hydraulic distribution in domestic space heating systems', *Energy and Buildings* **91**, 199–209.
- Costa, A., Keane, M. M., Torrens, J. I. & Corry, E. (2013), 'Building operation and energy performance: Monitoring, analysis and optimisation toolkit', *Applied Energy* **101**, 310–316.
- Dalamagkidis, K., Kolokotsa, D., Kalaitzakis, K. & Stavrakakis, G. S. (2007), 'Reinforcement learning for energy conservation and comfort in buildings', *Building and Environment* **42**(7), 2686–2698.
- Dounis, A. I. & Caraiscos, C. (2009), 'Advanced control systems engineering for energy and comfort management in a building environment-A review', *Renewable and Sustainable Energy Reviews* **13**(6-7), 1246–1261.
- EU (2010), 'Directive 2010/31/EU of the European Parliament and of the Council of 19 May 2010 on the energy performance of buildings (recast)', *Official Journal of the European Union* pp. 13–35.
- Ferreira, P. M., Ruano, a. E., Silva, S. & Conceição, E. Z. E. (2012), 'Neural networks based predictive control for thermal comfort and energy savings in public buildings', *Energy and Buildings* **55**, 238–251.
- Harish, V. & Kumar, A. (2016), 'A review on modeling and simulation of building energy systems', *Renewable and Sustainable Energy Reviews* **56**, 1272–1292.
- IEA (2015), 'Key World Energy Statistics 2015'.
- Liao, Z. & Dexter, A. (2010), 'An Inferential Model-Based Predictive Control Scheme for Optimizing the Operation of Boilers in Building Space-Heating Systems', *IEEE Transactions on Control Systems Technology* **18**(5), 1092–1102.

- Liao, Z., Swainson, M. & Dexter, A. L. (2005), 'On the control of heating systems in the UK', *Building and Environment* **40**(3), 343–351.
- López, A., Sánchez, L., Doctor, F., Hagrass, H. & Callaghan, V. (2004), An evolutionary algorithm for the off-line data driven generation of fuzzy controllers for intelligent buildings, in 'Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics', Vol. 1, pp. 42–47.
- Ma, Y., Kelman, A., Daly, A. & Borelli, F. (2012), 'Predictive control for energy efficient buildings with thermal storage', *IEEE Control Systems Magazine* **32**(February), 44–64.
- Paris, B., Eynard, J., Grieu, S., Talbert, T. & Polit, M. (2010), 'Heating control schemes for energy management in buildings', *Energy and Buildings* **42**(10), 1908–1917.
- Peeters, L., Van der Veken, J., Hens, H., Helsens, L. & D'haeseleer, W. (2008), 'Control of heating systems in residential buildings: Current practice', *Energy and Buildings* **40**, 1446–1455.
- Pérez-Lombard, L., Ortiz, J. & Pout, C. (2008), 'A review on buildings energy consumption information', *Energy and Buildings* **40**(3), 394–398.
- Shaikh, P. H., Nor, N. B. M., Nallagownden, P., Elamvazuthi, I. & Ibrahim, T. (2014), 'A review on optimized control systems for building energy and comfort management of smart sustainable buildings', *Renewable and Sustainable Energy Reviews* **34**, 409–429.
- Sturzenegger, D., Gyalistras, D., Morari, M. & Smith, R. S. (2015), 'Model Predictive Climate Control of a Swiss Office Building : Implementation , Results , and Cost – Benefit Analysis', *IEEE Transactions on Control Systems Technology* pp. 1–12.
- Yang, R. & Wang, L. (2013), 'Development of multi-agent system for building energy and comfort management based on occupant behaviors', *Energy and Buildings* **56**, 1–7.

Chapter 2

Technologies used in Building Energy Systems

2.1 Introduction

In this chapter, an overview is provided of some of the different methods which can be used to control building energy systems so as to achieve improvements in both energy efficiency and thermal comfort for occupants. Traditional approaches are first considered and compared to intelligent, learning-based approaches, which can allow for the misalignment between the complexities of the building structure and the simplicity of the classical linear control laws to be accounted for without an extensive modelling effort. Model based approaches which can exploit predictions of the future system states in the current control action are then outlined. A description of constrained optimisation and model predictive control (MPC) is given, with an overview of how they have been applied to the problem of building energy control in literature.

As the need to reduce energy consumption becomes clearer, many advancements in the building envelope have occurred in recent years in terms of both construction materials and design. The increased use of wireless sensor networks has provided the opportunity for more advanced techniques to be used for the control of building energy systems. With appropriate control methods, the additional data obtained from sensors can allow for some of the complexity present in the thermal requirements of a building to be allowed for in a way that was not previously possible with traditional techniques. Intelligent control strategies using fuzzy-logic control, neural networks and multi-agent systems can learn over time the appropriate control decisions to be made when a particular set of measurements is observed. These methods do not re-

quire the underlying dynamics of the system to be explicitly known, which, given the complexity of a large building structure, can be quite advantageous.

If a model of the system can be obtained however, optimal control techniques can be used to further improve energy efficiency by determining a set of control inputs which can achieve some desired performance criteria in the future. This predictive approach allows for the slow thermal response of the building to be taken into account, as well as allowing for weather forecasts and fuel-price forecasts to be incorporated. The most popular technique in the literature for optimal building energy control is model predictive control (MPC), which solves a constrained optimisation problem over a specified prediction horizon. These techniques, and how they are applied to buildings are outlined here.

2.2 Control of Building Energy Systems

2.2.1 Passive and Active Energy Efficiency

As the link between the burning of fossil fuels and the effects of climate change becomes evermore compelling (IPCC 2014), the need for large scale reductions in building energy consumption is clear. Regulations have been introduced to reflect this (Pérez-Lombard et al. 2011), while many technologies capable of providing a significant improvement in energy efficiency exist within both the building envelope and the heating/cooling system.

At the design stage, buildings are now often designed with passive energy efficiency in mind. Depending on the climate in which the building is located, walls can be constructed in such a way as to prevent heat transfer, to increase latent heat storage or to trap and transfer solar heat energy (Sadineni et al. 2011). Similarly, glazing can be chosen to suit the particular requirements of a given climate. Options range from Aero-gel glazing, which offers improved insulation properties, to windows with switchable properties such as electrochromic devices, liquid crystal devices or suspended-particle devices (Baetens et al. 2010), which can adapt to meet the needs of specific situations. The external form and internal layout of the building can also be designed to maximise the capture of solar energy as well as optimising airflow and ventilation through the building.

In terms of active energy efficiency, more efficient heating and cooling equipment also exists, though it is shown in (Peeters et al. 2008) that suboptimal control of the equip-

ment can often lead to inefficient operation. A system may be designed to be energy efficient, but may not be correctly adjusted to meet the specific load requirements of changing conditions. Advancements in control technology however, along with the proliferation of wireless sensor networks have enabled new building energy control methods to be brought into consideration (Guinard et al. 2009).

In addition to temperature sensing, occupancy can be measured using passive infrared (PIR) sensors, which measure the infrared radiation emitted by the human body, while pyranometers can measure solar radiation flux density. With measurements available for internal temperatures, external temperatures, solar irradiance and occupancy, as well as up-to-date forecasts of external weather conditions, more sophisticated control decisions can be taken. Moving a step further, using a combined heat and power (CHP) system, it is possible to integrate the thermal requirement of the building with the electrical load to obtain an overall more energy efficient or cost effective performance (Houwing et al. 2010).

By coordinating a distributed set of controls, a complete strategy can seek to control thermal comfort, air quality, lighting and various other electrical loads with a minimum financial or energy cost (Shaikh et al. 2014). A possible architecture, incorporating a boiler to supply heat, a CHP to supply heat and power (along with the grid) as well as adjustable glazing is shown in tandem with a sensor network providing illuminance, occupancy, temperature and solar radiation measurements in Fig. 2.1. While the ability to communicate easily between sensors and equipment should result in an improved control performance, perhaps of greater importance is the way in which the data is used. Knowledge of the state of the system alone is not beneficial. Appropriate methods for using the knowledge to make control decisions are required. In the following section, some of the current methods used to control a building heating system to satisfy the building's thermal requirements are outlined.

2.2.2 Traditional Approaches for Thermal Control of Buildings

To achieve thermal comfort in buildings, the most common technologies are heating ventilation and air conditioning (HVAC) systems, more simple hydronic radiator heating systems and electrical storage heater-based heating systems. In large buildings, HVAC systems are often favoured, as the air temperature and quality can be controlled simultaneously. In such a system, heat can be generated by boilers or CHPs powered by any of a variety of fuels, or by heat pumps. The heat is distributed to the building by water-based or air-based means, with the former typically employing a system of

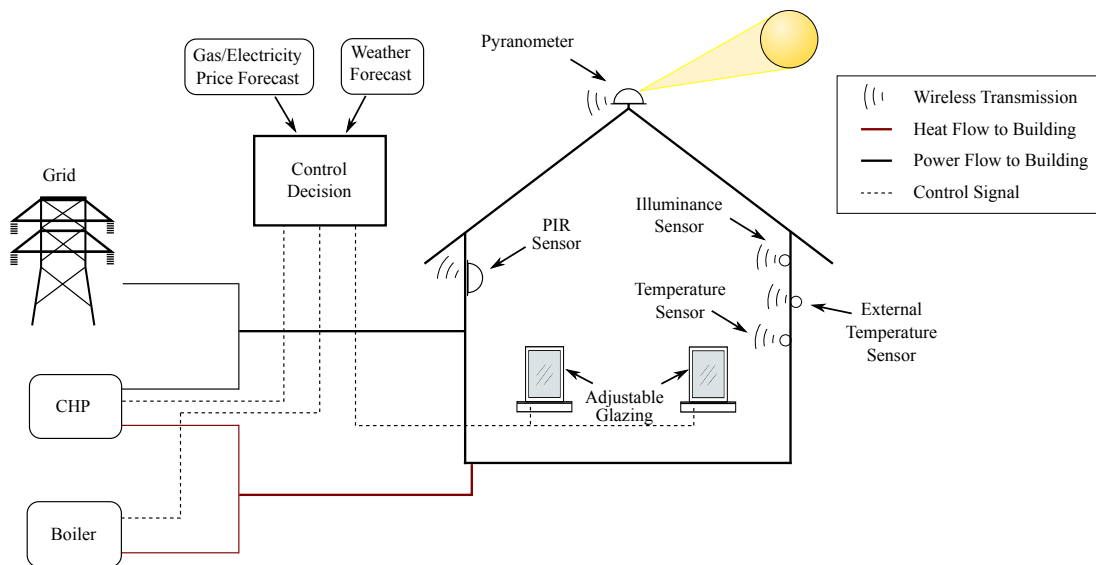


Figure 2.1: Active energy efficiency, using wireless sensors to determine electrical and thermal load requirements

radiators around which heating water is pumped, while in the latter, air is distributed through ductwork to different parts of the building. The ventilation element of the system can be natural (through louvres or windows) or forced using an air-handling unit (AHU). The objective of the ventilation component is to provide sufficient fresh air for occupants with guidelines provided in (CIBSE (Chartered Institution Building Services Engineers) 2006). The air conditioning component of the HVAC system can provide both cooling and humidity control.

In certain climates, cooling may not be paramount and a stand alone heating system may suffice. Using a hydronic heating system, heat supply to the building is regulated by the control of valves at different levels of the system. At zone level, thermostatic radiator valves (TRVs) can be used, which open and close depending on the temperature of the room in which they are located. At a higher level, mixing valves can control the temperature of water flowing to the radiators of multiple zones by mixing supply water, coming from the heating system, with return water coming from the zones. In modern office buildings, such heating systems are typically controlled centrally by a building management system (BMS). The BMS communicates with the heating/cooling system equipment, monitoring data measurements and sending control set-points.

Traditional control strategies are often carried out using empirically determined schedules, on/off controllers and classical control techniques such as proportional-integral (PI) or proportional-integral-derivative (PID) controllers, all of which adjust the equipment in the system so as to achieve the desired set-points. These set-points can be manually chosen, or, in modern systems, measurements of the external temperature

can be used to determine weather compensated set-points. In this way, when the external temperature increases, to reflect the fact that the heat requirement of the building should in some way decrease, the temperature set-points of the system should then be reduced. A schematic of a hydronic system with a PID controlled mixing valve with a weather compensated set-point delivering water from a heating system to multiple zones is shown in Fig. 2.2.

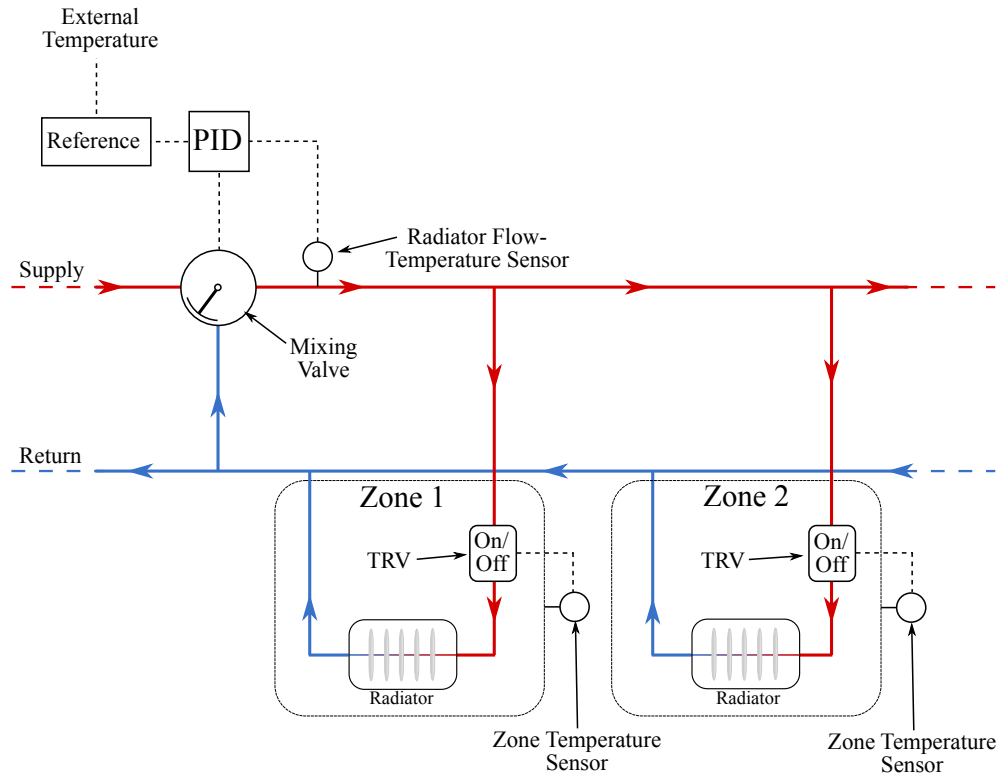


Figure 2.2: Weather-compensated mixing valve in hydronic heating system

To account for the anticipated set-point changes that occur over time (occupancy hours in an office building may have the same profile each day for example), empirically derived schedules tend to be used. Such schedules make an assumption that little change in load requirement will be observed from day to day. To improve the performance in the presence of measured disturbance, a feed-forward strategy is outlined in (Thomas et al. 2005), however, it is noted that a high-level of sensitivity to modelling error results in such an approach.

The slow, yet complex dynamics associated with a building's thermal behaviour can however make the standard techniques unsuitable (Chen 2002). No knowledge of the system or disturbances is incorporated, while control decisions are based only on the current state of the system. These techniques can then only react to changes in load requirement after the change has occurred. By the time the control has had the desired effect, a significant period of over-heating, under-heating or inefficient operation may

have transpired. While such approaches have been employed for many years, a better performance may be possible, though often at the expense of additional complexity.

2.2.3 Intelligent Model-Free Approaches for the Control of Building Energy Systems

2.2.3.1 Defining a Desired Thermal Performance

An issue which must be taken into account when considering building energy control, is the subjective nature of thermal comfort. Many factors can impact the thermal sensation of an occupant, which are often neglected in traditional control strategies. If a wide range of sensor data is available, it may be possible to characterise the comfort of an occupant with more detail. A commonly used index for quantifying comfort is the Predicted Mean Vote (PMV), initially developed in the late 1960s by Fanger and widely studied and expanded upon in the interim (Ole Fanger & Toftum 2002, Sherman 1985). PMV is determined based on the Metabolic rate of the occupants, the clothing insulation, relative humidity, air velocity and the mean radiant temperature. Methods for estimating these values can be found in (Ansi/Ashrae 2013, EN-ISO-7730 2005). The PMV (I_{pmv}) index is calculated as a point on a thermal sensation scale, ranging from -3 (cold) to $+3$ (hot), with a recommended operating range of 0 ± 0.5 (Castilla et al. 2011).

A simplified version of the PMV can be calculated as:

$$I_{pmv} = (0.303e^{-0.036M} + 0.028) L \quad (2.1)$$

where M is the metabolic rate and L is the thermal load on the body (which is defined as the difference between the internal heat production of a body and the heat loss to the environment). A related metric is the Predicted Percentage Dissatisfied (PPD), which, as the name suggests, predicts the percentage of occupants that will not be satisfied by a given set of thermal conditions. It is calculated as:

$$I_{ppd} = 100 - 95e^{-(0.03353I_{pmv}^4 + 0.2179I_{pmv}^2)} \quad (2.2)$$

It should be noted that the observed mean vote will not always align with the value predicted by the PMV formula (Humphreys & Fergus Nicol 2002). A certain degree of subjectivity is inevitable.

2.2.3.2 Obtaining a Desired Performance

The primary advantage of the traditional control strategies is simplicity. Without considering the complex thermal dynamics of a building and the disturbances which affect it, a particular control strategy can be easily set-up and scaled to suit a wide-range of building designs. As the application is not safety-critical, an adequate level of performance can be achieved. To achieve significant improvements in efficiency and comfort however, some of the complexity of the dynamics underpinning the control problem must be acknowledged. A review of techniques which attempt to incorporate this complexity with the intelligent use of data is provided in (Dounis & Caraiscos 2009).

Learning-based approaches using fuzzy-logic control (FLC) are proposed in (Paris et al. 2010, López et al. 2004) and (Gouda et al. 2001). Fuzzy logic controllers can map sensor measurements to sets of appropriate fuzzy membership functions that can allow for the state of the full system to be taken into account when calculating a controller action, without requiring a detailed model. By incorporating fuzzy rules with standard PID control, the inputs to the system can be corrected if the building states deviate from the operating region for which the PID control was tuned.

In (Karatasou et al. 2006) and (Morel et al. 2001), artificial neural networks (ANNs) are used to predict loads and develop appropriate control set-points, while in (Liang & Du 2005, Kanarachos & Geramanis 1998), ANNs are used to directly determine controller actions. These learning-based strategies can compensate for the inaccuracies resulting from the constant controller gains used in a classical control strategy, without an intensive modelling requirement.

In (Zhao et al. 2010, Yang & Wang 2013a) and (Yang & Wang 2013b), agent-based control systems are described for the application of building energy control. In artificial intelligence (AI), the best control results are obtained when agents operate rationally. By considering the building controls as a large network of agents, the collective behaviour of the agents can be trained to achieve some desired performance. In (Yang & Wang 2013b) for example, people in the building interact with zone-level agents which in turn interact with local-controller agents. A central agent then coordinates the local controller agents to achieve maximum comfort satisfaction using a particle swarm optimiser. In (Dalamagkidis et al. 2007) and (Anderson et al. 2004), reinforcement learning techniques are used to train agents to provide a particular performance by rewarding certain behaviour. Without needing to pre-define a desired behaviour, these methods can resolve conflicts that occur between competing objectives in different zones.

Beyond the domain of scientific literature, commercial technologies are available for controlling building energy systems which use learning methods to improve efficiency, the Nest Learning Thermostat (a product made by Nest Labs, a subsidiary of Alphabet) being the most prevalent example. The Nest thermostat chooses heating schedules based on historical profiles previously used by the occupants as well as automatically determining unoccupied periods based on motion and light sensors.

Perhaps the most commonly proposed methodology for improved control of building energy system is Model Predictive Control (Shaikh et al. 2014). Predictive strategies have the advantage of being able to incorporate predicted future conditions in the current control action. Given the ever-changing external and internal environments in which a building exists, as well as the slow dynamics with which a building reacts to a change in control input, this can allow for better energy efficiency with reduced discomfort. The main disadvantage is that a model of the thermal properties of the building is required. An outline of MPC is provided in the next section as well as how it has been applied to the application of building energy control.

2.3 Optimal Model-based Control Formulations

Optimal control strategies use mathematical models of a plant to predict the future outputs over a prediction horizon in terms of the inputs that are applied. By minimising a cost function describing some performance aspects of the control strategy subject to constraints associated with the plant model and the allowable range of inputs, an optimal sequence of future inputs is calculated.

Where, at time t , the system inputs are given by $u(t)$ and the system states are given by $x(t)$, a general formulation for the cost function and constraints is as follows:

$$J = M(x(t_f), u(t_f), t_f) + \int_{t_0}^{t_f} \mathcal{L}(x(t), u(t), t) dt \quad (2.3)$$

where t_0 and t_f are initial and final times of the horizon respectively. The constraints can then be typically formulated as:

$$\dot{x} = f(x, u, t) \quad (2.4)$$

$$g(x(t), u(t)) \leq 0 \quad (2.5)$$

$$h(x(t_0), x(t_f)) = 0 \quad (2.6)$$

The system dynamics are described by (2.4), while the set of inequality constraints,

given by (2.5), describes a set of hard constraints on the possible inputs which can be applied to the system as well as the possible outputs that the system can produce. The boundary conditions are given by (2.6).

2.3.1 Solving Constrained Optimisation Problems

For most constrained control problems, analytically determining an optimal solution is not possible, and so some kind of numerical optimisation is required. The type of algorithm used to find a solution is very much dependent on the problem. Broadly speaking, minimisation problems can be separated into convex and non-convex problems. For a function $f(\theta)$ to be convex, it is necessary that for any two candidate solutions $\theta_1 \neq \theta_2$, the following inequality holds:

$$\begin{aligned} \varphi f(\theta_1) + (1 - \varphi)f(\theta_2) &\geq f(\varphi\theta_1 + (1 - \varphi)\theta_2) \\ \forall \varphi &\in [0, 1] \end{aligned} \quad (2.7)$$

A convex function contains no local minima, though the global minimum is not necessarily unique. The test for strict convexity (a function with a unique minimum is strictly convex) is then:

$$\begin{aligned} \varphi f(\theta_1) + (1 - \varphi)f(\theta_2) &> f(\varphi\theta_1 + (1 - \varphi)\theta_2) \\ \forall \varphi &\in (0, 1) \end{aligned} \quad (2.8)$$

In Fig. 2.3, strictly convex, convex and non-convex functions are shown:




Figure 2.3: Convex (strict), convex and non-convex surface plots

A commonly encountered (and particularly straightforward) form of convex problem uses linear or quadratic cost functions with a set of constraints expressed as affine functions; it should be noted however that certain classes of nonlinear problems can

also be convex. Linear problems are classified as LP, while non-linear problems are classified as NLP. A quadratic cost function with linear constraints, though technically an NLP problem, is often classified separately as a QP problem. Methods for solving LP and QP problems are particularly well understood (Nocedal & Wright 1999).

While problems for which the cost function and the constraints are convex are generally the simplest to implement, non-convex problems can often arise. Such problems can be continuous or discontinuous. Discontinuities can occur when variables are required to be assigned value taken from a discrete interval (a switch may be represented as a binary variable for example). These problems are classed as mixed-integer programs, which can be linear (mixed-integer linear programming - MILP) or non-linear (mixed-integer non-linear programming - MINLP).

If a cost function and the set of associated constraints is convex, a solution is optimal if it satisfies the Karush-Kuhn-Tucker (KKT) conditions (provided a number of regularity conditions are also met). Satisfaction of the KKT conditions is both sufficient and necessary for optimality. To illustrate the conditions, the constrained optimisation problem of (2.3)-(2.6) is now considered as a minimisation problem where the optimisation variables are given as θ :

$$\theta^* = \arg \min J(\theta) \quad (2.9)$$

s.t.

$$g_i(\theta) \leq 0 \quad (2.10)$$

$$h_j(\theta) = 0 \quad (2.11)$$

$$\forall i = 1, \dots, m_{leq}$$

$$\forall j = 1, \dots, m_{eq}$$

where m_{leq} and m_{eq} are the number of inequality and equality constraints respectively. The Lagrangian of the problem is then given as:

$$\mathcal{L}(\theta, \rho, \lambda) = J(\theta) + \rho g(\theta) + \lambda h(\theta) \quad (2.12)$$

where ρ and λ are the Lagrange multipliers associated with the inequality and equality constraints respectively. It should be noted that if no inequality constraints exist, a solution can be found by the method of Lagrange multipliers. The first KKT condition, to guarantee stationarity, requires that the gradient of the Lagrange function of the

problem be equal to 0:

$$\nabla J(\boldsymbol{\theta}) = - \sum_{i=1}^{m_{leq}} \rho_i \nabla g_i(\boldsymbol{\theta}) - \sum_{j=1}^{m_{eq}} \lambda_j \nabla h_j(\boldsymbol{\theta}) \quad (2.13)$$

This condition reflects the fact that at a stationary point, the gradients of the objective and the constraints are parallel. The second KKT condition, for primal feasibility is as follows:

$$g_i(\boldsymbol{\theta}) \leq 0 \quad (2.14)$$

$$h_j(\boldsymbol{\theta}) = 0 \quad (2.15)$$

$$\forall i = 1, \dots, m_{leq}$$

$$\forall j = 1, \dots, m_{eq}$$

For the dual of the problem to be feasible, the following is required:

$$\rho_i \geq 0 \quad (2.16)$$

$$\forall i = 1, \dots, m_{leq}$$

The final condition, for complementary slackness, is given as:

$$\rho_i g_i(\mathbf{x}^*, \mathbf{u}^*) = 0 \quad (2.17)$$

$$\forall i = 1, \dots, m_{leq}$$

Different methods can be used to find a solution for which these conditions are satisfied depending upon the class of the problem. For LP problems, the Simplex method is widely used, having been originally conceived by George Dantzig in the 1940s. It can be shown that for constrained linear problems, the optimal solution lies on one of the vertices of the feasible polytope. Exploiting this fact, the simplex algorithm iteratively improves the objective by moving along the constraint edges until an optimal solution is achieved.

For larger LP problems, the simplex algorithm may not be suitable as it cannot be shown to scale with polynomial-time (the worst case scenario can scale exponentially (Nocedal & Wright 1999)). Interior-point algorithms have been derived which can be solved more efficiently (Karmarkar 1984), and as such, these methods have often been used in place of simplex methods since the 1980s. Rather than travelling along the edges of the feasible region, an interior-point algorithm (in its original form)

moves through the interior of the feasible region. Many variations of the interior-point algorithm have been derived which can also handle NLP problems, often with the requirement that the objective and constraints are convex and twice continuously differentiable (Wächter & Biegler 2006).

A standard interior-point algorithm follows a primal-dual barrier approach, which seeks to calculate solutions for a sequence of barrier problems of the form:

$$\min_{\theta} \left(f(\theta) - \mu \sum_{i=1}^{m_{eq}} \ln(-g_i(\theta)) \right) \quad (2.18)$$

s.t.

$$h_j(\theta) = 0, \quad \forall j = 1, \dots, m_{eq} \quad (2.19)$$

The logarithmic function acts as a barrier to the search. As an inequality constraint boundary is approached, the cost approaches infinity. Iteratively, the algorithm decreases the barrier parameter μ to allow the search to move towards the constraints. As $\mu \rightarrow 0$, it can be shown that the optimal solution is approached. In Fig. 2.4, the progression of the algorithm through the interior of a constrained search space from an initial solution θ_0 to a solution θ^* as the barrier weight is reduced to 0 is shown.




Figure 2.4: Interior-point algorithm pathway as barrier weight is reduced to 0

An additional method which can be useful for QP problems with low numbers of inequality constraints is the Active-set method. Observing the complementary slackness condition of (2.17), if an inequality constraint is inactive ($g_i(\theta) < 0$), the associated Lagrange multiplier must be equal to zero and the constraint can then be removed

from the Lagrangian. Conversely, if an inequality constraint is active ($g_i(\theta) = 0$), it becomes equivalent to an equality constraint. The active-set method seeks to iteratively improve the solution by activating different combinations of constraints and testing for optimality.

Different methodologies are required to solve MILP problems (Floudas 1995), which if a large number of integer variables exist, can quickly become highly complex. By considering each combination of possible discrete intervals for each of the variables, the problem grows exponentially as the number of integer variables increases. To avoid the need for testing each combination, the branch and bound method is often employed. With this method, the problem is first divided into a set of separable subproblems, each of which is further divided to a lower-level set of problems. By a process of problem-relaxation and fathoming (determining if an improved feasible solution to a relaxed candidate subproblem exists), a subproblem and all lower level divisions coming from the subproblem can be removed from the problem if an improvement is not available. The method allows for a solution to be found without the need to determine the optimal value for every possible combination of variables.

The MINLP class of problem is more challenging again, due to the possible existence of multiple solutions across a potentially large combination of intervals. Methods that exist for solving such problems are outlined in (Floudas 1995), but are not covered here.

2.3.2 Model Predictive Control

2.3.2.1 General Formulation

MPC is an optimal control strategy in which a constrained numerical optimisation problem is typically solved at regular intervals to determine the future control inputs and predicted plant responses which optimise some performance index over a specified prediction horizon. Most commonly, a receding horizon approach is taken. This involves minimising a cost function at each time-step (by numerical optimisation), and applying the first element of the calculated optimal sequence of control inputs to the plant. At the next time-step, the plant states are measured and updated and the optimisation is repeated (Maciejowski 2002, Camacho & Alba 2007). An illustration of the receding horizon concept is shown in Fig. 2.5 for a horizon length of 5 samples.

After originally coming to prominence as a method for solving multi-variable constrained control problems in the process and petrochemical industries, applications for

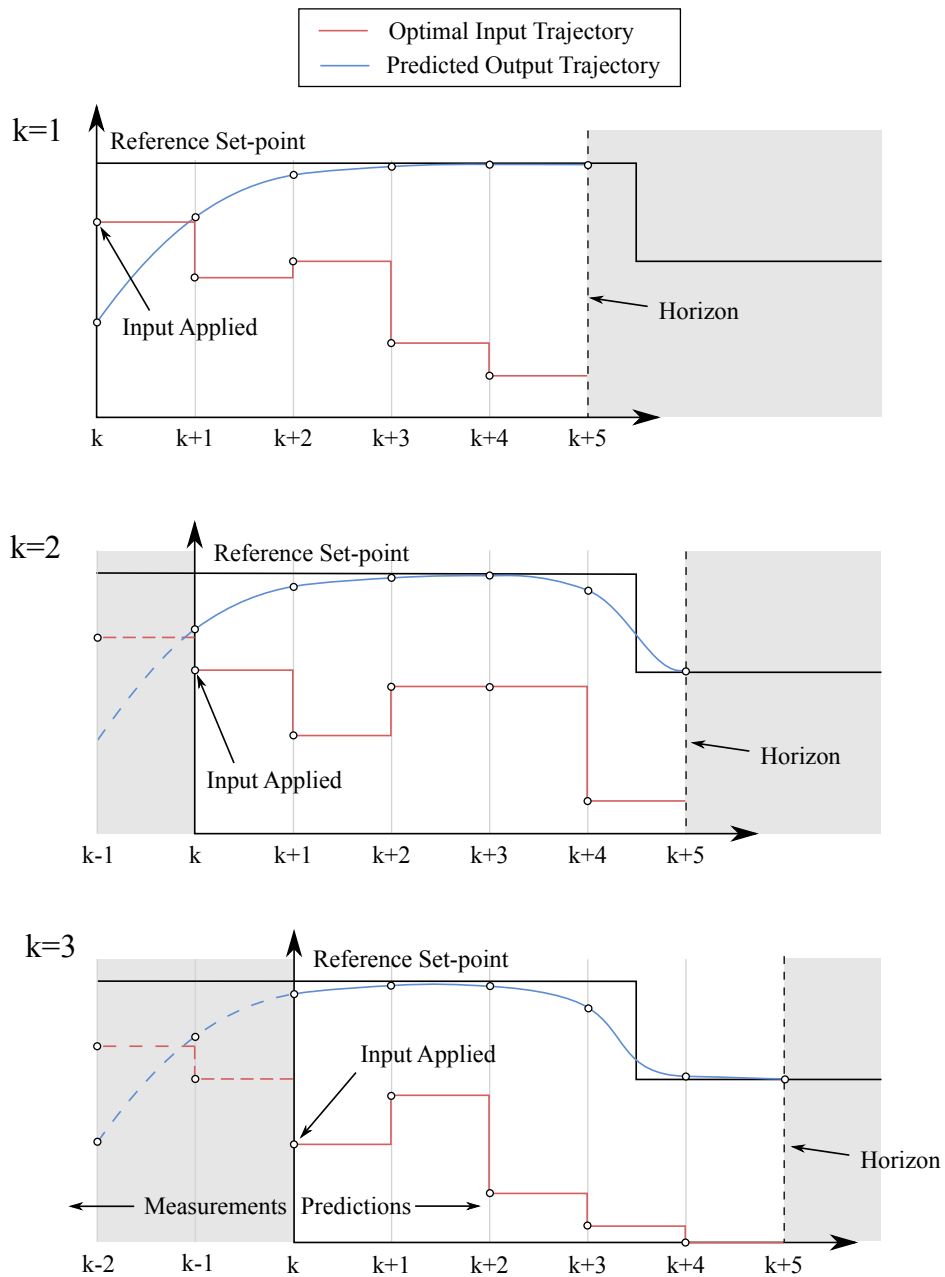


Figure 2.5: Receding Horizon for 3 samples, using a horizon length of 5 samples

MPC have grown more varied (Morari & Lee 1999). By recognising the slow nature of a building's thermal dynamics and using future weather and building usage forecasts, more efficient thermal control can be achieved by use of an MPC strategy (Karlsson & Hagetoft 2011, Abbas et al. 2011, Chen 2002).

To allow for numerical optimisation to be carried out quickly on-line, convex linear or quadratic forms with linear discrete-time models are often favoured. A typical cost function may seek to minimise a combination of a tracking error component and the input increment. Minimising the input increment introduces a type of integral action

to the formulation, removing steady-state tracking error. A terminal penalty can also be included to allow for the finite horizon solution to coincide with what would be obtained using an infinite horizon (assuming a perfect model with no disturbances).

Constraints can be described as hard or soft. Hard constraints define the feasible solution set and must be satisfied by the numerical optimisation. These constraints are used to define limits on inputs and other physical limitations associated with the system. A valve cannot operate in a region outside 0-100% open for example, and so, for a solution to make sense, the set of possible solutions must be hard-constrained to this range. Soft constraints are a little different. It may be desirable for an output variable to remain within a specific range for example, which due to physical constraints on the system, is impossible to achieve. As such, placing a hard constraints on such an output will simply result in an infeasible optimisation problem. A more sensible approach is to soften the constraint by allowing it to be broken, but heavily penalizing the infraction by including an additional penalty term in the cost function. In this way, if a constraint cannot be satisfied, the control should seek to obtain a solution that is as close as possible to the constraint. Classes of MPC that exist which handle constraint uncertainty differently include robust MPC, which allows for some performance specification to be met, given a certain amount of prediction inaccuracy, and stochastic MPC which can allow constraints to be broken with a specified probability.

2.3.2.2 Applying MPC to Building Energy Systems

While many forms of MPC exist, a method is typically chosen to suit the performance index to be optimised as well as the underlying model used for prediction. In the application of building energy control, a number of considerations must be taken into account. First of all, the complex thermodynamic properties of a large building do not easily align with the convexity desired for efficient numerical optimisation (as outlined in Section 2.3.1). Furthermore, even if accurate, detailed models are available, the validity of predictions into the future will still be sensitive to the impact of disturbances. Changing weather outside a building and moving occupants inside a building, even if measured, cannot be predicted with perfect accuracy. As an additional layer of complexity, the performance index must be taken into account. The many zones of a building, allied to the subjective nature of thermal comfort, result in an overall objective which consists of many competing comfort satisfaction and energy reduction objectives.

To make MPC a practicable alternative, what is intrinsically a highly complex prob-

lem must be simplified for numerical optimisation while maintaining a performance improvement over traditional approaches with a competitive level of commissioning and reconfiguration effort. This conflict between complexity (for performance) and simplicity (for practicability) characterises much of the literature on MPC for building energy applications, with the trade-off most explicitly covered in (Ma et al. 2014). Uncertainties associated with the building model are considered in (Maasoumy et al. 2014), while the incorporation of uncertain internal and external disturbance terms are considered respectively in (Oldewurtel et al. 2013) and (Oldewurtel et al. 2012). The issue of large-scale model complexity is considered in (Ma et al. 2011), (Ma et al. 2012) and (Moroşan et al. 2010) through the use of distributed MPC, while in (Hazyuk et al. 2012a) and (Hazyuk et al. 2012b), the models are simplified by methods which allow for nonlinearities associated with radiator heat transfer to be incorporated in an LP problem. An inferential building model is used in (Liao & Dexter 2010) to enable more straightforward commissioning. To handle the many objectives in a characterisable fashion, while acknowledging the uncertainty of the problem, stochastic MPC strategies are developed in (Sadeghi 2014), (Oldewurtel et al. 2010) and (Oldewurtel et al. 2014).

Despite extensive literature on the topic, implementation of MPC strategies in occupied buildings are quite rare. In (Sturzenegger et al. 2015), it is noted that at present, the initial investment may be too high to justify implementation when compared to the operational savings which would be achieved. Examples in which experimentation has been carried out in buildings that are occupied can be found in (Ferkl et al. 2010), (Široký et al. 2011) and (Sturzenegger et al. 2015). To compete with traditional strategies, the commissioning effort both in terms of developing models of the system, and effective control configurations which can be easily set up and adjusted, must not be excessive.

2.4 Conclusion

This chapter provides a background to some of the technologies used for control of building energy systems. Current strategies that are commonly used were first covered, noting how the addition of ever-more data can facilitate opportunities for an improved performance to be achieved, if suitable techniques are used, by incorporating different controllers and allowing for more of the complex dynamics associated with a building to be included. A brief overview was provided of model-free methods which are used for intelligent building control, including fuzzy logic control, artificial neural net-

works and agent-based methods. Optimal model-based methods were then introduced, focussing specifically on model predictive control (MPC), which can use predictions of the future needs of the building to make more informed control decisions in the present.

A description of how constrained optimisation problems are solved and how MPC is formulated was first given, followed by an outline of how MPC is applied to the problem of building energy control. The difficulty of including enough complexity to achieve a performance improvement over traditional approaches, while maintaining enough simplicity to be comparable in terms of commissioning effort and computational requirement is emphasised.

References

- Abbas, M. A., Naughton, R. & Eklund, J. M. (2011), System identification and predictive control of a building heating system with multiple boilers, *in* ‘Canadian Conference on Electrical and Computer Engineering’, pp. 001483–001488.
- Anderson, C. W., Hittle, D., Matt, K. & Young, P. (2004), Robust Reinforcement Learning for Heating, Ventilation, and Air Conditioning Control of Buildings, *in* ‘Handbook of Learning and Approximate Dynamic Programming’, Wiley-IEEE Press, chapter 20, pp. 517–534.
- Ansi/Ashrae (2013), ‘ANSI/ASHRAE 55:2013 Thermal Environmental Conditions for Human Occupancy’, *Ashrae* p. 30.
- Baetens, R., Jelle, B. P. & Gustavsen, A. (2010), ‘Properties, requirements and possibilities of smart windows for dynamic daylight and solar energy control in buildings: A state-of-the-art review’, *Solar Energy Materials and Solar Cells* **94**(2), 87–105.
- Camacho, E. & Alba, C. (2007), *Model predictive control*, second edn, Springer-Verlag, London.
- Castilla, M., Álvarez, J., Berenguel, M., Rodríguez, F., Guzmán, J. & Pérez, M. (2011), ‘A comparison of thermal comfort predictive control strategies’, *Energy and Buildings* **43**(10), 2737–2746.
- Chen, T. (2002), ‘Application of adaptive predictive control to a floor heating system with a large thermal lag’, *Energy and Buildings* **34**, 45–51.
- CIBSE (Chartered Institution Building Services Engineers) (2006), *Environment Design: CIBSE Guide A*, number CIBSE Publications.
- Dalamagkidis, K., Kolokotsa, D., Kalaitzakis, K. & Stavrakakis, G. S. (2007), ‘Reinforcement learning for energy conservation and comfort in buildings’, *Building and Environment* **42**(7), 2686–2698.
- Dounis, A. I. & Caraiscos, C. (2009), ‘Advanced control systems engineering for en-

- ergy and comfort management in a building environment-A review', *Renewable and Sustainable Energy Reviews* **13**(6-7), 1246–1261.
- EN-ISO-7730 (2005), ISO 7730:2005, Technical Report 015.
- Ferkl, L., Široký, J. & Přívara, S. (2010), Model predictive control of buildings: The efficient way of heating, in 'Proceedings of the IEEE International Conference on Control Applications', Yokohama, pp. 1922–1926.
- Floudas, C. A. (1995), *Nonlinear and Mixed-Integer Optimization: Fundamentals and Applications*, illustrate edn, Oxford University Press, USA.
- Gouda, M. M., Danaher, S. & Underwood, C. P. (2001), 'Thermal comfort based fuzzy logic controller', *Building Service Engineering Research and Technology* **22**(4), 237–253.
- Guinard, A., McGibney, A. & Pesch, D. (2009), 'A wireless sensor network design tool to support building energy management', *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings* pp. 25–30.
- Hazyuk, I., Ghiaus, C. & Penhouet, D. (2012a), 'Optimal temperature control of intermittently heated buildings using Model Predictive Control: Part I - Building modeling', *Building and Environment* **51**, 379–387.
- Hazyuk, I., Ghiaus, C. & Penhouet, D. (2012b), 'Optimal temperature control of intermittently heated buildings using Model Predictive Control: Part II - Control algorithm', *Building and Environment* **51**, 388–394.
- Houwing, M., Negenborn, R. R. & De Schutter, B. (2010), 'Demand Response With Micro-CHP Systems', *Proceedings of the IEEE* **99**(1).
- Humphreys, M. A. & Fergus Nicol, J. (2002), 'The validity of ISO-PMV for predicting comfort votes in every-day thermal environments', *Energy and Buildings* **34**(6), 667–684.
- IPCC (2014), 'Climate Change 2014 Synthesis Report', *Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* pp. 1–151.
- Kanarachos, a. & Geramanis, K. (1998), 'Multivariable control of single zone hydronic heating systems with neural networks', *Energy Conversion and Management* **39**(13), 1317–1336.
- Karatasou, S., Santamouris, M. & Geros, V. (2006), 'Modeling and predicting build-

- ing's energy use with artificial neural networks: Methods and results', *Energy and Buildings* **38**(8), 949–958.
- Karlsson, H. & Hagentoft, C.-E. (2011), 'Application of model based predictive control for water-based floor heating in low energy residential buildings', *Building and Environment* **46**(3), 556–569.
- Karmarkar, N. (1984), 'A New Polynomial-Time Algorithm for Linear Programming', *Combinatorica* **4**(April), 373–395.
- Liang, J. & Du, R. (2005), 'Thermal comfort control based on neural network for HVAC application', *Proceedings of 2005 IEEE Conference on Control Applications, 2005. CCA 2005*. pp. 819–824.
- Liao, Z. & Dexter, A. (2010), 'An Inferential Model-Based Predictive Control Scheme for Optimizing the Operation of Boilers in Building Space-Heating Systems', *IEEE Transactions on Control Systems Technology* **18**(5), 1092–1102.
- López, A., Sánchez, L., Doctor, F., Hagrass, H. & Callaghan, V. (2004), An evolutionary algorithm for the off-line data driven generation of fuzzy controllers for intelligent buildings, in 'Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics', Vol. 1, pp. 42–47.
- Ma, Y., Anderson, G. & Borrelli, F. (2011), A Distributed Predictive Control Approach to Building Temperature Regulation, in 'American Control Conference', pp. 2089–2094.
- Ma, Y., Kelman, A., Daly, A. & Borelli, F. (2012), 'Predictive control for energy efficient buildings with thermal storage', *IEEE Control Systems Magazine* **32**(February), 44–64.
- Ma, Y., Matusko, J. & Borrelli, F. (2014), 'Stochastic Model Predictive Control for Building HVAC Systems: Complexity and Conservatism', *IEEE Transactions on Control Systems Technology* **23**(1), 1–1.
- Maasoumy, M., Razmara, M., Shahbakhti, M. & Vincentelli, a. S. (2014), 'Handling model uncertainty in model predictive control for energy efficient buildings', *Energy and Buildings* **77**, 377–392.
- Maciejowski, J. (2002), *Predictive Control: With Constraints*, Prentice-Hall, UK.
- Morari, M. & Lee, J. (1999), 'Model predictive control: past, present and future', *Computers & Chemical Engineering* **23**(4-5), 667–682.

- Morel, N., Bauer, M., El-Khoury, M. & Krauss, J. (2001), 'Neurobat, a Predictive and Adaptive Heating Control System Using Artificial Neural Networks', *International Journal of Solar Energy* **21**(2-3), 161–201.
- Moroşan, P.-D., Bourdais, R., Dumur, D. & Buisson, J. (2010), 'Building temperature regulation using a distributed model predictive control', *Energy and Buildings* **42**(9), 1445–1452.
- Nocedal, J. & Wright, S. (1999), *Numerical Optimization*, 2 edn, Springer-Verlag.
- Oldewurtel, F., Jones, C. N., Parisio, A. & Morari, M. (2014), 'Stochastic Model Predictive Control for Building Climate Control', *IEEE Transactions on Control Systems Technology* **22**(3), 1198–1205.
- Oldewurtel, F., Parisio, A., Jones, C., Morari, M., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B. & Wirth, K. (2010), Energy efficient building climate control using Stochastic Model Predictive Control and weather predictions, in 'American Control Conference', pp. 5100–5105.
- Oldewurtel, F., Parisio, A., Jones, C. N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B. & Morari, M. (2012), 'Use of model predictive control and weather forecasts for energy efficient building climate control', *Energy and Buildings* **45**, 15–27.
- Oldewurtel, F., Sturzenegger, D. & Morari, M. (2013), 'Importance of occupancy information for building climate control', *Applied Energy* **101**, 521–532.
- Ole Fanger, P. & Toftum, J. (2002), Extension of the PMV model to non-air-conditioned buildings in warm climates, in 'Energy and Buildings', Vol. 34, pp. 533–536.
- Paris, B., Eynard, J., Grieu, S., Talbert, T. & Polit, M. (2010), 'Heating control schemes for energy management in buildings', *Energy and Buildings* **42**(10), 1908–1917.
- Peeters, L., Van der Veken, J., Hens, H., Helsen, L. & D'haeseleer, W. (2008), 'Control of heating systems in residential buildings: Current practice', *Energy and Buildings* **40**, 1446–1455.
- Pérez-Lombard, L., Ortiz, J., Coronel, J. F. & Maestre, I. R. (2011), 'A review of HVAC systems requirements in building energy regulations', *Energy and Buildings* **43**(2-3), 255–268.
- Sadeghi, S. A. (2014), Stochastic Model Predictive Control of Mixed-mode Buildings

- Based on Probabilistic Interactions of Occupants With Window Blinds, in 'International High Performance Buildings Conference'.
- Sadineni, S. B., Madala, S. & Boehm, R. F. (2011), 'Passive building energy savings: A review of building envelope components', *Renewable and Sustainable Energy Reviews* **15**(8), 3617–3631.
- Shaikh, P. H., Nor, N. B. M., Nallagownden, P., Elamvazuthi, I. & Ibrahim, T. (2014), 'A review on optimized control systems for building energy and comfort management of smart sustainable buildings', *Renewable and Sustainable Energy Reviews* **34**, 409–429.
- Sherman, M. (1985), 'A simplified model of thermal comfort', *Energy and Buildings* **8**(1), 37–50.
- Široký, J., Oldewurtel, F., Cigler, J. & Prívara, S. (2011), 'Experimental analysis of model predictive control for an energy efficient building heating system', *Applied Energy* **88**(9), 3079–3087.
- Sturzenegger, D., Gyalistras, D., Morari, M. & Smith, R. S. (2015), 'Model Predictive Climate Control of a Swiss Office Building : Implementation , Results , and Cost – Benefit Analysis', *IEEE Transactions on Control Systems Technology* pp. 1–12.
- Thomas, B., Soleimani, M. & Fahlén, P. (2005), 'Feed-forward in temperature control of buildings', *Energy and Buildings* **37**(7), 755–761.
- Wächter, A. & Biegler, L. (2006), 'On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming', *Mathematical programming* **106**(1), 25–57.
- Yang, R. & Wang, L. (2013a), 'Development of multi-agent system for building energy and comfort management based on occupant behaviors', *Energy and Buildings* **56**, 1–7.
- Yang, R. & Wang, L. (2013b), 'Multi-zone building energy management using intelligent control and optimization', *Sustainable Cities and Society* **6**(1), 16–21.
- Zhao, P., Suryanarayanan, S. & Simoes, M. G. (2010), 'An Energy Management System for Building Structures Using a Multi-Agent Decision-Making Control Methodology', *Industry Applications Society Annual Meeting (IAS), 2010 IEEE* pp. 1–8.

Chapter 3

Representing a Building's Thermodynamics in a Validated Simulation Platform

3.1 Introduction

This chapter introduces a methodology for developing a simulation platform, designed to represent the thermal dynamics of a building in Simulink. The purpose of the resulting simulation platform is to allow for analysis and comparison of different control strategies. Different metaheuristic search algorithms are assessed for parameter identification purposes and a disturbance estimation technique is introduced to assist in the identification process using Principal Component Analysis.

To obtain a meaningful comparison of the performance of different building control strategies (particularly in terms of energy consumption), the external conditions, disturbances and initial conditions for each experiment should be consistent (Ghiaus 2006). Due to the ever-changing external environment of a real building, this may not be trivial. A reduction in energy use may be more attributable to the sun, for example, than a particular change in the heating system control strategy. A possible option to achieve consistent experimental conditions could be to run the experiments for an extended period of time, long enough for short-term trends to be cancelled out. Given the slow nature of the changing external environment (seasonal changes would have to be classified as short-term trends), such an approach is not realistic in this context.

A more practical method for control performance comparison is to design an appropri-

ate simulation platform in which all experimental conditions can be set (Hazyuk et al. 2012a). Different approaches can be used to represent a building as a large scale simulation model for this purpose. In this work, due to the linear structure, flexibility and availability of building dimensions and material specifications, an RC-network type strategy is used. The components of the network are initially estimated using physical material properties and dimensions. A novel method for training the system to better represent the building's thermodynamics is introduced using measured input and output data and spatially filtered disturbance estimates.

An appraisal of the concept is carried out using a high-order test network designed to represent the dynamics of a hypothetical building. Using such a test network allows for a more complete assessment of the identification techniques developed. The strategy is then implemented using measured data from the Nimbus Centre to develop a platform that represents its thermal dynamics. Finally, a representation of the heating system of the Nimbus Centre is established.

3.2 Modelling the Thermal Properties of Buildings

An extensive range of literature on issues associated with developing models which can capture the complex characteristics of a building exists, reviews of which can be found in (Coakley et al. 2014) and (Harish & Kumar 2016). While many different methods are possible (artificial neural networks are used, for example, in (Kumar et al. 2013, Neto & Fiorelli 2008) and (Karatasou et al. 2006)), the most common technique is to represent the structure as an RC-network (Oldewurtel et al. 2012), (Ma et al. 2012), particularly when a knowledge of the physical composition of the building is available (Široký et al. 2011). In such a circuit, each wall is represented by a system of resistances and capacitances, each room or airspace is represented by a single capacitance, while windows are modelled by resistances. In the context of this analogy, temperatures and heat flows are viewed as voltages and currents respectively. The arrangement of the network depends on the structure of the specific building being modelled, while resistance and capacitance parameters depend on building materials and dimensions.

A simple 3R2C network describing the thermal dynamics of a single wall connecting the zone to the external environment including a window and a heat source is shown in Fig. 3.1. The zone capacitance is given by c_1 , the three wall resistances are denoted r_1 , r_2 and r_3 , while the wall capacitances are denoted c_2 and c_3 . The window is represented by the resistance r_4 and the heat supply to the zone at temperature T_z is given by Q_{in} .

T_{W1} and T_{W2} are temperatures in the interior of the wall. This network could be




Figure 3.1: Equivalent 3R2C RC-network model of the thermal connection between a zone and the external environment

described by the following state-space equations:

$$\begin{bmatrix} \dot{T}_z(t) \\ \dot{T}_{W1}(t) \\ \dot{T}_{W2}(t) \end{bmatrix} = \begin{bmatrix} \frac{-1}{c_1} \left(\frac{1}{r_1} + \frac{1}{r_4} \right) & \frac{1}{c_1 r_1} & 0 \\ \frac{1}{c_2 r_1} & \frac{-1}{c_2} \left(\frac{1}{r_1} + \frac{1}{r_2} \right) & \frac{1}{c_2 r_2} \\ 0 & \frac{1}{c_3 r_2} & \frac{-1}{c_3} \left(\frac{1}{r_2} + \frac{1}{r_3} \right) \end{bmatrix} \begin{bmatrix} T_z(t) \\ T_{W1}(t) \\ T_{W2}(t) \end{bmatrix} + \begin{bmatrix} \frac{1}{c_1} \\ 0 \\ 0 \end{bmatrix} Q_{in}(t) + \begin{bmatrix} \frac{1}{c_1 r_4} \\ 0 \\ \frac{1}{c_3 r_3} \end{bmatrix} T_{External}(t) \quad (3.1)$$

$$T_z(t) = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} T_z(t) \\ T_{W1}(t) \\ T_{W2}(t) \end{bmatrix} \quad (3.2)$$

In Fig. 3.2, the network is extended to a two zone building with one partition wall and windows in each of the external walls (ignoring floors and ceilings), where the zone temperatures are denoted T_{z1} and T_{z2} , and the radiator heat inputs are denoted Q_{in1} and Q_{in2} . Such a network can also then be represented in a standard linear state-space format as follows:

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) + \mathbf{E}\mathbf{d}(t) \quad (3.3)$$

$$\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t), \quad (3.4)$$

where \mathbf{x} is a vector of states, \mathbf{y} is a vector of zone temperatures, \mathbf{u} is a vector of heat inputs, \mathbf{d} is a vector of disturbances (including external temperature and heat gains) and \mathbf{A} , \mathbf{B} and \mathbf{E} are matrices made up of combinations of resistances and capacitances. \mathbf{C} is the observation matrix.

Data-based approaches are used to train model parameters in (Privara et al. 2011, Yang




Figure 3.2: Equivalent RC Network model of two zones with one internal wall, six external walls and six windows

& Wang 2012) and (O'Dwyer et al. 2016), while in (Zhao & Magoulès 2012) and (Pérez-Lombard et al. 2008), models are developed to accurately predict energy consumption. Aside from the thermal properties of the building structure, modelling of the impact of occupants on the internal building temperatures is considered in (Oldewurtel et al. 2013) and (Page et al. 2008). A review of computer packages which have been designed for the purpose of building energy simulation (both in terms of heat transfer through the structure and how the external environment impacts the internal temperatures) is provided in (Crawley et al. 2008).

3.3 Representative RC-Network for Analysis of Modelling Methodologies

To first investigate the different methodologies for training RC-networks to fit data, a network consisting of 231 resistors and 172 capacitors was set up to simulate the dynamics of a notional 30 zone building. A 3R-2C arrangement was used to represent all the walls (both internally and externally) in the building, while windows and zones were modelled as single resistances and capacitances respectively. The hypothetical building was configured so that each zone is a different size, but rectangular and enclosed by four walls. The input to each zone is taken as the heat supplied by its radiator, while the outputs are the zone temperatures. The external temperature is used as a measured disturbance and treated as an additional input.

The arrangement of the zones of the notional building, overlaid by the RC-network

representing it, can be seen in Fig. 3.3.

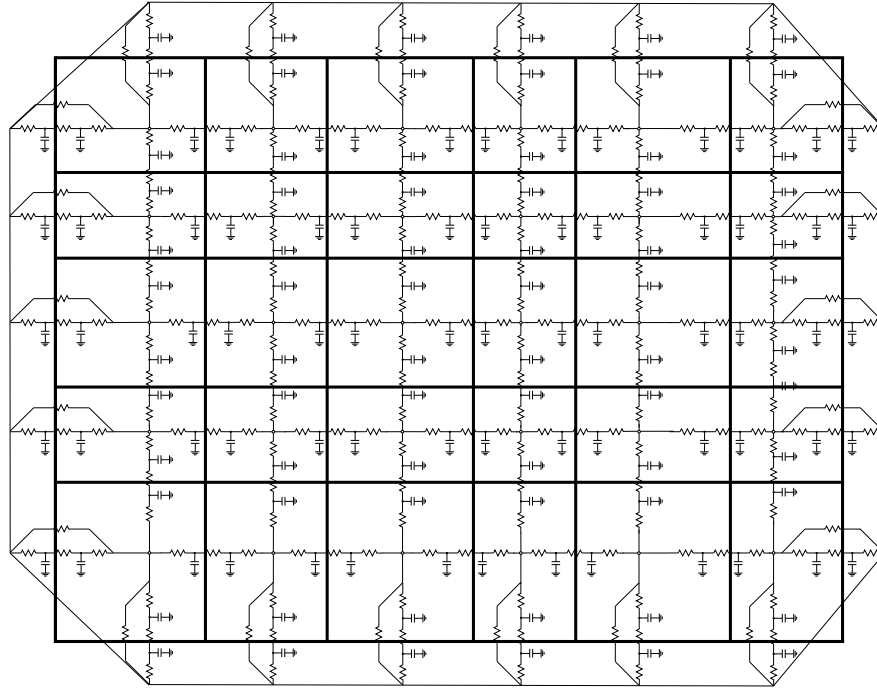


Figure 3.3: RC-network used for testing identification strategies

A model of the network can be represented in a linear state-space structure as:

$$\dot{\mathbf{x}}(t) = A(\boldsymbol{\theta})\mathbf{x}(t) + B(\boldsymbol{\theta})\mathbf{u}(t) \quad (3.5)$$

$$\mathbf{y}(t) = C\mathbf{x}(t) \quad (3.6)$$

where, at time t , $\mathbf{y}(t) \in \mathbb{R}^{30}$ is the vector of temperature measurements (one for each zone), $\mathbf{x}(t) \in \mathbb{R}^{172}$ is the vector of system states and $\mathbf{u}(t) \in \mathbb{R}^{31}$ represents the inputs and the measured disturbance (in this case, the external temperature, $T_{amb}(t)$), while A , B and C are the RC-network coefficient matrices. The set of network parameters (the R and C values) are given by the vector $\boldsymbol{\theta}$.

To represent an initial (possibly inaccurate) physics-based estimate of the system, another vector of parameters given by $\boldsymbol{\theta}_0$ was generated, deliberately chosen to be different than the actual set of network parameters. The network $G(s, \boldsymbol{\theta}_0)$ was then taken to be the initial, approximate estimate of the network so that:

$$G(s, \boldsymbol{\theta}_0) \neq G(s, \boldsymbol{\theta}) \quad (3.7)$$

$$G(s, \boldsymbol{\theta}) = C(s\mathbf{I} - A(\boldsymbol{\theta}))^{-1}B(\boldsymbol{\theta}), \quad (3.8)$$

where s is the Laplace operator and $\boldsymbol{\theta}$ is the set of correct parameter values so that

$G(s, \theta)$ can be regarded as the target system.

The objective of the identification process was to adjust the parameters of this initial estimate, from $\theta_0 \rightarrow \tilde{\theta}$, so as to obtain a better approximation of the building model, $G(s, \tilde{\theta}) \approx G(s, \theta)$, using only input and output data with realistic excitation. The test network allows for the performance to be characterised (as opposed to using the real plant where details of the target dynamics are unknown), thus enabling an appropriate selection of tuning parameters and fitness functions.

To begin with, no unmeasured disturbances were applied to the system - all inputs and outputs were assumed to be known. In a real building however, such an assumption would most likely be invalid. Heat gains from building occupants, equipment and solar radiation, as well as other factors such as wind and occupant interaction (open windows or doors for example) will tend to have a large corrupting impact on any identified parameters. As such, the identification approach is updated to incorporate an estimate of these factors at a later point in the chapter.

Input data was generated and applied to the test network, $G(s, \theta)$, to replicate a realistic heat supply that would occur in normal circumstances using a weather-compensated control strategy, without functional tests. The heat supply profile followed a schedule in which the system was switched on for 4 hours in the morning, 2 hours in the afternoon and 30 minutes in the evening. Measured external temperature data was also included as an input to obtain the output data-set over a training period of 5 days.

In Fig. 3.4, the convective input for Zone 1 can be seen, while Fig. 3.5 shows the external temperature.

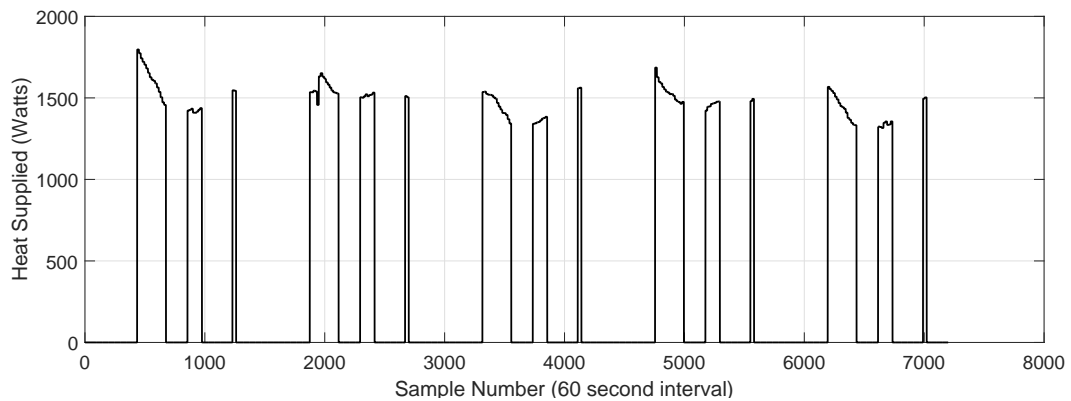


Figure 3.4: Convective heat supply profile applied to Zone 1 of test network used for training process

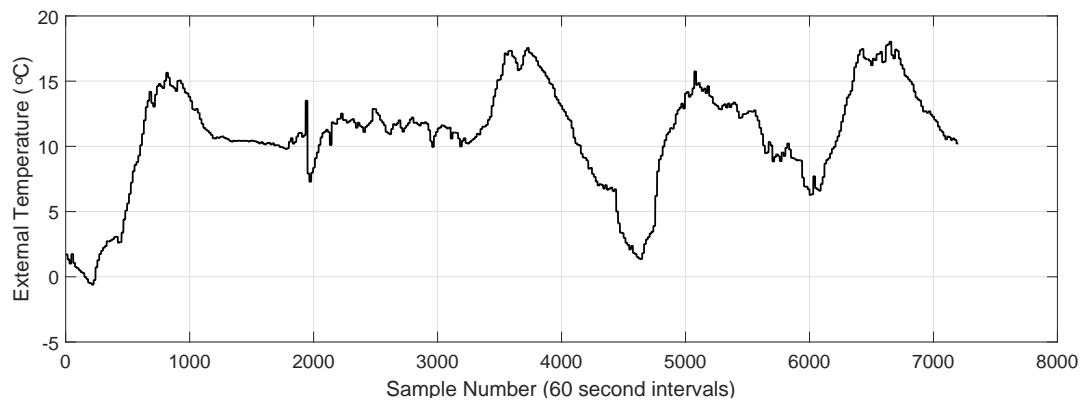


Figure 3.5: External temperature profile applied to test network used for training process

3.4 Identification of Model Parameters from Measured Data

3.4.1 Outlining the Identification Problem - Complexity and Identifiability Considerations

3.4.1.1 Excitation of the System

The quality of data used to train the network is of critical importance. Any data driven modelling approach used to ascertain the thermal dynamics of a building is likely to be limited by the range of possible input excitation. A building's heating system will typically not be capable of pushing the temperature of zones in the building outside of a relatively narrow operating region. Additionally, higher frequency inputs may not be permissible.

Finding the relationship between the external temperature and the measured outputs is more restrictive again, with the excitation in this case being entirely dependent on the weather. A final issue is the paucity of measurements, relative to the number of unmeasured states. Such an under-determined problem will tend to produce a non-unique solution (Coakley et al. 2014).

Structural and data-based identifiability definitions and tests are outlined in (Van Doren et al. 2008, 2009). In (Agbi & Krogh 2014) and (Agbi et al. 2012), the required conditions are investigated in the area of building and zone modelling. In (Agbi et al. 2012) it is noted that the required experimental conditions to allow for identifiable models fall outside normal building operation.

3.4.1.2 Complexity of the Parameter Selection Process

A full scale building will generally consist of many zones connected by an intricate system of walls, floors and ceilings. A network representing such a system will then be composed of a large number of parameters. Adjustment of these parameters to better replicate the dynamics observed in measured data results in an identification problem of high dimensionality.

The structure of the model must also be considered here. Though an RC-network can be expressed in a linear state-space format, a nonlinear relationship exists between the state-space parameters and the individual R and C components. Taking the one wall RC-network of (3.1)-(3.2) as an example, the identification problem could first be viewed as a linear grey-box problem, where the model matrices are parameterized by θ as follows:

$$\dot{\mathbf{x}}(t) = A(\theta) \mathbf{x}(t) + B(\theta) u(t) + E(\theta) d(t) \quad (3.9)$$

$$y(t) = C\mathbf{x}(t) \quad (3.10)$$

$$A = \begin{bmatrix} \theta_1 & \theta_2 & 0 \\ \theta_3 & \theta_4 & \theta_5 \\ 0 & \theta_6 & \theta_7 \end{bmatrix} \quad (3.11)$$

$$B = \begin{bmatrix} \theta_8 \\ 0 \\ 0 \end{bmatrix} \quad (3.12)$$

$$E = \begin{bmatrix} \theta_9 \\ 0 \\ \theta_{10} \end{bmatrix} \quad (3.13)$$

$$C = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \quad (3.14)$$

In this form, the RC-network structure is lost, along with the link to the physical reality of heat flows and temperature differences. Notably, thermodynamic passivity is no longer certain. To maintain the integrity of the network, the identification process must seek to directly identify resistance and capacitance values rather than directly estimating the parameters of the state-space system. The model matrices, in terms of the adjustable parameter vector θ would then have the form:

$$A = \begin{bmatrix} \frac{-1}{c_1} \left(\frac{1}{r_1} + \frac{1}{r_4} \right) & \frac{1}{c_1 r_1} & 0 \\ \frac{1}{c_2 r_1} & \frac{-1}{c_2} \left(\frac{1}{r_1} + \frac{1}{r_2} \right) & \frac{1}{c_2 r_2} \\ 0 & \frac{1}{c_3 r_2} & \frac{-1}{c_3} \left(\frac{1}{r_2} + \frac{1}{r_3} \right) \end{bmatrix} \quad (3.15)$$

$$B = \begin{bmatrix} \frac{1}{c_1} \\ 0 \\ 0 \end{bmatrix} \quad (3.16)$$

$$E = \begin{bmatrix} \frac{1}{c_1 r_4} \\ 0 \\ \frac{1}{c_3 r_3} \end{bmatrix} \quad (3.17)$$

$$C = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \quad (3.18)$$

$$\theta = \begin{bmatrix} c_1 & c_2 & c_3 & r_1 & r_2 & r_3 & r_4 & r_5 \end{bmatrix}^T \quad (3.19)$$

This identification process is nonlinear. For a large building with many zones, there could potentially be thousands of parameters to be identified. Furthermore, as these parameters should preferably correspond to physically sensible resistance and capacitance values, the search space should be in some way constrained to favour such results. Negative values for example, clearly would not make sense in the context of passive heat flow.

To incorporate this kind of flexibility in a high dimensional nonlinear problem, metaheuristic optimisation approaches were deemed to be most appropriate for training the RC-network.

3.4.2 Metaheuristic Search Algorithms for Complex Optimisation Problems

Deterministic optimisation algorithms, such as those underpinning the MPC formulations and system-identification approaches introduced in later chapters, use knowledge of the objective function gradients, and possibly second-derivatives, to determine a globally optimal solution (details are provided in Section 2.3.1). In certain applications, for reasons of complexity and structure, such knowledge may not be practically attainable. In contrast to the deterministic approaches, metaheuristic algorithms only require that the objective can be evaluated within a designated search-space and typically do not require the function's derivatives. Broadly speaking, a metaheuristic

search algorithm seeks to find a sufficiently close solution to an optimisation problem in a non-deterministic fashion.

Often taking inspiration from the behaviour of complex systems found in nature, different algorithms iteratively guide the search process using simple assumptions made about the system. A set of termination criteria are specified to designate the location at which a candidate solution is considered sufficiently close to optimal. Four common approaches based on different frameworks are briefly outlined here.

Chemotaxis

The chemotaxis algorithm, as outlined in (Willis et al. 1991) as a method of selecting weights for artificial neural networks, seeks to iteratively move from the initial D-dimensional parameter vector estimate, to candidate solutions of progressively lower cost, until some termination criterion is satisfied. The approach could be thought of as a single agent following an ever-decreasing path (found by trial and error as it randomly explores its local neighbourhood) until a solution is deemed to be sufficiently close to optimal.

For a non-convex cost function, convergence to a local minimum is possible as no mechanism exists for moving to a position in the search-space of higher cost. In Fig. 3.6 a contour plot of a hypothetical nonlinear objective function is shown to highlight the issue. The position of the global minimum is desired. In Fig. 3.6(a), by randomly testing the local neighbourhood, an ever-decreasing pathway of candidate solutions is shown, leading towards the global minimum of the objective. In Fig. 3.6(b) however, an equivalently valid pathway in which the cost decreases with each move is shown to lead towards a local minimum which is not globally optimum.




Figure 3.6: Chemotaxis search on a nonlinear surface

Simulated Annealing

Simulated Annealing (SA) is an alternative algorithm in which higher-cost candidate solutions can be accepted. Developed in (Kirkpatrick et al. 1983), SA is an iterative process in which the probability of moving from one candidate solution to another is higher if the movement results in a reduction in cost. As the search progresses, the probability of accepting higher-cost candidate solutions diminishes. The algorithm is based on the process of annealing, used in metallurgy in which a material is heated to a high temperature, at which the physical properties of the material are altered. The material is then slowly allowed to cool, resulting in an improved level of ductility in the material. The change in energy experienced by the material in the process is analogous to the change in the probability of candidate solution acceptance in the algorithm.

Genetic Algorithms

Using concepts originally proposed in (Turing 1950) and later developed in the 1960s and 1970s, genetic algorithms belong to the more general class of evolutionary algorithms, seeking to imitate some of the characteristics of an evolving gene-pool. The search progresses by evaluating a set of candidate solutions, stochastically selecting those of greatest fitness and introducing random mutations to modify this remaining candidate set. This selection and mutation process is repeated iteratively (each iteration is then considered a new *generation*) until termination criteria are satisfied. The mutation property of the search allows for higher-cost solutions to be selected, which may ultimately lead to a lower final solution. The likelihood of convergence at local minima is then reduced.

Particle Swarm Optimisation

Introduced in (Kennedy & Eberhart 1995), particle swarm optimisation (PSO) is another gradient-free optimisation algorithm in which a population of particles, or agents, move through a multi-dimensional space. Many variations on the original algorithm exist with different convergence properties (Valle 2008, Kulkarni & Venayagamoorthy 2011). Algorithms based on swarm intelligence, such as PSO and Ant Colony Optimisation (ACO), permit individual agents to exploit information of a global population of agents to move towards globally lower-cost regions of the search-space, even if such a movement produces a higher local cost. PSO seeks to emulate the dynamics of a flock of birds or a shoal of fish, in which the collective knowledge of the swarm is available to (and can be exploited by) each individual member. In a standard PSO algorithm, the velocity of each agent in the population is determined at every iteration, subject to the previous velocity of the agent, the previous local-best (minimum cost) position and the

global-best previous position of the entire population (Gazi & Passino 2011).

In Fig. 3.7(a), a population of 10 agents is shown on the same contour plot as in Fig. 3.6. The agents explore the search-space, over time being drawn towards the lowest cost candidate. If an adequate level of exploration of the search-space is carried out, the swarm should converge towards the global minimum, as shown in Fig. 3.7(b).

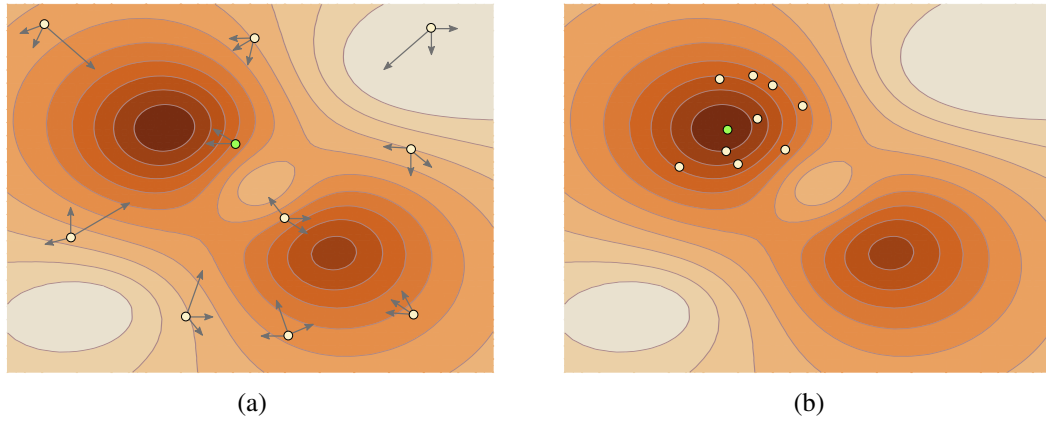


Figure 3.7: Particle Swarm Optimisation with 10 agents

In the no free lunch theorems of (Wolpert & Macready 1997), it is shown that if an algorithm outperforms another over one class of problems, the opposite must be true for all other classes. As such, to select an algorithm, the cost function (referred to in this chapter as the fitness function to avoid confusion with costs and objectives of later control-oriented chapters) to be minimised must be considered. For this problem, the chemotaxis search algorithm is first analysed.

3.5 Chemotaxis Search for Identifying Model Parameters

3.5.1 The Chemotaxis Algorithm

Where the fitness function is given as $f(\theta)$, the pseudo-code for the chemotaxis search algorithm is shown in Algorithm 1.

The scalar parameter γ_{ch} determines the overall aggression of the search. The main diagonal of the square matrix $\mu \in \mathbb{R}^{N \times N}$, is comprised of scalars, each of which determines how aggressively the search is carried out in each of the D dimensions of

Algorithm 1 Chemotaxis Algorithm

```

1:  $\tilde{\theta} = \theta_0$ 
2: for  $II = 1..T$  do ▷ Number of iterations
3:    $\phi \sim \mathcal{N}(0, 1)$ ;
4:    $\hat{\theta} = \tilde{\theta} + \gamma_{ch}\mu\phi$ ;
5:   if  $f(\hat{\theta}) < f(\tilde{\theta})$  then
6:      $\tilde{\theta} = \hat{\theta}$ ; ▷ Update candidate solution
7:   end if
8: end for

```

θ , relative to the other dimensions. If certain parameters of the initial physics-based network estimate are likely to be known with greater accuracy than others, such knowledge can be incorporated by correctly weighting μ . Otherwise, a sensible approach is to set the main diagonal equal to θ_0 , so as to account for the difference in scale of different components.

Aside from μ and γ_{ch} , the behaviour of the algorithm can be adjusted by changing the termination criteria and the fitness function used. If a root-mean-square (rms) error is used as the fitness function for example, the low frequency and steady-state gain will be emphasised, whereas using a k-step ahead prediction error will result in more accurate higher frequencies, possibly at the expense of steady-state gain.

The dynamics of the building are slow and the objective of any control strategy will be to maintain comfort over long periods of time. As such, the steady-state gain accuracy of the simulation model is seen here as being more important than accuracy in the higher end of the frequency spectrum. Taking this into account, the cumulative squared output-error across all zones for the full training period was used as the fitness function as follows:

$$f(\theta) = \sum_{j=1}^M \sum_{i=1}^N (y_i(j) - \hat{y}_i(j))^2, \quad (3.20)$$

where N is the number of zones, M is the number of samples in the batch of training data, $y_i(j)$ is the measurement of the i^{th} zone temperature at the j^{th} sample and $\hat{y}_i(k)$ represents the model output of the same zone at the j^{th} sample when the model is parameterized by θ .

Acknowledging the outlined identifiability limitations, to reduce the dimension of the problem, only parameters directly linked to a measured state were adjusted - all other parameters were fixed at the initial physics-based estimate. In terms of the RC-network, as states associated with the internal wall temperatures cannot be measured

this amounts to only the zone capacitances and internal wall resistances being adjusted. As the effect of other components are filtered through these before measurement, it is more difficult to determine the impact they might have on the output, particularly given an under-excited set of data.

Using the input and output data described in Section 3.3, the chemotaxis algorithm performance was investigated. For different choices of γ_{ch} , the parameters of the estimated network were adjusted until the algorithm converged.

The main diagonal of the tuning parameter μ was set equal to θ_0 . In Fig. 3.8, the cost function is plotted for each iterate as the search progresses, using different values of γ_{ch} (values of 0.001, 0.01, 0.1 and a fourth set in which the value is linearly reduced from 0.01 to 0.001 at each iteration are shown).




Figure 3.8: Fitness-function reduction as algorithm progresses for different levels of search aggression

The most conservative search ($\gamma_{ch} = 0.001$) converges quite quickly to a local minimum. Though initially achieving a significant cost improvement, the most aggressive search did not provide a final cost as low as the search with $\gamma_{ch} = 0.01$.

3.5.2 Evaluation of solutions

To appropriately evaluate the quality of the solution, different metrics are possible. First, to ensure the network does not over-fit the training data-set, the output error is viewed when applying a new unseen set of inputs to the target system $G(s, \theta)$, the initial system $G(s, \theta_0)$ and the identified system $G(s, \tilde{\theta})$. This validation data-set simulates a period of 10 days.

If the data is not rich enough to provide a unique solution, a low cost function value may conceal inconsistencies in the solution. Though the output of a system may be similar to that of the target system, the component values may be very different, no longer corresponding to physically realistic thermal resistance and capacitance values. To examine this, a comparison is made between the original estimated parameter values, the target values and the adjusted network values as an additional evaluation metric.

In Table 3.1 the average daily accumulated absolute output errors for the both the training $Train_{err}$ and validation Val_{err} data-sets are shown, given in units of $^{\circ}C.hr$ (an average zone output deviation of $1^{\circ}C$ from the target for 1 hour would contribute $1^{\circ}C.hr$ to the total for the day).

Additionally, the mean relative error of the network component values (compared to those of the target network and given as a percentage) are shown, denoted by $Param_{err}$. It was found that of these parameters, the capacitance values relating to the thermal mass of the zones appear to have the largest impact on the performance. As such, the mean relative error of these are also shown, given by C_{err} .

Table 3.1: Model error present before and after chemotaxis adjustment using test network for different algorithm configurations

NETWORK ADJUSTMENT USING CHEMOTAXIS					
	Initial Network	$\gamma_{ch} :$ 0.001	$\gamma_{ch} :$ 0.01	$\gamma_{ch} :$ 0.1	$\gamma_{ch} :$ 0.01 \rightarrow 0.001
$Train_{err}$ ($^{\circ}C.hr$)	216.00	211.2	17.51	54.94	159.71
Val_{err} ($^{\circ}C.hr$)	276.35	270.58	19.40	58.26	203.98
$Param_{err}$ (%)	31.62	34.23	26.33	38.59	31.02
C_{err} (%)	34.48	32.02	11.91	26.69	31.65

In Fig. 3.9, the outputs of the system for the validation-set of inputs is shown for Zone 8 of the networks identified using the chemotaxis algorithm. It can be seen that the output obtained from the network, as identified with $\gamma_{ch} = 0.01$, almost exactly matches the target.

Two important conclusions can be drawn from the above results. First of all, the importance of appropriately tuning this class of search algorithm is clear. The set of results in which $\gamma_{ch} = 0.01$ outperforms the other configurations in all measures.




Figure 3.9: Validation data-set output comparison (temperature of Zone 8) of identified networks with target network

Secondly, achieving an improved output accuracy does not necessarily correspond to a more accurate set of network parameters. In the best performing set of results, the output of the identified model is almost identical to the target yet the percentage difference between the adjusted parameters in the two networks is on average approximately 26%. When using the most aggressive search, though the input/output behaviour was captured quite well for the set of inputs used for training and validation, and the most dominant components in the network (the zone airspace capacitances) were improved, the network parameters were, on average, further from the target than the original estimate. Given the large number of parameters present, the low level of excitation in the data will not result in a unique solution.

The inputs of the validation set were different to the training set (and unseen by the search algorithm). In both cases, standard operating conditions were simulated and so the operating range was quite similar. While the operating range of a real building will most likely be comparably narrow, the accuracy of the network in all possible scenarios can not be guaranteed. An improved model accuracy over the standard operating range is achieved however, in both training and validation data-sets.

As a final comparison, for Zone 1 and Zone 25, bode plots showing the frequency response of the transfer function between the input heat flow to the zone (from the radiator) and the output temperature of the zone can be seen in Fig. 3.10 - Fig. 3.13 for the identified systems. For the same zones, bode plots showing the transfer function between the external temperature and the internal temperature are also shown. The bode plots show that despite the difference in network component parameters, with appropriate configuration of tuning parameters, the dynamics of the derived system can closely match the target over a wide frequency band.




Figure 3.10: Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 1 using target system, initial system and identified systems (for different values of γ_{ch})




Figure 3.11: Frequency response comparison of transfer function between external temperature and output temperature for Zone 1 using target system, initial system and identified systems (for different values of γ_{ch})

3.6 Platform Parameter Identification in the Presence of Unmeasured Disturbances

Up to this point, as only input and output data is used in the search algorithm, the effects of unmeasured disturbances would have to be considered negligible. Such an assumption may not be valid. Without knowledge of the disturbances, the derived RC-network parameters will often be corrupted (O'Dwyer et al. 2016). Sources of disturbance in buildings include external factors such as wind and solar gain as well as internal factors such as occupancy and equipment gains. Typically, these types of disturbance will not be zero-mean and may significantly impact the model accuracy.




Figure 3.12: Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 25 using target system, initial system and identified systems (for different values of γ_{ch})




Figure 3.13: Frequency response comparison of transfer function between external temperature and output temperature for Zone 25 using target system, initial system and identified systems (for different values of γ_{ch})

In the identification procedure, the network parameters are adjusted so as to minimise the difference between output data from the model and measured output data from the plant. As a considerable proportion of this difference may be due to unmeasured disturbances as opposed to model error, improved parameter selection is only possible if the effects of inputs on the system are separated from those of disturbances.

It should be noted at this point that it is possible to measure some such disturbances - a pyranometer can provide information of the solar intensity, while occupancy sensors (passive infrared sensors for example) can monitor the movement of occupants through the building. Such sensors may not be available in a building however. Furthermore, the measurements obtained may not be useful to each of the zones - a single measure-

ment of solar intensity for example, will not take into account the different sources of shade which may exist in various zones. The presence of misleading measurements can be detrimental to the modelling process. Additionally, even if certain sources are measured accurately, the individual effect of each source on the output may be difficult to ascertain from data (functional testing is clearly not an option for environmental disturbances).

3.6.1 Generation of disturbance profiles

Using the same test RC-network, the performance of the identification procedure was assessed in the presence of unmeasured disturbances. Additional heat inputs were applied to the system to represent solar, occupancy and equipment gains, with the assumption that no measurement of these was available.

Equipment gains (the heating effect of computers, lighting etc.) were taken as constant inputs in each zone, with a step-change increase during office hours. Occupancy gains are a result of the heating effect of people inside a building. For each zone of the building a random-walk based occupancy profile was generated during office hours with the starting point for the profile chosen to be proportional to the area of the zone. Each occupant was given an assumed metabolic heat rate of 100W. A 5 day profile of the internal gain of Zone 2, used for training is plotted in Fig. 3.14.

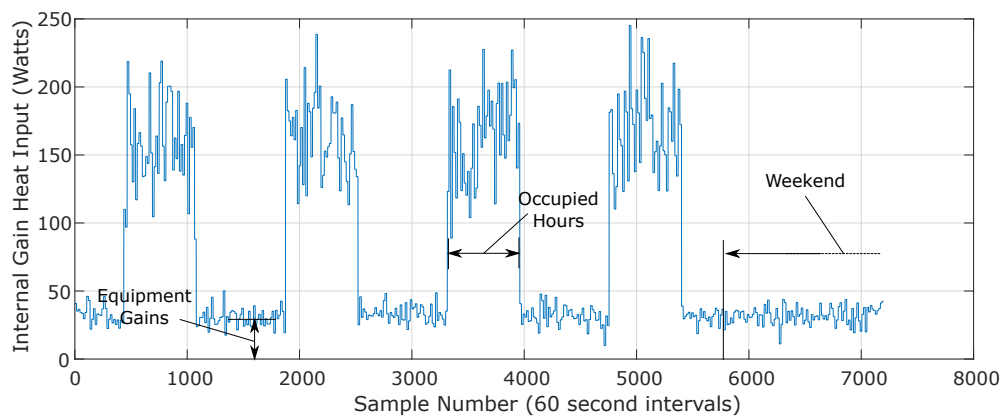


Figure 3.14: Zone 2 internal gain profile (occupancy and equipment gains) for 5 day period

The solar heat gains (being radiative in nature) were applied to the surfaces of the walls of the building, divided into direct and diffuse components. The direct component represents the beam radiation, which is received from the sun without having been scattered, whereas the diffuse component has been scattered by the atmosphere (Duffie & Beckman 2013). The intensity of the radiation was determined as a function of the

wall orientation and the angle of incidence of the solar radiation. This was found using the time of day, time of year and the latitude (Φ) of the building location in the following manner (Duffie & Beckman 2013):

Where t is the time of day on the n^{th} day of the year, the hour angle ω_h and approximate declination δ_d are expressed (in degrees) as:

$$\omega_h = 15(t - 12) \quad (3.21)$$

$$\delta_d = 23.45 \sin \left(360 \frac{284 + n}{365} \right) \quad (3.22)$$

Other angles used for calculating the angle of incidence are the surface azimuth γ_{az} , solar azimuth γ_{sz} , solar altitude α_s and zenith θ_z . These are shown in Fig. 3.15.




Figure 3.15: Relevant angles for determining intensity of solar heat gain

The relationship between these properties and the angle of incidence on a tilted surface given by θ_t and a horizontal surface given by θ_h is as follows (Duffie & Beckman 2013):

$$\begin{aligned} \cos \theta_t = & \sin \delta_d \sin \Phi \cos \beta - \sin \delta_d \cos \Phi \sin \beta \cos \gamma + \cos \delta_d \cos \Phi \cos \beta \cos \omega_h \\ & + \cos \delta_d \sin \Phi \sin \beta \cos \gamma \cos \omega_h + \cos \delta_d \sin \beta \sin \gamma \sin \omega_h \end{aligned} \quad (3.23)$$

$$\cos \theta_h = \sin \delta_d \sin \Phi + \cos \delta_d \cos \Phi \cos \omega_h \quad (3.24)$$

The intensity of the solar radiation on each vertical wall is then found by (Hazyuk et al. 2012a):

$$I_{tot} = I_{dir} \left(\frac{\cos \theta_t}{\cos \theta_h} + \frac{\rho}{2} \right) + I_{dif} \left(\frac{1 + \rho}{2} \right), \quad (3.25)$$

where ρ is the albedo of the ground surface and I_{dir} and I_{dif} denote the direct and diffuse components of the total solar irradiance. The profile for these irradiances is taken from measured pyranometer data where the direct component is assumed to be 1.5 times larger than the diffuse component. The effect of clouds was simulated by subtracting a scaled random term from the solar gain. The direct solar radiation applied to each $1m^2$ area of north, south, east and west facing wall is shown in Fig. 3.16:




Figure 3.16: Direct solar radiation profile for each orientation over 5 day period

3.6.2 Chemotaxis in the Presence of Unmeasured Disturbances

Using the same input and external temperature profile as in the training procedure of Section 3.3, the chemotaxis algorithm was once again implemented. This time, however, the unmeasured disturbances were also applied to the system. The evaluated fitness function (the cumulative squared output-error) at each iteration is shown in Fig. 3.17, with the value of $\gamma_{ch} = 0.01$, chosen, based on the results taken from the undisturbed scenario. It can be seen that the cost is far higher than in the undisturbed case.

To assess the quality of the solution, the output accuracy must once again be compared. Though the output error for the training data-set is shown to reduce in Fig. 3.17, it is unclear if the input/output relationship has been more accurately captured. To better determine the output accuracy of the solution, without the corrupting influence of the disturbances, the average daily absolute output error is found using the training and validation data-sets *without* disturbances included. These results can be seen in Table 3.2. The mean relative component error and mean relative zone capacitance error are also shown.

By these metrics, the adjusted network provides a less accurate representation of the




Figure 3.17: Fitness-function reduction as algorithm progresses - unmeasured disturbances included

Table 3.2: Model error present before and after chemotaxis adjustment using test network with unmeasured disturbances present in the system

NETWORK ADJUSTMENT		
UNMEASURED DISTURBANCE INCLUDED		
	Initial Network	Chemotaxis: $\gamma_{ch} : 0.01$
$Train_{err}$ (°C.hr)	216.00	909.16
Val_{err} (°C.hr)	276.35	1025.5
$Param_{err}$ (%)	31.62	99.41
C_{err} (%)	34.48	45.68

target than the initial estimate. This is further supported by the bode plots shown in Fig. 3.18 - Fig. 3.21. Once again, these show the frequency response of the transfer function between the input heat flow to the zone (from the radiator) and the zone temperature, as well as that of the transfer function between the external temperature and zone temperature for Zone 1 and Zone 25. To improve the solution, the unmeasured disturbance must in some way be characterised.

3.7 Methods of Disturbance Estimation for Improved System-Identification

Techniques for estimating unmeasured disturbances from data while improving model parameter accuracy are introduced here. By estimating these disturbances, it is possible




Figure 3.18: Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 1 using target system, initial system and identified systems (with disturbances included)




Figure 3.19: Frequency response comparison of transfer function between external temperature and output temperature for Zone 1 using target system, initial system and chemotaxis-identified systems (with disturbances included)

to derive a less biased simulation platform (O'Dwyer et al. 2016).

3.7.1 Kalman Filtered Disturbance Estimation

Estimated disturbance states are initially generated using an augmented discrete-time Kalman filter. The initial estimate of the RC network (i.e. the model parameterized by θ_0) is used as the underlying Kalman filter model. It is assumed that these disturbance states directly affect the measured states (which typically will correspond to the zone temperatures). The building model, with m states, n zones and n disturbance estimates (one for each zone), can be expressed as:




Figure 3.20: Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 25 using target system, initial system and chemotaxis-identified systems (with disturbances included)




Figure 3.21: Frequency response comparison of transfer function between external temperature and output temperature for Zone 25 using target system, initial system and chemotaxis-identified systems (with disturbances included)

$$\begin{bmatrix} \hat{\mathbf{x}}(k+1) \\ \hat{\mathbf{x}}_d(k+1) \end{bmatrix} = \begin{bmatrix} A(\theta_q) & C^T \\ \mathbf{0}_{n \times m} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}}(k) \\ \hat{\mathbf{x}}_d(k) \end{bmatrix} + \begin{bmatrix} B(\theta_q) \\ \mathbf{0}_{n \times 2} \end{bmatrix} \mathbf{u}(k) + K_1(\mathbf{y}(k) - \hat{\mathbf{y}}(k)) \quad (3.26)$$

$$\hat{\mathbf{y}}(k) = \begin{bmatrix} C & \mathbf{0}_{n \times n} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}}(k) \\ \hat{\mathbf{x}}_d(k) \end{bmatrix}, \quad (3.27)$$

where $\hat{\mathbf{x}}_d \in \mathbb{R}^n$ is a sequence of states representing the disturbance estimates for each zone, A , B and C are the parameter matrices discretized by zero-order hold and $K \in \mathbb{R}^{(n+m) \times n}$ is the Kalman filter gain. To tune the Kalman filter, it is assumed that the model error covariance is a diagonal matrix, in which each of the non-zero elements is 10 times larger than those of the diagonal sensor error covariance matrix.

These temporally-filtered disturbance state estimates will be both a result of modelling error and unmeasured disturbance (as both impact the output error). To improve the identification process by incorporating these disturbance estimates in the formulation, the modelling error portion must first be suppressed.

3.7.2 Spatial Filtering using Principal Component Analysis

In many cases, a single disturbance source will affect several zones at once - the solar gain for example will simultaneously impact all zones of similar orientation, likewise, occupancy hours may be the same in many zones. The individual disturbance estimates will then be in some way correlated. The dominant underlying patterns common to multiple disturbance states are extracted by Principal Component Analysis (PCA) and taken to represent the overall building disturbances.

Consider, for illustrative purposes, a data-set of 10 correlated variables, each with 500 samples, given by $X = [\mathbf{X}_1 \cdots \mathbf{X}_{10}]$. For this example, each of the variables is made up of two parts. The first is common to all variables, while the second is unique to each individual variable. The variables in this case are defined (arbitrarily) as follows:

$$X_i(k) = \sin(0.02\pi k) + \phi_{PCA_i} \cos(\omega_{PCA_i} k + \varphi_{PCA_i}) + \gamma_{PCA_i}(k) \quad (3.28)$$

$$\forall k = 1, 2, \dots, 500,$$

where ω_{PCA} , φ_{PCA} and ϕ_{PCA} are drawn for each variable from a uniform random distribution, while γ_{PCA} is additional Gaussian-driven noise. Using PCA, it is possible to isolate the constituent common to all variables without a priori knowledge of what it is, only that it exists.

The unfiltered variables are shown in Fig. 3.22.

Decomposing the sample covariance of the data into a set of eigenvectors and eigenvalues, the principal components can be found using the eigenvectors corresponding to the largest eigenvalues. Often, a large proportion of the variance of the data can be described by a number principal components far lower than the original dimension (I.T. Jolliffe 2002).

In this example, using only the primary principal component, an approximation of the complete data-set is reconstructed and shown in Fig. 3.23. The term that was common to each of the variables is also plotted as a dashed line. It can be seen from the reconstructed data-set that individual anomalies are largely removed, while the




Figure 3.22: Set of 10 correlated variables comprising common and individual constituent parts

common element dominates what remains.




Figure 3.23: Set of 10 variables after filtering process (PCA)

The same concept can be applied to the disturbance estimates obtained by Kalman filtering. Through selection of appropriate subspaces within the multidimensional set of disturbance estimates using PCA, the common disturbances can be accentuated. Conversely, model errors found in individual zones should be suppressed.

At each sample, elements of the vector \hat{x}_d are grouped by geographical location and orientation. Disturbance estimates from south-facing zones, for example, can be categorized in a set indexed by the indicial set S . Where $\{s_1, \dots, s_{|S|}\} \in S$ and $|S|$ is the cardinal of S , the disturbance estimates of the south-facing zones at sample k can be expressed as $\chi_s(k) = [\hat{x}_{d_{s_1}}(k), \dots, \hat{x}_{d_{s_{|S|}}}(k)]^T$.

The vectors which best describe the variance of the data in each set are then found through eigendecomposition. For a batch of M samples, the sample covariance $\Sigma_s \in$

$\Re^{|S| \times |S|}$ of the south-facing zone disturbance estimates $\chi_s \in \Re^{|S| \times M}$ is found as:

$$\bar{\chi}_s = \frac{1}{M} \sum_{i=1}^M \chi_s(i) \quad (3.29)$$

$$\tilde{\chi}_s(k) = \chi_s(k) - \bar{\chi}_s \quad (3.30)$$

$$\mathbf{X}_s = [\tilde{\chi}_s(1) \cdots \tilde{\chi}_s(M)] \quad (3.31)$$

$$\Sigma_s = \frac{1}{M-1} \mathbf{X}_s \mathbf{X}_s^T \quad (3.32)$$

The set of orthonormal eigenvectors $\mathbf{\Gamma}_s = [\mathbf{v}_1 \dots \mathbf{v}_{|S|}]$ and corresponding eigenvalues $\mathbf{\Lambda}_s = \text{diag}[\lambda_1 \dots \lambda_{|S|}]$ is found to satisfy:

$$\Sigma_s = \mathbf{\Gamma}_s \mathbf{\Lambda}_s \mathbf{\Gamma}_s^T, \quad (3.33)$$

where $\mathbf{\Lambda}_s$ is arranged so that $\lambda_1 \geq \lambda_2 \geq \dots \lambda_{|S|}$. The first p principal components at the k^{th} sample can then be constructed from a linear combination of the first p eigenvectors by:

$$\begin{bmatrix} \zeta_1 \\ \vdots \\ \zeta_p \end{bmatrix}_k = [\mathbf{v}_1 \quad \cdots \quad \mathbf{v}_p]^T \tilde{\chi}_s(k) \quad (3.34)$$

The number of principal components, p , can be determined by observing the spectral decomposition of Σ_s . The sample-variance of the j^{th} row of χ_s is given by $\sum_{i=1}^{|S|} \lambda_i v_{j,i}^2$. The objective is to isolate a low number of components which account for as much of the variance as possible. Appropriate selection of this number can be achieved by considering the contribution of the r^{th} component to this variance, given by:

$$\psi_{j,r} = \frac{\lambda_r v_{j,r}^2}{\sum_{i=1}^{|S|} \lambda_i v_{j,i}^2} \quad (3.35)$$

With a suitable selection of p , an approximation of the original disturbances can be defined as:

$$\begin{bmatrix} \tilde{\mathbf{x}}_{d_{S_1}} \\ \vdots \\ \tilde{\mathbf{x}}_{d_{S_{|S|}}} \end{bmatrix}_k \triangleq [\mathbf{v}_1, \dots, \mathbf{v}_p] \begin{bmatrix} \zeta_1 \\ \vdots \\ \zeta_p \end{bmatrix}_k + \bar{\chi}_s \quad (3.36)$$

The spatially filtered disturbance estimates from each group are then reassembled at

each sample and arranged in the same order as the original Kalman filtered disturbance estimate ($\hat{\mathbf{x}}_d(k)$). This vector of spatio-temporally filtered disturbances is denoted $\hat{\mathbf{x}}'_d(k)$.

3.7.3 Incorporating Disturbance Estimation in the Identification Process

The disturbance estimate vector $\hat{\mathbf{x}}'_d(k)$ can be used to refine the model parameter estimates by including it as an additional input in the identification process, the model used now having the form:

$$\hat{\mathbf{x}}(k+1) = A(\boldsymbol{\theta}_q)\hat{\mathbf{x}}(k) + \begin{bmatrix} B(\boldsymbol{\theta}_q) & C^T \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \hat{\mathbf{x}}'_d \end{bmatrix}_k \quad (3.37)$$

$$\hat{\mathbf{y}}(k) = C\hat{\mathbf{x}}(k) \quad (3.38)$$

A new estimate is obtained for the building model, which is then in turn used to recompute the Kalman filtered disturbance estimates. The PCA procedure is carried out once again and the identification is repeated. By iteratively following the process of identification, estimation and spatial filtering until some convergence criteria have been met, a more accurate model can be extracted from the data.

Fig. 3.24 shows the process of identification, grouping and reduction for a set of 30 zones, grouped here in 4 orientation categories.

An additional aspect that requires consideration is that as the identification process progresses, model errors common to several zones begin to appear in the spatio-temporally filtered disturbance estimates. It was previously assumed that model error would be specific to each individual zone and as such would be suppressed by the PCA reduction. As the same external temperature affects all the zones however, errors in the gains associated with it will have a common shape in all zones and as such, will not be suppressed. The steps taken to avoid this are next outlined.

The external temperature profile (T_{amb}) is included in each set of Kalman filtered disturbance estimates (as well as any other known common inputs). The new set of

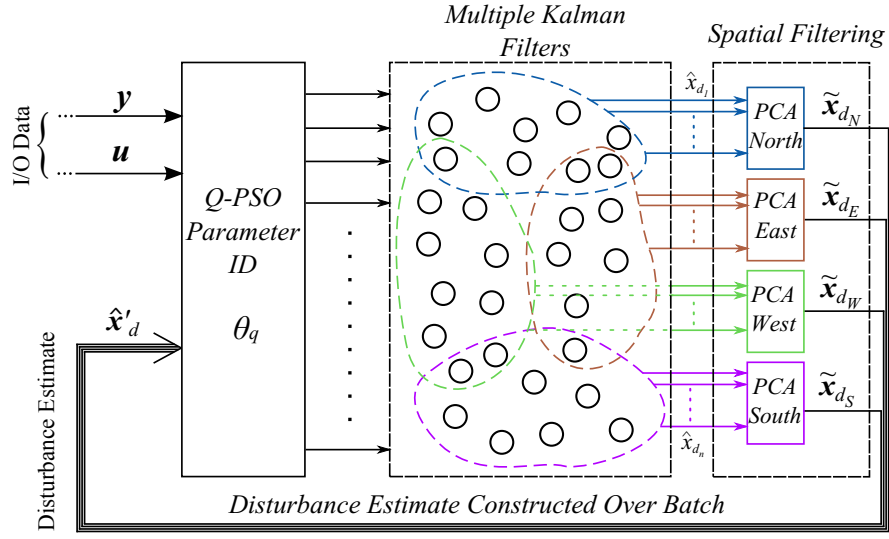


Figure 3.24: Disturbance Estimation - Zones are grouped by geographical similarity. Zone models and disturbance estimates are determined iteratively

south facing disturbances for example, would be given by:

$$\chi'_s(k) = \begin{bmatrix} \chi_s(k) \\ T_{amb}(k) \end{bmatrix} \quad (3.39)$$

The component contributing the most to the variance of the external temperature profile is found by determining the largest element of:

$$\psi_{(|S|+1),r} = \frac{\lambda_r v_{(|S|+1),r}^2}{\sum_{i=1}^{(|S|+1)} \lambda_i v_{(|S|+1),i}^2} \quad (3.40)$$

$$\forall r = 1, 2, \dots, (|S| + 1)$$

The associated component is then removed from the disturbance rebuilding process of (3.34) and (3.36), and replaced by $(p + 1)^{th}$ component.

A conservative approach is taken to the identification process - a low number of iterations is carried out for each update. This prevents the algorithm from over-fitting to the initial disturbance estimate (which, as the initial model is incorrect, must also be assumed to be incorrect). Likewise, if the disturbance estimates are directly included in the identification algorithm without being spatially filtered, the search may not seek to improve upon the initial estimate, as the effect of the errors will appear to be fully accounted for by the disturbances. The gradual removal of the modelling errors from the disturbance estimates through successive iterations is an important aspect of the strategy.

3.7.4 Chemotaxis Algorithm with Disturbance Estimation

The inputs, outputs and disturbances used in *Section 3.4* are implemented with the test network once again to determine the performance of the chemotaxis algorithm. In this case, the disturbance-estimation procedure outlined above is now included. Zones are categorized in terms of orientation within the building, i.e. facing south, north, east or west with an additional group included comprising of the centrally located zones.

The number of principal components used for each disturbance estimate was determined using the spectral decomposition from (3.35). As a graphical example of the selection process, the contribution of the first 8 components to the initial disturbance estimates from the central set of zones is shown in Fig. 3.25. It can be seen that the first three components alone describe over 90% of the variance of most of the individual disturbances.

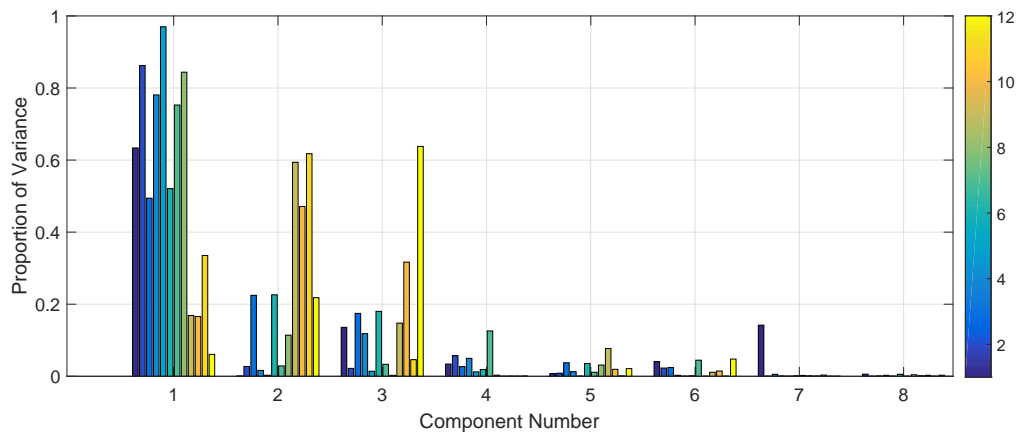


Figure 3.25: Spectrum of variance described by first 8 principal components for central zone disturbances

Two principal components were used here for the central set of zones, while only the primary component was used for all other groups. After every 10 iterations of the identification process, the disturbance estimate was updated. The identification and estimation process was then repeated until the search algorithm converged. The solution is evaluated in Table 3.3, once again using a value for γ_{ch} of 0.01, using the same performance criteria as in the previous sections.

A significant improvement can be seen in terms of the output error. The mean relative error of the network parameters has not improved, though the dominant components (the zone capacitors) are more accurately represented. As no mechanism exists to allow a move which increases the cost, an algorithm of this class will be prone to convergence at local extrema. To achieve better results, a more exploratory search is

Table 3.3: Model error present before and after chemotaxis adjustment using test network with disturbance estimation

NETWORK ADJUSTMENT USING CHEMOTAXIS			
	Initial Network	Disturbance Estimate Temporally Filtered	Disturbance Estimate Spatio-Temporally Filtered
$Train_{err}$ ($^{\circ}C.hr$)	216.00	140.85	126.82
Val_{err} ($^{\circ}C.hr$)	276.35	152.21	126.23
$Param_{err}$ (%)	31.62	37.00	33.34
C_{err} (%)	34.48	30.08	24.44

required.

3.8 Identification Algorithm Based on Particle Swarm Optimisation and Disturbance Estimation

In place of the chemotaxis search algorithm used in the identification procedure, a standard Particle Swarm Optimisation (PSO) algorithm is implemented, incorporating disturbance estimates so as to achieve a better estimate of the test network.

The pseudo-code for PSO is outlined in Algorithm 2, where the position of each D -dimensional agent i , within a population of size Ω , is denoted z_i , its previous lowest cost position (or local best) is denoted P_{b_i} , vel_i is the agent's velocity parameter, the global minimum cost of all previous positions of the swarm is denoted G_b , and $f(z_i)$ is the fitness function to be minimised.

The adjustable parameters ω , ϕ_p and ϕ_g represent an agents inertia, attraction to P_b and attraction to G_b respectively. These control the convergence speed of the algorithm as well as providing an insight into the exploration/exploitation balance of the search. A higher value of ϕ_g relative to ϕ_p indicates that the swarm will be more strongly drawn towards the global best position. A larger value of ϕ_p encourages a greater level of exploration for particles away from G_b . The inertia parameter can be reduced as the algorithm progresses to allow for the search to become finer as the agents begin to converge.

Algorithm 2 Particle Swarm Optimisation

```

1: for i in 1.. $\Omega$  do                                ▷ Initialize population position and velocity
2:   for j in 1.. $D$  do
3:      $z_{i,j} \sim U(z_{j_{lo}}, z_{j_{hi}})$ ;
4:      $vel_{i,j} \sim U(-|z_{j_{hi}} - z_{j_{lo}}|, |z_{j_{hi}} - z_{j_{lo}}|)$ ;
5:      $P_{b_{i,j}} = z_{i,j}$ ;
6:   end for
7:    $G_b = \arg \min f(P_{b_i})$ ;
8: end for
9: for ii in 1.. $T$  do                                ▷ Number of Iterations
10:  for i in 1.. $\Omega$  do                                ▷ Number of agents
11:    for j in 1.. $D$  do                                ▷ Dimension of agent
12:       $r_p \sim U(0, 1)$ ;  $r_g \sim U(0, 1)$ ;
13:       $vel_{i,j} = \omega vel_{i,j} + \phi_p r_p (P_{b_{i,j}} - z_{i,j}) + \phi_g r_p (G_{b_j} - z_{i,j})$ ;
14:       $z_{i,j} = z_{i,j} + vel_{i,j}$ ;
15:    end for
16:    if  $f(z_i) < f(P_{b_i})$  then  $P_{b_i} \leftarrow z_i$ 
17:    end if
18:    if  $f(P_{b_i}) < f(G_b)$  then  $G_b \leftarrow P_{b_i}$ 
19:    end if
20:  end for
21: end for

```

3.8.1 PSO - Implementation with test network

Using the same training data as previously outlined, a PSO algorithm was implemented with 20 agents. The inertia ω was set to decrease linearly from 0.9-0.45, while the tuning parameters ϕ_p and ϕ_g were set conservatively at 0.04. The PCA-based disturbance estimates were incorporated into the algorithm in the same manner (and with the same groupings and principal component numbers) as for the chemotaxis approach. Candidate solutions produced by agents that resulted in unstable model poles or negative component values were discarded, with the agent returning to its previous permissible position and velocity.

The performance of the algorithm is evaluated in Table 3.4.

The results here show no clear improvement when compared to those from the chemotaxis approach. The high dimensional search-space, along with the low excitation of the input data results in a multitude of potential lower-cost directions for each agent. Remaining within a stable, physically realistic search-space necessitated a conservative search, which tended to converge quickly in a similar manner to the chemotaxis algorithm. Given the deliberately incremental separation of disturbance and model

Table 3.4: Model error present before and after PSO adjustment incorporating disturbance estimation using test network

NETWORK ADJUSTMENT USING PSO		
	Initial Network	PSO/PCA
$Train_{err}$ (°C.hr)	216.00	110.12
Val_{err} (°C.hr)	276.35	124.78
$Param_{err}$ (%)	31.62	33.13
C_{err} (%)	34.48	27.51

error from the Kalman filtered disturbance estimates, maintaining the stability of the population while slowly decreasing the cost is of more importance than immediately locating a lower cost position. When a more aggressive strategy is adopted, if an outlying agent locates a globally optimum position, the population tends to be quickly drawn towards a position that potentially lacks physical meaning.

To achieve a more expansive search, without compromising the stability of the population, a modified version of traditional PSO, namely Quantum-behaved Particle Swarm Optimisation (Q-PSO) as developed in (Sun, Xu & Feng 2004, Sun, Feng & Xu 2004) was investigated.

3.8.2 Quantum-behaved Particle Swarm Optimisation

Using traditional PSO, agents are attracted only to the local and global best positions. Agents in close proximity to the global best converge quite quickly, while outlying agents have no influence on the population. In addition to the local and global best attractor points of PSO, agents in a population of a Q-PSO search are influenced by a mean-best term, M_b , which represents the mean local best position of all Ω agents in the population, given for the j^{th} dimension by:

$$M_{b_j} = \sum_{i=1}^{\Omega} (P_{b_{i,j}}) \quad (3.41)$$

The agents in the vicinity of G_b are now also influenced by outlying agents through this mean best term, enabling a more extensive search of the region around G_b as the outlying agents are slowly drawn in. Furthermore, the mean best term has the ef-

fect of anchoring the population in the event of an outlying agent discovering a global optimum in an unrealistic location. A wider, more aggressive search can then be undertaken without compromising the stability of the general population. The pseudo-code for Q-PSO is outlined in Algorithm 3. Here, α_0 and α_1 are tuning parameters which determine the initial and final values of the contraction-expansion coefficient, which controls the convergence speed of the algorithm. A further analysis of the algorithm is given in Appendix B.

Algorithm 3 Quantum-behaved Particle Swarm Optimisation

```

1: for i in 1.. $\Omega$  do                                ▷ Initialize population
2:    $z_i \sim U(z_{lo}, z_{hi});$ 
3:    $P_{b_i} = z_i; G_b = \arg \min f(P_{b_i});$ 
4: end for
5: for ii in 1.. $T$  do                                ▷ Number of Iterations
6:   for j in 1.. $D$  do                                ▷ Dimension of Agents
7:      $M_{b_j} = \frac{1}{\Omega}(\sum_{i=1}^{\Omega} P_{b_{i,j}});$ 
8:   end for
9:    $\alpha = \alpha_0 + \frac{1}{T}(\alpha_1 - \alpha_0)(T - ii);$ 
10:  for i in 1.. $\Omega$  do
11:    for j in 1.. $D$  do
12:       $\phi \sim U(0, 1); \rho \sim U(0, 1); \delta \sim U(0, 1);$ 
13:       $\gamma_{i,j} = \phi P_{b_{i,j}} + (1 - \phi)G_{b_j};$ 
14:      if  $\delta > 0.5$  then
15:         $z_{i,j} = \gamma_{i,j} + \alpha|M_{b_j} - z_{i,j}|\log(\frac{1}{\rho});$ 
16:      else
17:         $z_{i,j} = \gamma_{i,j} - \alpha|M_{b_j} - z_{i,j}|\log(\frac{1}{\rho});$ 
18:      end if
19:    end for
20:    if  $f(z_i) < f(P_{b_i})$  then  $P_{b_i} \leftarrow z_i$ 
21:    end if
22:    if  $f(P_{b_i}) < f(G_b)$  then  $G_b \leftarrow P_{b_i}$ 
23:    end if
24:  end for
25: end for

```

3.8.2.1 Q-PSO - Implementation with test network

Initializing the algorithm with a contraction-expansion coefficient of 1, linearly reducing to 0.4, the Q-PSO algorithm (with a population of 20 agents) was implemented by applying the training data to the test network. Once again, the spatially filtered disturbance estimates were incorporated with the same PCA groupings used. The performance of this approach is evaluated in Table 3.5.

Table 3.5: Model error present before and after Q-PSO adjustment using test network

NETWORK ADJUSTMENT USING Q-PSO		
	Initial Network	Q-PSO/PCA
$Train_{err}$ (°C.hr)	216.00	100.32
Val_{err} (°C.hr)	276.35	104.74
$Param_{err}$ (%)	31.62	30.27
C_{err} (%)	34.48	24.17

This algorithm shows an improved performance in both output error and parameter accuracy when compared to the PSO and chemotaxis approaches. Observing the resulting network however, while on average a slight improvement in the relative error of the component values was obtained, a certain number of individual resistance values deviated significantly from the initial estimate. In some cases, the identified parameters were over 25 times larger than the initial value, an indication that the link to the physical dimensions and materials of the buildings was lost. With such a large number of parameters to be adjusted, those of lower dominance are inclined to drift. As such, the next step taken was to break-up the search.

Given the low number of iterations along with the high dimensional search space, certain parameters may deviate over time from physically realistic values, provided the overall cost continues to decrease. This can be seen in the previous results, in which, though the output error is reduced, certain parameters can move quite far from the target values. To discourage this behaviour, in place of a centralised identification process whereby all parameters are adjusted simultaneously, a decentralised strategy was adopted.

Progressing through the building, only parameters associated with one zone at a time were adjusted. As there is interaction between the different zones, changing the parameters of one zone may affect the selection of parameters in another. To allow for this, only a low number of iterations were completed in one zone before moving to the next. After all zones in the building were adjusted, the disturbance estimates were updated and the process of identification and estimation was repeated iteratively as before.

3.8.2.2 Q-PSO Decentralised Search

For the decentralised search, as the dimension of each adjustment was greatly reduced, the number of agents used was lowered to 10. Table 3.6 shows the results obtained for the search.

Table 3.6: Model error present before and after decentralised Q-PSO adjustment using test network

NETWORK ADJUSTMENT USING DECENTRALISED Q-PSO		
	Initial Network	Q-PSO/PCA
$Train_{err}$ (°C.hr)	216.00	104.58
Val_{err} (°C.hr)	276.35	123.51
$Param_{err}$ (%)	31.62	28.09
C_{err} (%)	34.48	26.58

The output error of the adjusted network is slightly higher than for the centralised Q-PSO approach. Similarly, the relative error of the capacitance values is higher, though the relative error of the overall network is reduced. Crucially, the largest deviation in component value, at 70%, is far lower than that of the centralised search. The bode plots in Fig. 3.26-Fig. 3.29 show the frequency response of the transfer function between the input heat flow and output temperature in Zone 1 and Zone 25. The transfer function between the external temperature and output for the same two zones is also shown.

In Fig. 3.30 the output of Zone 17 is shown for the undisturbed validation data-set using the target, adjusted and initial networks. There is a particularly noticeable improvement in the steady-state gain of the identified system.

For a real building, the "target network" is more of an abstraction than a physical reality, while data quality is likely to be reduced. Misuse of the flexibility present in a large scale RC-network must be avoided. As such, the stability of the decentralised approach, despite a slightly compromised performance, is favoured here as the most suitable candidate for use with real building data.




Figure 3.26: Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 1 using target system, initial system and decentralised Q-PSO-identified systems (with disturbances and disturbance estimation included)




Figure 3.27: Frequency response comparison of transfer function between external temperature and output temperature for Zone 1 using target system, initial system and decentralised Q-PSO-identified systems (with disturbances and disturbance estimation included)

3.9 Case Study - Modelling the Nimbus Centre from Process Data

The techniques developed in this chapter were implemented so as to obtain a simulation platform representative of the Nimbus Centre, a two-storey office building located within the campus of Cork Institute of Technology. The building is the site of the National Sustainable Building Energy Test-bed, and includes a solar-powered wireless sensor network and a micro-grid used for research purposes. The building has an inte-




Figure 3.28: Frequency response comparison of transfer function between heat supply (from radiator) and output temperature of Zone 25 using target system, initial system and decentralised Q-PSO-identified systems (with disturbances and disturbance estimation included)




Figure 3.29: Frequency response comparison of transfer function between external temperature and output temperature for Zone 25 using target system, initial system and decentralised Q-PSO-identified systems (with disturbances and disturbance estimation included)

rior floor area of $1868m^2$, comprising office and laboratory space for approximately 80 people. A hydronic heating system uses radiators to deliver heat to the different zones.

Each zone of the building can communicate to a BMS system via the wireless temperature sensors. The external temperature is also measured and communicated. Additional temperature sensors provide measurements of the flow and return temperature of the water flowing through the network between the mixing valves and zones as well as the supply and return temperatures of the water to the boiler. Flow rates of the water in the system are also measured as well as set-points and mixing-valve positions.




Figure 3.30: Validation data-set output (temperature of Zone 17) using identified network and initial estimated network compared to the target network



Figure 3.31: The Nimbus Centre

The standard control strategy of the heating system in the Nimbus Centre can be viewed as a three layer hierarchy. At the highest level, boilers supply sufficient heat to maintain a weather-compensated (varying linearly with the external temperature) set-point in the system header. The header supplies heat to the entire building. In the middle layer of the hierarchy, PI controlled mixing valves on each floor mix supply-water from the header with return-water from the radiators to achieve a floor-level temperature set-point (once again, weather-compensated). The lowest control layer (zone-level) consists of thermostatically controlled radiator valves which are actuated depending on the temperature within the zone. These valves operate on a hysteresis loop, remaining open until the temperature in the zone exceeds 0.5°C above the set-point before closing and remaining closed until the zone temperature falls 0.5°C below the set-point. A schematic of the ground floor section of the heating system can be seen in Fig. 3.32.




Figure 3.32: Schematic of ground floor section of heating system used in the Nimbus Centre

3.9.1 Processing Raw Data for Use in Identification Strategy

3.9.1.1 Heating System Sensor Data

Measurement data for the 28 day period of February 2015 was used to train and validate the models. Heating system measurements (flow rates, flow-temperatures and set-points) were available at one minute samples throughout. For illustrative purposes, in Fig. 3.33, a plot of the boiler supply and return temperatures can be seen for a three day period:




Figure 3.33: Supply and demand boiler water temperature measurements taken for a period of three days

As the building has been in operation with many occupants for the past number of years, long term functional tests were not possible (carrying out of functional tests over

a significant time-scale would have a large negative impact on users of the building). As such, only data-sets from periods of standard operation were available.

Measurements of the heat flow to each zone were not available, nor were the positions of the thermostatic valves in each zone and so the inputs which, for the RC-network, were desired to be heat flows in Watts, had to be inferred from other data. The valve positions were found by observing the temperature measurements in each zone as the set-points and thresholds were known.

The relationship between the radiator flow-temperature in a single zone is found in the following way, with the radiator heat transfer coefficient assumed to be constant. This assumption greatly simplifies the modelling process (a method for incorporating the nonlinearity of the heat transfer coefficient in a linear model is given in (Hazyuk et al. 2012b)). Where this coefficient, multiplied by the radiator area is denoted UA , and the supply and return temperatures of the water to the radiators are denoted T_{fl} and T_{ret} respectively, the heat Q_z transferred from a radiator to a zone of temperature T_z can be expressed as:

$$Q_z(t) = UA \left(\frac{T_{fl}(t) + T_{ret}(t)}{2} - T_z(t) \right) \quad (3.42)$$

The heat transferred to the zone could also be taken as the heat removed from the water flowing through the radiator. Where $\dot{m}_{fl}(t)$ is the flow rate and C_p is the specific heat capacity of water, this could be expressed as:

$$Q_z(t) = \dot{m}_{fl}(t)C_p (T_{fl}(t) - T_{ret}(t)) \quad (3.43)$$

(3.42) and 3.43 can be equated and rearranged as follows to produce an equation for T_{ret} :

$$UA \left(\frac{T_{fl}(t) + T_{ret}(t)}{2} - T_z(t) \right) = \dot{m}_{fl}(t)C_p (T_{fl}(t) - T_{ret}(t)) \quad (3.44)$$

$$T_{ret}(t) = \frac{T_{fl}(t) (2\dot{m}_{fl}(t)C_p - UA) + T_z(t) (2UA)}{2\dot{m}_{fl}(t)C_p + UA} \quad (3.45)$$

Substituting this representation of T_{ret} back into 3.42 or 3.43 produces:

$$Q_z(t) = \frac{2UA\dot{m}_{fl}(t)C_p}{2\dot{m}_{fl}(t)C_p + UA} (T_{fl}(t) - T_z(t)) \quad (3.46)$$

As all the quantities on the right-hand-side of 3.46 are known or measured, the heat inputs to each zone $Q_z(t)$ can be found.

3.9.1.2 Zone and External Temperature Sensor Data

The data quality obtained from the zone temperature sensors was less consistent. Communication between the sensors and the BMS occurred once again every minute, but the signals only updated in increments of 0.1568°C . In certain cases, despite the availability of hundreds of data-points, this resolution could result in an unchanged temperature reading for as long as 2-3 hours.

To avoid long periods of unexcited data, a cubic spline interpolation was used between successive distinct points. The resulting signal is then used as the output for that particular zone. The sensor measuring the external temperature had a one minute sample time, updating in temperature increments of $10^{-4^{\circ}}\text{C}$. The signal was used as a measured disturbance without interpolating the data.

3.9.2 Physics-based RC-Network Representing Initial Estimate of Nimbus Centre

The building was first described as an RC-network with 2 floors and 25 zones (11 on the ground floor, 14 on the first). Walls, floors, ceilings and windows were all included to produce a network of 835 components. Once again, a 3R-2C configuration was used to represent each wall, ceiling and floor, single resistances were used for windows and single capacitances were used for airspaces.

Conversion of the dimensions and material properties of the building into suitable resistance and capacitance values was completed by United Technologies Research Centre Ireland (UTRCI).

3.9.3 Training of Physics-based Network Using Measured Data

The strategy developed in this chapter - that of a decentralised Q-PSO identification process allied to a PCA-based form of disturbance estimation - was implemented using the data obtained from the Nimbus Centre to train the physics-based RC-network. In this case, the target system is clearly unknown, as well as the unmeasured disturbance. As such, the performance is assessed only in terms of the similarity of the model outputs to the measurements. The tuning parameters that performed best when using the test network were used once again here.

A training period of 15 days was used, with a sample time of 600 seconds (the larger

sample time relative to the test network was chosen as a consequence of the poor output measurement resolution). In the same manner as the test network, the Q-PSO procedure was used to gradually adjust parameters of one zone at a time, with disturbance estimation carried out after all zones were updated. 20 iterations of this Q-PSO/PCA procedure were completed. To estimate the disturbances by PCA, zones were categorized into groups representing the north-facing, south-facing, east-facing, west-facing and central zones respectively. 2 principal components were used to filter the central zones, while a single component was used for the rest.

The output accuracy of two zones (specifically the ground floor meeting room and corridor) of the system are compared here to those of the real building as well as the original physics based network in Fig. 3.34 and Fig. 3.35. These outputs are determined (or measured) for an 8 day period of unseen data.




Figure 3.34: Comparison of real data (Nimbus Centre) with simulation platform: Q-PSO adjusted vs. original RC network comparison with real data - Ground Floor Meeting Room

In Table 3.7, the average daily output error for all the adjusted zones of the building is shown for this validation data-set (once again in $^{\circ}\text{C}.\text{hr}$). For this real-data identification (as an undisturbed data-set cannot be obtained from the real building) the disturbances are included in the output comparison. A far more accurate output is obtained after the identification process, particularly in terms of steady state gain.

3.10 Modelling of the Nimbus Centre Heating System

To simulate different heating strategies in the Nimbus Centre, a model of the heating system itself was also required in the simulation platform, as well as the RC-network




Figure 3.35: Comparison of real data (Nimbus Centre) with simulation platform: Q-PSO adjusted vs. original RC network comparison with real data - Ground Floor Corridor

Table 3.7: Mean daily output error present before and after Q-PSO adjustment of RC-network compared to data measured in Nimbus Centre

NETWORK ADJUSTMENT USING Q-PSO		
	Initial Network	Q-PSO/PCA
Val_{err} (°C.hr)	1629.38	553.04

model of the thermal dynamics. The development in Simulink of the three levels of the heating system hierarchy are described in this section.

3.10.1 Boiler & Header Thermodynamic Models

The heat flow from the boiler $Q_{Bo}(t)$ to the header, as well as the heat removed from the header and delivered to the building $Q_{Dem}(t)$ can be expressed in terms of the header temperature $T_h(t)$ and the input power to the boiler $P_{in}(t)$ as:

$$Q_{Bo}(t) = Q_{Dem}(t) + m_h C_p \dot{T}_h(t) \quad (3.47)$$

$$P_{in}(t)\eta(t) = Q_{Bo}(t) \quad (3.48)$$

where m_h is the mass of water stored in the header and the efficiency, η , is based on manufacturer's efficiency curves. The block diagram of this level of the system is shown in Fig. 3.36.

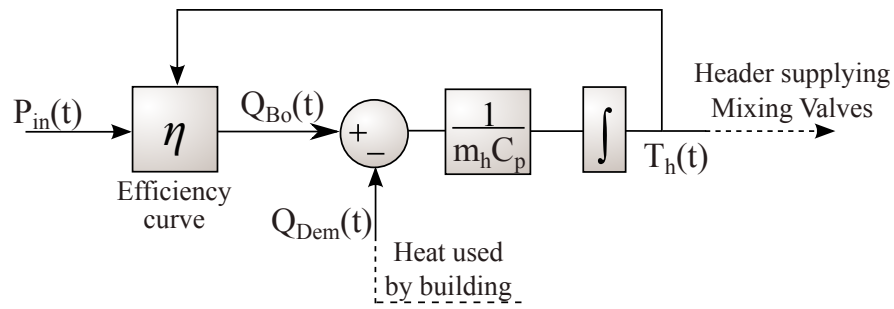


Figure 3.36: Block diagram of header dynamics with heat added by the boiler and heat removed by the building

3.10.2 Modelling the PI-Controlled Mixing Valves

The mixing valve uses a PI controller to adjust its position, $M(t)$, so that the output $T_{fl}(t)$ is as close as possible to the set-point $T_{fl_{sp}}$. If the valve is in the fully open position ($M = 1$), the output is equal to the supply coming from the header $T_h(t)$. On the other hand, when the valve is fully closed, with $M = 0$, the output is equal to the water returning from the building given in terms of the individual return temperatures from each zone as:

$$T_{Ret}(t) = \frac{1}{\sum_{j=1}^N \dot{m}_{fl_j}} \sum_{i=1}^N \dot{m}_{fl_i} T_{ret_i}(t) \quad (3.49)$$

The relationship between the valve position and the input and output flow-temperatures is as follows:

$$T_{fl}(t) = M(t)T_h(t) + (1 - M) T_{Ret}(t), \quad (3.50)$$

where $M(t)$ can vary between 0 and 1.

Using (3.45) and (3.49), at time t , the return water temperature $T_{Ret}(t)$ can be found in terms of the flow-temperature $T_{fl}(t)$ and the zone temperatures. Observing (3.50) and considering the dynamics of the mixing valve however, the flow-temperature at time t clearly depends on the return temperature at time t , not the other way around. For the sake of both realism and causality, the time required for the water to get from the output of the mixing valve, through the radiators of the building and back to the return inlet of the mixing valve must be taken into account.

As a detailed description of the individual time delays associated with each section of the system could become quite complex, with little noticeable effect, for simplicity the entire delay is assigned to the flow from the radiators back to the mixing valve. In

place of (3.49), the return temperature is given by:

$$T_{Ret}(t) = \frac{1}{\sum_{j=1}^N \dot{m}_{fl_j}} \sum_{i=1}^N \dot{m}_{fl_i} T_{ret_i}(t - \tau_{ret}), \quad (3.51)$$

where the delay τ_{ret} is chosen as 50 seconds to allow for the water to fully circulate (determining a value for this was not possible given the available resolution of the data).

The valve itself is assumed to have first order dynamics. The relatively long sample-time for any set of measurements taken from the building prevented any proper determination of these dynamics and so the time constant (given here by τ) is chosen to be 20 seconds, giving the valve relatively fast dynamics.

The control loop for the valve considers T_h and T_{Ret} as disturbances to maintain the linearity of the loop. The proportional and integral gains are chosen as 0.05 and 3×10^{-4} to attain a reasonable balance of overshoot and rise time for typical values of T_h and T_{Ret} . An anti-windup mechanism is also included to prevent integral windup.

In Fig. 3.37, the control loop for one of the mixing valves is shown, where p is the differentiation operator.

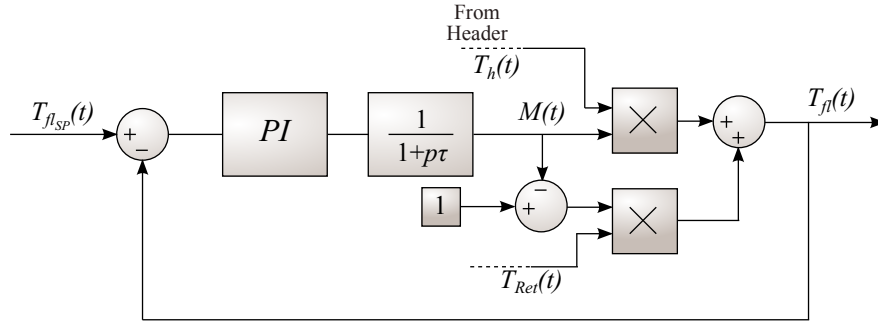


Figure 3.37: Block diagram of mixing valve control loop

3.10.3 Modelling the Zone-level Thermostatic Valves

The zone-level thermostatic valves operate on a hysteresis loop. These valves are implemented in Simulink by comparing the zone temperature $T_z(t)$ and the zone temperature set-point T_{zsp} . If the zone temperature exceeds a designated threshold above the set-point, the valve closes and remains closed until the temperature has fallen under a threshold below the set-point.

Using (3.46), the temperature of water flowing from the mixing valve $T_{flsp}(t)$ and the zone temperature $T_z(t)$ are used to determine the heat flow in Watts to the building.

In Fig. 3.38, the block diagram for the zone valve associated with the i^{th} zone is illustrated.




Figure 3.38: Block diagram of zone valve associated with i^{th} zone

3.10.4 System Pumps and Flow Rate Considerations

The mass flow-rate of the water through the heating system is determined by pumps on each floor, which are controlled to maintain the same pressure in the system at all times so that the water flowing to each zone remains the same regardless of the number of open or closed zone valves on a floor.

In the simulation platform, the dynamics of the pump are not included, with the mass flow-rate assumed to be equal to the set-point. The justification for this is that the flow of heat to the building is quite insensitive to flow-rate change. This is shown in the following analysis.

If the heat flow to a zone using the designed mass flow-rate, \dot{m}_{fl_1} , is denoted Q_1 , while the hypothetical heat flow to the same zone for a different mass flow rate \dot{m}_{fl_2} is denoted Q_2 , using (3.46), the ratio of the heat flows could be expressed as:

$$\frac{Q_1}{Q_2} = \frac{2\dot{m}_{fl_1}\dot{m}_{fl_2}C_p + UA\dot{m}_{fl_1}}{2\dot{m}_{fl_1}\dot{m}_{fl_2}C_p + UA\dot{m}_{fl_2}} \quad (3.52)$$

By setting the relationship between the two mass flow-rates as $\dot{m}_{fl_2} = \alpha_m \dot{m}_{fl_1}$ (so that, for example, double the original flow-rate would require $\alpha_m = 2$), (3.52) can be rearranged as:

$$\frac{Q_1}{Q_2} = \frac{2C_p \frac{\dot{m}_{fl_1}}{UA} + \frac{1}{\alpha_m}}{2C_p \frac{\dot{m}_{fl_1}}{UA} + 1} \quad (3.53)$$

As the pressure drop across the radiator of any zone is approximately proportional to

the area of the radiator, the ratio between the mass flow rates in each zone and the radiator area should be constant in all zones. The heat transfer coefficient U is also set as a constant (a simplification mentioned above), and so the ratio $\frac{\dot{m}_{fl_1}}{UA}$ can be thought of as a design constant for choosing the standard flow-rate \dot{m}_{fl_1} . For the simulation platform, the ratio is set at 0.0013.

With this ratio, the change in Q_2 relative to Q_1 is plotted, with the ratio between the mass flow-rates (α_m) increasing from 0.5 – 2 in Fig. 3.39 to show the change in heat output obtained by changing the mass flow-rate. It can be seen that even if the pumps are operated to produce a mass flow-rate twice as large (or small) as the design flow-rate, the heat transmitted to the building is not significantly changed ($< 10\%$).




Figure 3.39: Change in heat flow to zone as mass flow rate through radiators is increased

3.11 Conclusion

In this chapter, methodologies were introduced for the development of a simulation platform, capable of representing the thermodynamic characteristics of a building for use in control strategy design. Metaheuristic search algorithms were assessed for use in tandem with a disturbance estimation procedure to allow for models to be identified from data in which significant unmeasured disturbances are present.

To develop suitable approaches, a high-order RC-network test model was first created, representing a notional 30 zone building. It was found that, if appropriately configured, a chemotaxis algorithm could identify an accurate representation of the test network from standard input and output disturbance-free data if an initial estimate of the system is available. In the presence of disturbances however, this was shown to not be the case.

A disturbance estimation strategy was next derived to improve the identification process. Unmeasured disturbances were estimated using a centralised Kalman filter, and

spatially filtered by a PCA procedure. By including the spatially filtered estimates as an additional input to the model in the identification process, an improved representation of the target system was obtained.

The chemotaxis algorithm was then replaced by more sophisticated approaches. A particle swarm optimisation (PSO) was first evaluated, followed by a quantum-behaved particle swarm optimisation (Q-PSO). The most accurate approximation of the test network was obtained by using Q-PSO. By decentralising the search, a more robust solution could be obtained, at the expense of output accuracy.

The Q-PSO/PCA identification and disturbance estimation strategy was implemented using measured data from the Nimbus Centre, to refine the parameters of a physics-based RC-network estimate of the thermal dynamics of the building. A model of the heating system used in the building was also developed. The heating system, in tandem with the RC-network constitute a simulation platform which allows for different control strategies to be compared and analysed.

References

- Agbi, C. & Krogh, B. (2014), ‘Decentralized identification of building models’, *2014 American Control Conference* pp. 1070–1075.
- Agbi, C., Song, Z. & Krogh, B. (2012), ‘Parameter identifiability for multi-zone building models’, *Proceedings of the IEEE Conference on Decision and Control* pp. 6951–6956.
- Coakley, D., Raftery, P. & Keane, M. (2014), ‘A review of methods to match building energy simulation models to measured data’, *Renewable and Sustainable Energy Reviews* **37**, 123–141.
- Crawley, D. B., Hand, J. W., Kummert, M. & Griffith, B. T. (2008), ‘Contrasting the capabilities of building energy performance simulation programs’, *Building and Environment* **43**(4), 661–673.
- Duffie, J. A. & Beckman, W. A. (2013), *Solar Engineering of Thermal Processes*, 4th edn, John Wiley & Sons, Inc., Hoboken, NJ, USA.
- Gazi, V. & Passino, K. M. (2011), *Swarm stability and optimization*, Springer.
- Ghiaus, C. (2006), ‘Experimental estimation of building energy performance by robust regression’, *Energy and Buildings* **38**(6), 582–587.
- Harish, V. & Kumar, A. (2016), ‘A review on modeling and simulation of building energy systems’, *Renewable and Sustainable Energy Reviews* **56**, 1272–1292.
- Hazyuk, I., Ghiaus, C. & Penhouet, D. (2012a), ‘Optimal temperature control of intermittently heated buildings using Model Predictive Control: Part I - Building modeling’, *Building and Environment* **51**, 379–387.
- Hazyuk, I., Ghiaus, C. & Penhouet, D. (2012b), ‘Optimal temperature control of intermittently heated buildings using Model Predictive Control: Part II - Control algorithm’, *Building and Environment* **51**, 388–394.

- I.T. Jolliffe (2002), *Principal component analysis*, Vol. 2, 2nd edn, Springer, New York.
- Karatasou, S., Santamouris, M. & Geros, V. (2006), 'Modeling and predicting building's energy use with artificial neural networks: Methods and results', *Energy and Buildings* **38**(8), 949–958.
- Kennedy, J. & Eberhart, R. C. (1995), 'Particle swarm optimization', in *Proc. IEEE International Conference on Neural Networks* pp. 39–43.
- Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P. (1983), 'Optimization by Simulated Annealing', *Science* **220**(4598), pp. 671–680.
- Kulkarni, R. V. & Venayagamoorthy, G. K. (2011), 'Particle Swarm Optimization in Wireless-Sensor Networks: A Brief Survey', *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* **41**(2), 262–267.
- Kumar, R., Aggarwal, R. & Sharma, J. (2013), 'Energy analysis of a building using artificial neural network: A review', *Energy and Buildings* **65**(2), 352–358.
- Ma, Y., Kelman, A., Daly, A. & Borelli, F. (2012), 'Predictive control for energy efficient buildings with thermal storage', *IEEE Control Systems Magazine* **32**(February), 44–64.
- Neto, A. H. & Fiorelli, F. A. S. (2008), 'Comparison between detailed model simulation and artificial neural network for forecasting building energy consumption', *Energy and Buildings* **40**(12), 2169–2176.
- O'Dwyer, E., Tommasi, L. D., Kouramas, K., Cychowski, M. & Lightbody, G. (2016), 'Modelling and Disturbance Estimation for Model Predictive Control in Building Heating Systems', *Energy and Buildings* p. (Under Review).
- Oldewurtel, F., Parisio, A., Jones, C. N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B. & Morari, M. (2012), 'Use of model predictive control and weather forecasts for energy efficient building climate control', *Energy and Buildings* **45**, 15–27.
- Oldewurtel, F., Sturzenegger, D. & Morari, M. (2013), 'Importance of occupancy information for building climate control', *Applied Energy* **101**, 521–532.
- Page, J., Robinson, D., Morel, N. & Scartezzini, J. L. (2008), 'A generalised stochastic model for the simulation of occupant presence', *Energy and Buildings* **40**(2), 83–98.

- Pérez-Lombard, L., Ortiz, J. & Pout, C. (2008), 'A review on buildings energy consumption information', *Energy and Buildings* **40**(3), 394–398.
- Privara, S., Vana, Z., Gyalistras, D., Cigler, J., Sagerschnig, C., Morari, M. & Ferkl, L. (2011), 'Modeling and identification of a large multi-zone office building', *2011 IEEE International Conference on Control Applications (CCA)* pp. 55–60.
- Široký, J., Oldewurtel, F., Cigler, J. & Prívara, S. (2011), 'Experimental analysis of model predictive control for an energy efficient building heating system', *Applied Energy* **88**(9), 3079–3087.
- Sun, J., Feng, B. & Xu, W. (2004), 'Particle swarm optimization with particles having quantum behavior', *Proceedings of the 2004 Congress on Evolutionary Computation (IEEE Cat. No.04TH8753)* **1**, 325–331.
- Sun, J., Xu, W. & Feng, B. (2004), 'A global search strategy of quantum-behaved particle swarm optimization', *IEEE Conference on Cybernetics and Intelligent Systems, 2004.* **1**, 1–3.
- Turing, A. M. (1950), 'Computing Machinery and Intelligence', *Mind* **49**, 433–460.
- Valle, Y. D. (2008), 'Particle swarm optimization: basic concepts, variants and applications in power systems', *Evolutionary Computation* **12**(2), 171–195.
- Van Doren, J. F., Douma, S. G., Van den Hof, P. M., Jansen, J. D. & Bosgra, O. (2009), 'Identifiability: From Qualitative Analysis to Model Structure Approximation', *15th IFAC Symposium on System Identification* pp. 664–669.
- Van Doren, J. F. M., Van den Hof, P. M. J., Jansen, J. D. & Bosgra, O. H. (2008), 'Determining identifiable parameterizations for large-scale physical models in reservoir engineering', *Proceedings of the 17th IFAC World Congress, Seoul (2005)*, 6–11.
- Willis, M., Di Massimo, C., Montague, G., Tham, M. & Morris, A. (1991), 'Artificial neural networks in process engineering', *IEE Proceedings D Control Theory and Applications* **138**(3), 256.
- Wolpert, D. & Macready, W. (1997), 'No free lunch theorems for optimization', *IEEE Transactions on Evolutionary Computation* **1**(1), 67–82.
- Yang, R. & Wang, L. (2012), 'Multi-objective optimization for decision-making of energy and comfort management in building automation and control', *Sustainable Cities and Society* **2**(1), 1–7.

- Zhao, H. X. & Magoulès, F. (2012), ‘A review on the prediction of building energy consumption’, *Renewable and Sustainable Energy Reviews* **16**(6), 3586–3592.

Chapter 4

Development of Low-Order Building Zone Models Suitable for Predictive Control Applications

4.1 Introduction

In this chapter, identification methods for models designed to be used within a building heating system MPC strategy are investigated. Physics-based approaches which take into account the physical make-up of the building are first considered. As a scalable alternative, data-driven formulations are then evaluated. The disturbance estimation techniques developed in Chapter 3 are incorporated in the modelling process to reduce the corrupting impact of unmeasured disturbances. The different approaches are assessed using a high order RC-network simulating the thermodynamics of a building. Real data from the Nimbus Centre is then used to determine suitable zone models.

When considering the type of model to be used within an MPC strategy, the focus must be different than that of a model used for detailed simulation purposes. For control, particularly when numerical optimisation is employed (as is the case with standard MPC formulations (Maciejowski 2002, Camacho & Alba 2007)), the emphasis must be placed on simplicity. Though accuracy is inevitably compromised as the complexity is reduced, if appropriate models are used, the retained accuracy can still be sufficient to provide useful predictions of the future state of the building. A review of different methodologies for building energy modelling can be found in (Li & Wen 2014).

Broadly speaking, two approaches can be taken for development of lower order mod-

els which can be used for control purposes. The first method requires that a high order model is initially available. Model reduction techniques can be implemented to define low-order subspaces which can capture the dominant dynamics of the high order model. Examples of these approaches can be found in (Cigler & Privara 2010, Privara et al. 2010) and (Gouda et al. 2002). The second approach is to derive models using system identification methods with measured data taken from the building, such as in (Goyal et al. 2011)(Ríos-Moreno et al. 2007) and (Radecki & Hencsey 2013). This approach can be advantageous as a high order model is not required and little knowledge of the building structure needs to be taken into account. Black-box models (in which no knowledge of the system is incorporated), or grey-box models (which exploit a certain amount of knowledge of the building's thermal properties) can be used. Identifiability issues are noted however in (Agbi et al. 2012) and (Lin et al. 2012), in which the importance of well-excited data is emphasised. The impact of modelling error on the closed-loop control performance is analysed in (Bengea et al. 2011). Comparisons between some of these approaches can be found in (Afram & Janabi-Sharifi 2014)

Initially considering a scenario in which a complex physics-based model is available (such as that developed in the previous chapter), a model reduction technique is outlined for extracting low-order representations of high-order systems using Hankel singular value decomposition (HSVD). Using the test network developed in Chapter 3, HSVD is used to determine which system states dominate the dynamics of each individual zone. Discarding all other states, a low-order approximation is derived to represent the input/output relationship of each zone.

Data-driven methodologies are then investigated based on an ARMAX structure, such as those outlined in (Ríos-Moreno et al. 2007). Prediction-error methods are used to fit decentralised zone models to data taken from the test network. The impact of unmeasured disturbance is next considered. It is shown that the disturbance estimation and PCA-based spatial filtering techniques (introduced in Chapter 3) can be incorporated into the zone model identification process to derive models that better represent the dynamics of the system. The data-driven techniques are applied to measured data taken from the Nimbus Centre. It is shown that the data used must be of sufficient quality of the identification methods to be effective.

4.2 Model Suitability for Predictive Control Formulations

4.2.1 Decomposing the Full-Scale Building Model

In a typical heating strategy, thermal control of several zones is often carried out by a single system (this is also true of cooling systems, HVAC and fan coil units). In the case of the Nimbus Centre for example, all zones on the same floor are supplied by the same radiator flow-temperature. The full building envelope must be in some way considered in the control strategy. Crucially however, a fully centralised framework may not be suitable for a numerical optimisation-based control strategy, due to the complexity that may result. In a large multi-storey building the problem can quickly become too complex for on-line optimisation to be a realistic option, even if reduced-order approximations are used (Lazos et al. 2014). The system architecture must be considered.

Due to the physical divide naturally arising between separate rooms or zones, a block diagonal structure primarily governs the full-scale RC-network. This implies that a sensible approach may be to first consider the building as a set of decentralised or distributed zone models (Agbi & Krogh 2014). The complete control problem can then be considered as an optimisation across a set of smaller, less complex models. Furthermore, with a decentralised or distributed approach, faults are less likely to propagate through the control. The system architecture, as well as the grouping and partitioning of zones is considered in detail in (Chandan & Alleyne 2011, 2012, 2013) and (Chandan & Alleyne 2014).

Given the choice between decentralised and distributed framework, the complexity of both the control and model identification must be considered. Typically, a distributed strategy can achieve a control performance that approaches that of a centralised strategy, with a computational cost that is almost as low as the decentralised case (Moroşan et al. 2010). Identification of the thermodynamic coupling between the individual zone models is required in a distributed structure, in addition to communication between the resulting zone control problems. Alternatively, in a decentralised structure, the individual zone models can be taken to be thermodynamically independent with inputs from one zone not affecting the output of another. The respective merits of the two approaches have been studied in (Gouda et al. 2002).

4.2.2 Evaluating Performance of the Zone Models

It can be assumed that by reducing the complexity of the building models, a certain amount of accuracy will, in general, be lost. To determine what constitutes a sufficiently accurate low-order model, the purpose of the model must be considered. In this case, the objective is to determine a set of zone models which can be used within an MPC strategy. In such a strategy, a receding horizon approach is typically used, whereby model states are updated using measured outputs at each sample and the future states over a set horizon are predicted.

If the prediction horizon used is two hours, for example, the steady-state accuracy of the zone models over several days may not be relevant. A model output may provide a good representation of the zone dynamics over the horizon, despite drifting slowly over time. As such, using the output error as a quality metric will not be instructive. Conversely, if only the higher frequency dynamics of the system are accurately captured, the model may begin to drift before the end of the prediction horizon. In such a circumstance, the predicted behaviour of the plant towards the end of the horizon will be incorrect. This can significantly inhibit the controller performance. Depending on the class of model used, different methods can be used to determine more appropriate measures of frequency-relevant accuracy.

4.3 Physics-based Techniques for Identifying Reduced-Order Models

While intricate simulation models can provide a high level of accuracy, the complexity can be prohibitive when considering useful models for control purposes. Nonetheless, if a simulation model is available, lower order approximations can be extracted. Though the order of the complete RC-network may be high, the relationship between the inputs and outputs may be dominated by a far lower number of poles, particularly from the perspective of an individual zone.

In (Cole et al. 2013), a model developed using EnergyPlus is perturbed to produce data from which a linear reduced-order model is derived, while in (Hazyuk et al. 2012), a detailed simulation model is used to produce low order models, identified using least squares methods.

Examples in which model reduction techniques are used to derive reduced-order zone models from the resistance and capacitance parameters of a full scale RC-network can

be found in (Deng et al. 2014, Gouda et al. 2002, Palomo Del Barrio 2000). Such approaches allow for both the individual zone thermodynamics and the interactions between different zones to be characterised in a way that maintains a certain link to physical knowledge of the structure. An intensive modelling effort is required to produce such a network however, such as that outlined in the previous chapter.

4.3.1 Model Reduction - Hankel Singular Value Decomposition

If a high order simulation model of the building thermal dynamics is available, model reduction techniques may be implemented to reduce the order, while preserving the dominant dynamics. One such approach is illustrated here using the test network outlined in Chapter 3. This network model consists of 172 states resulting in a 172-order transfer function between the input of each zone and its output. Observing the structure of the network however, this input/output relationship will be largely unaffected by states associated with zones on the other side of the building (neglecting interaction through control of the heating system). For each zone, a lower-order approximation of the associated transfer function should be capable of capturing the most relevant dynamics in the system.

Hankel singular values can be defined as the square-root of the eigenvalues of the product of the controllability and observability gramians of a system. By linearly transforming the system so as to equate the gramians, the new model is said to be balanced, with the Hankel singular values corresponding to the terms of the main diagonal of the gramians. The states corresponding to the (relatively) smaller Hankel singular values can be thought of as *low energy* states. Removal of an appropriate number of these states can then result in a system of lower order in which the dominant dynamics remain. Such an approach is referred to as a balanced model-order reduction. The implementation of the strategy with a building model is outlined in (Ma et al. 2012).

Taking the equations describing the test network from (3.6), the observability and controllability gramians of the system, given by W_c and W_o , are found to satisfy the Lyapunov equations:

$$AW_c + W_cA^T = -BB^T \quad (4.1)$$

$$AW_o + W_oA^T = -CC^T \quad (4.2)$$

Using the transformation $\tilde{x} = Tx$, where $T \in \mathbb{R}^{172 \times 172}$ is invertible, a transformed

system could be defined as:

$$\dot{\tilde{\mathbf{x}}}(t) = \tilde{A}\tilde{\mathbf{x}}(t) + \tilde{B}\mathbf{u}(t) \quad (4.3)$$

$$\tilde{\mathbf{y}}(t) = \tilde{C}\tilde{\mathbf{x}}(t), \quad (4.4)$$

where

$$(\tilde{A}, \tilde{B}, \tilde{C}) = (TAT^{-1}, TB, CT^{-1}) \quad (4.5)$$

The gramians of the transformed system are given as:

$$\tilde{W}_c = TW_cT^T \quad (4.6)$$

$$\tilde{W}_o = (T^{-1})^T W_o T^{-1} \quad (4.7)$$

By appropriate selection of T , the system can be transformed to a balanced system, in which the transformed gramians (\tilde{W}_c and \tilde{W}_o) are equal and diagonal:

$$\tilde{W}_c = \tilde{W}_o = \text{diag}(\sigma_{hsv_1}, \sigma_{hsv_2} \dots \sigma_{hsv_{172}}), \quad (4.8)$$

which can be arranged so that $\sigma_{hsv_1} \geq \sigma_{hsv_2} \geq \dots \geq \sigma_{hsv_{172}} \geq 0$. In this form, the i^{th} Hankel singular value is then given as σ_{hsv_i} . Taking the transformed (balanced) system, states corresponding to an appropriately selected set of relatively small Hankel singular values can be removed from the system, as well as the associated parameters.

Where the set of retained states in the reduced balanced system is denoted $\tilde{\mathbf{x}}_1$, while the discarded states are denoted $\tilde{\mathbf{x}}_2$, the full-order balanced system of (4.3-4.4) could be expressed as:

$$\begin{bmatrix} \dot{\tilde{\mathbf{x}}}_1(t) \\ \dot{\tilde{\mathbf{x}}}_2(t) \end{bmatrix} = \begin{bmatrix} \tilde{A}_{1,1} & \tilde{A}_{1,2} \\ \tilde{A}_{2,1} & \tilde{A}_{2,2} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{x}}_1(t) \\ \tilde{\mathbf{x}}_2(t) \end{bmatrix} + \begin{bmatrix} \tilde{B}_1 \\ \tilde{B}_2 \end{bmatrix} \mathbf{u}(t) \quad (4.9)$$

$$\tilde{\mathbf{y}}(t) = \begin{bmatrix} \tilde{C}_1 & \tilde{C}_2 \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{x}}_1(t) \\ \tilde{\mathbf{x}}_2(t) \end{bmatrix} \quad (4.10)$$

The reduced system is then given as:

$$\dot{\hat{\mathbf{x}}}_1(t) = \tilde{A}_{1,1}\tilde{\mathbf{x}}_1(t) + \tilde{B}_1\mathbf{u}(t) \quad (4.11)$$

$$\hat{\mathbf{y}}(t) = \tilde{C}_1\tilde{\mathbf{x}}_1(t) \quad (4.12)$$

In Fig. 4.1, the first 40 Hankel singular values are shown (on a log scale) for the transfer

function between the heat input and zone temperature of Zone 1 of the 30 zone representative test network. These show the contribution of the 15 most dominant states in the transfer function. By viewing this, an informed decision may be made as to the number of states that should be retained in the truncated system. It can be seen for example, that the first element is approximately 350 times larger than the fifth element and will impact the system dynamics with an equivalent dominance.




Figure 4.1: 15 largest Hankel singular values for transfer function between the heat input and zone temperature output of Zone 1 of test network (plotted on log scale)

The transfer function between the heat supply and zone temperature for Zone 1, as well as the transfer function between the external and internal temperatures for the same zone is truncated for illustrative purposes to 1st, 2nd, 3rd and 4th order approximations. The bode plots of the different systems are compared in Fig. 4.2 and Fig. 4.3.




Figure 4.2: Frequency responses of full-order system and reduced-order approximations for transfer function between heat input and internal temperature for Zone 1 of the test network




Figure 4.3: Frequency responses of full-order system and reduced-order approximations for transfer function between external and internal temperatures for Zone 1 of the test network

It can be seen that for model-orders higher than second order, the truncated models represent the dynamics of the full-order system across a wide frequency spectrum. Directly truncating the model in this manner does not guarantee the preservation of the DC-gain however, which in certain circumstances may be required. If a long term projection of energy use is desired for example, a system approximation with an accurate DC-gain may provide more insight, at the expense of accuracy in a higher frequency band.

Preserving the DC-gain of the system described by (4.9)-(4.10) can be accomplished by setting $\dot{\tilde{x}}_2 = 0$, and solving for $\dot{\tilde{x}}_1$. The reduced system can be then expressed as:

$$\dot{\tilde{x}}_1(t) = [\tilde{A}_{1,1} - \tilde{A}_{1,2}\tilde{A}_{2,2}^{-1}\tilde{A}_{2,1}] \tilde{x}_1(t) + [\tilde{B}_1 - \tilde{A}_{1,2}\tilde{A}_{2,2}^{-1}\tilde{B}_2] u(t) \quad (4.13)$$

$$\hat{y}_{dc}(t) = [\tilde{C}_1 - \tilde{C}_2\tilde{A}_{2,2}^{-1}\tilde{A}_{2,1}] \tilde{x}_1(t) \quad (4.14)$$

In Fig. 4.4-Fig. 4.5 the frequency responses are shown for the same transfer functions Fig. 4.2-Fig. 4.3. In this case, the reduced order approximations are designed to maintain the DC-gain of the full-order system.

The bode plots show that, though the DC-gain is preserved, the higher frequency response of the reduced system may be quite inaccurate if the chosen order is too low. If the reduced models are to be used in an predictive control formulation, the length of prediction horizon must be considered when selecting the reduction method.

The ability to characterise the accuracy of the reduced order zone models (provided the high order model can be considered to be an accurate representation of the actual




Figure 4.4: Frequency responses of full-order system and reduced-order approximations for transfer function between heat input and internal temperature for Zone 1 of the test network - preserving DC-gain




Figure 4.5: Frequency responses of full-order system and reduced-order approximations for transfer function between external and internal temperatures for Zone 1 of the test network - preserving DC gain

building's dynamics) is a significant advantage of this type of approach. Furthermore, reduced order transfer functions of the heat flow between zone temperatures and the inputs of neighbouring zones can be developed, if a centralised or distributed control strategy is desired.

4.4 Deriving Zone Models from Measured Data

A significant disadvantage to the model reduction approach is the requirement of a full-scale high-order model to begin with. Using knowledge of physical characteris-

tics of a building allows for certain advantages such as thermal passivity as described in previous chapters. Developing an initial high-fidelity representation of a building requires an extensive commissioning effort, specific to each building. For implementation of a strategy in several different buildings, it should not be necessary to produce a high-order simulation model for each. While a simulation platform is a useful tool for controller design and analysis, once the strategy has been designed, given the reduction in required modelling effort and a priori knowledge of the system, a data-based modelling approach would appear to be a more attractive prospect for widespread implementation (Ghiaus 2006, Berthou et al. 2014, Ma et al. 2014, Freire et al. 2008). System identification methodologies which allow this are outlined in (Ljung 1999).

Black-box and grey-box data driven techniques for building modelling are described in (Berthou et al. 2014, Ríos-Moreno et al. 2007, Ferkl et al. 2010). These models can be developed for individual zones at orders low enough for efficient optimisation. A drawback to using these approaches is that input and output data from the plant, rich enough for identification purposes is required, as is referred to in (Privara et al. 2010, Rojas et al. 2008).

4.4.1 System-Identification Methodology

A possible approach for relating the measured inputs and outputs of each zone is by use of an ARX structure. Taking the heat supply to the zone and the external temperature as the inputs denoted by u_1 and u_2 respectively with the internal zone temperature measurements taken as the output y , a linear difference equation using the previous n_o outputs and n_i inputs can be developed to represent a predictor for the current estimated output $\hat{y}(k)$ as follows:

$$\hat{y}(k|\boldsymbol{\vartheta}_{ar}) = \boldsymbol{\varphi}_{ar}^T(k)\boldsymbol{\vartheta}_{ar}, \quad (4.15)$$

where

$$\boldsymbol{\vartheta}_{ar} = [a_1, \dots, a_{n_o}, b_1, \dots, b_{n_i}, b_{n_i+1}, \dots, b_{2n_i}]^T \quad (4.16)$$

$$\boldsymbol{\varphi}_{ar}(k) = [-y(k-1), \dots, -y(k-n_o), u_1(k-1), \dots, u_1(k-n_i), \\ u_2(k-1), \dots, u_2(k-n_i)]^T \quad (4.17)$$

Using a Least-Squares (LS) approach, the squared error between the measured outputs and the estimated model outputs over a batch of N_{ls} samples, the parameter values can

be computed. Where the model residual e at sample i is given by $e(i) = (y(i) - \hat{y}(i))$:

$$\hat{\boldsymbol{\vartheta}}_{ar} = \arg \min_{\boldsymbol{\vartheta}_{ar}} \sum_{i=i}^{N_{ls}} e(i)^2 \quad (4.18)$$

This structure implicitly assumes that the residual has a constant power spectral density (white noise). Given the large amount of unmeasured non-zero-mean disturbances acting on a building (as outlined in chapter 3), the assumption of white residuals will typically result in a biased model. To achieve greater flexibility in the description of the disturbance term, an ARMAX structure could instead be adopted. A comparative analysis between the two structures for use in building modelling can be seen in (Ríos-Moreno et al. 2007).

In an ARMAX structure, the residual is assumed to be a moving average white-noise term. The dynamics of this residual term must also be identified, which is not possible using the LS approach of the ARX model. By minimising the prediction accuracy as opposed to just model mismatch, Prediction Error Methods (PEM), as outlined in (Ljung 2002), can be used to identify discrete-time models suitable for use in model predictive control strategies (in which prediction accuracy is paramount).

While for an ARX model, the prediction error and model mismatch are equivalent, in an ARMAX model, the predicted output will depend on n_e past values of the residual. Using past values of the inputs and outputs, a predictor for the current output can be formed in terms of the parameter vector $\boldsymbol{\vartheta}_{arm}$. A general ARMAX model can be expressed as:

$$\hat{y}(k|\boldsymbol{\vartheta}_{arm}) = \boldsymbol{\varphi}_{arm}^T(k, \boldsymbol{\vartheta}_{arm})\boldsymbol{\vartheta}_{arm}, \quad (4.19)$$

where

$$\boldsymbol{\vartheta}_{arm} = [\boldsymbol{\vartheta}_{ar}^T, c_1, \dots, c_{n_e}]^T \quad (4.20)$$

$$\boldsymbol{\varphi}_{arm}(k, \boldsymbol{\vartheta}_{arm}) = [\boldsymbol{\varphi}_{ar}^T, \epsilon_{arm}(k-1, \boldsymbol{\vartheta}_{arm}), \dots, \epsilon_{arm}(k-n_e, \boldsymbol{\vartheta}_{arm})]^T \quad (4.21)$$

$$\epsilon_{arm}(k, \boldsymbol{\vartheta}_{arm}) = y(k) - \hat{y}(k, \boldsymbol{\vartheta}_{arm}) \quad (4.22)$$

By selecting a suitable norm, the difference between the predicted outputs and the measured outputs can be in some way minimised. In the zone models developed in this chapter, a quadratic norm is used. For a batch of N_{pem} samples, the cost function

is:

$$\hat{\boldsymbol{\vartheta}}_{arm} = \arg \min_{\boldsymbol{\vartheta}_{arm}} \sum_{i=1}^{N_{pem}} \epsilon_{arm}^2(i, \boldsymbol{\vartheta}_{arm}) \quad (4.23)$$

The minimisation of a cost function of this form can be referred to as a nonlinear least-squares problem and is solved here using a Gauss-Newton algorithm. This is a type of iterative search whereby, if at the q^{th} iteration, the estimated parameter vector is given as $\boldsymbol{\vartheta}_{arm}^q$, the $(q + 1)^{th}$ estimate is given by:

$$\boldsymbol{\vartheta}_{arm}^{q+1} = \boldsymbol{\vartheta}_{arm}^q + \left(J_{arm}^T J_{arm} \right)^{-1} J_{arm}^T \begin{bmatrix} \epsilon(1, \boldsymbol{\vartheta}_{arm}^q) \\ \vdots \\ \epsilon(N_{pem}, \boldsymbol{\vartheta}_{arm}^q) \end{bmatrix} \quad (4.24)$$

J_{arm} is the Jacobian of $f_{GN}(\boldsymbol{\vartheta}_{arm}^q)$, where:

$$f_{GN}(i, \boldsymbol{\vartheta}_{arm}^q) = [\boldsymbol{\varphi}_{arm}^q(i, \boldsymbol{\vartheta}_{arm}^q)]^T \boldsymbol{\vartheta}_{arm}^q \quad (4.25)$$

$$(J_{arm})_{i,j} = \frac{\partial f_{GN}(i, \boldsymbol{\vartheta}_{arm}^q)}{\partial \vartheta_{arm,j}^q} \quad (4.26)$$

$$\forall i \in 1 \cdots N_{pem}, \quad \forall j \in 1 \cdots |\boldsymbol{\vartheta}_{arm}^q|$$

Once defined, both the ARX and ARMAX structures can be converted to linear state-space models. Depending on the purpose of the models, different approaches can be taken for identifying the free ARMAX parameters. It is noted that a quadratic norm is used in this chapter for identification purposes. To focus the identification process on time-scales that correspond to the relevant dynamics of the processes being modelled, as well as to allow for different norms to be used, the cost for different numbers of steps ahead can be minimised. It was found here that when shorter horizons were used, certain longer time-scale dynamics were not captured adequately (the choice of typical time-scales associated with building energy control is expanded on in (Touretzky & Baldea 2013)).

4.4.2 Identifiability - Limitations Due to Excitation

In the previous chapter, the structural and data-dependent identifiability issues were highlighted for training a large scale RC-network to fit data. For low order zone models, given the greatly reduced number of parameters, structural identifiability becomes less of an issue. Adequate excitation in the data-sets used for identification remains

vital to the validity of the solution however, as outlined in (Lin et al. 2012).

Achieving this level of excitation in the building domain may not be trivial. As the external temperature cannot be controlled (and will typically change slowly), the range of excitation is limited. Furthermore, under standard operation, the heat transferred between neighbouring zones through well insulated internal walls will be small relative to the heat input from the heating system and heat lost to the environment. Using data-based approaches to determine the link between inputs and outputs from different zones would require neighbouring zones to be driven to divergent temperatures. Clearly, the building would have to be unoccupied for an extended period of time for this to be achieved.

In the data-driven strategies outlined here, to allow for scenarios in which invasive functional tests are not possible, the heat transfer between zones through internal walls is neglected. The resulting decentralised control strategy can then regard the zones as being thermally independent, only connected through the similar inputs of the heating system.

While the decision to use such a decentralised approach will result in sub-optimal control decisions in each zone, it can be seen to be a practical decision. It is noted in (Chinde et al. 2015) that when data-driven approaches are used to develop zone models, artificial correlations between measurements from different zones that are inevitable in the operation of a building can result in over-fitting, negatively affecting the accuracy of the derived models. The extensive functional testing effort required to avoid these issues may not be desirable, outweighing the benefits that could result. In (Wu & Sun 2012) for example, in which physics-based ARMAX zone models are developed, the authors observed that the difference between the external temperature and room temperature was on average 2200% larger than the temperature difference between any two rooms during experimentation. As a result, the heat convection between a room and the outside air dominated the thermal interaction between a room and its surroundings. Determining accurate representations of these zone interactions may then not be particularly advantageous.

4.4.3 Models Derived Using Undisturbed Data from the Test Network

In Section 3.3, a five day disturbance-free training data-set for use with the 30 zone test network was outlined. Using this data and the PEM techniques outlined above, a

set of low-order, decentralised ARMAX models are derived and transformed to a state-space format for each of the individual zones in the system. The sample-time for the training data is now increased to 300 seconds, to reflect the data quality issues observed when using real data from the Nimbus Centre. As before, the heat to the zones and the external temperature are the inputs, while the zone temperatures are the outputs to each of the systems. The zone models are considered thermally independent, with no heat transfer through the internal walls assumed. Initially, no unmeasured disturbances are included, so as to assess the ability of low-order models to capture the relevant system dynamics and to determine suitable model orders.

In Fig. 4.6 and Fig. 4.7, bode plots can be seen for 1^{st} – 4^{th} order approximations of the thermal dynamics of Zone 1. Fig. 4.6 shows the relationship between the heat input (in Watts) and the internal temperature, while in Fig. 4.7, the external temperature is taken as the input. It can be seen that the low frequency response of the 1^{st} and 2^{nd} order systems is quite different to the full-order network.




Figure 4.6: Frequency responses of full-order system and data-driven low-order transfer functions between heat input and internal temperature for Zone 1 of the test network

To observe the impact of the previously outlined identifiability issues, a second set of models was developed in which the heat supply input for the training data was replaced by a pseudo-random binary sequence (PRBS) shaped input. Such an input would not be feasible in reality (due to the dynamics of the heating system), but can be used as a comparative study. The bode plots for the same zone using the models derived with this set of inputs are shown in Fig. 4.8 and Fig. 4.9. It can be seen that the responses are now closer to the response of the full-order system, though not significantly.

In Section 4.2.2, the difficulty in comparing the simulated output of a model to the target output for these control models is highlighted. The ARMAX model structure




Figure 4.7: Frequency responses of full-order system and data-driven low-order transfer functions between external and internal temperatures for Zone 1 of the test network




Figure 4.8: Frequency responses of full-order system and data-driven low-order transfer functions between heat input and internal temperature for Zone 1 of the test network using PRBS input for training

used here allows for a p -step ahead model prediction to be made for the output, which can provide a more useful comparison than the simulated output. The simulated output (\hat{y}_{sim}) of a model at sample k is only a function of the previous inputs:

$$\hat{y}_{sim}(k) = f(u(k-1), \dots, u(1)) \quad (4.27)$$

A p -step ahead prediction (\hat{y}_p) of the output at time $k+p$ on the other hand also incorporates the residuals (which are just functions of the measured outputs) up to k :

$$\hat{y}_p(k+p|k) = f(u(k+p-1), \dots, u(1), y(k), \dots, y(1)) \quad (4.28)$$




Figure 4.9: Frequency responses of full-order system and data-driven low-order transfer functions between external and internal temperature for Zone 1 of the test network using PRBS input for training

This predicted output can provide a more instructive indication of the accuracy of predictions taken from the model over a specific time-scale, which is appropriate, given the MPC-focused purpose of the models. The accuracy of a prediction several days into the future for example, may not be relevant if the MPC prediction horizon for which the models are required is two hours. For the models developed in this chapter, a 6-hour prediction (72 steps ahead) is used. At this time-scale, the mid-frequency dynamics are most important. To measure how well the predicted model output *fits* that of the target system, the coefficient of determination (R^2) is used as a metric, given for a set of N_{val} samples as (Kutner et al. 1996):

$$R^2 = 1 - \frac{\sum_{i=1}^{N_{val}} (y(i) - \hat{y}_p(i))^2}{\sum_{i=1}^{N_{val}} (y(i) - \bar{y})^2}, \quad (4.29)$$

where \bar{y} is the sample mean of the measured data.

Using an unseen validation data-set simulating a period of 9 days, the 72-step ahead predicted output of the 1st–4th order model approximations derived using the standard set of inputs is shown along with the target output data for Zone 1 in Fig. 4.10.

In Table 4.1, the R^2 values of the models derived using both the standard input and the PRBS input are shown. It can be seen that a high prediction accuracy is achieved for all models in which the order is greater than 1. Additionally, it illustrates that, although low frequency error could be seen in the bode plots of the 2nd order models, the prediction accuracy is still high. With the exception of the 1st order case, the models derived using the PRBS input do not show an improvement.




Figure 4.10: p -step (6 hours) ahead predicted output for $1^{st} - 4^{th}$ order models of Zone 1

Table 4.1: Coefficient of determination for p -step ahead predicted outputs of $1^{st} - 4^{th}$ order models of Zone 1

COEFFICIENT OF DETERMINATION (DISTURBANCE-FREE)				
	1^{st} Order	2^{nd} Order	3^{rd} Order	4^{th} Order
Standard Input (R^2)	0.7654	0.9790	0.9979	0.9982
PRBS Input (R^2)	0.8423	0.9631	0.9888	0.9914

4.5 Including Unmeasured Disturbances in the Identification Data

As was the case for the simulation model identification in Chapter 3, unmeasured disturbances can significantly impact any attempt to derive models from input and output data (Maasoumy et al. 2014). To examine this impact, the PEM-based zone model identification process is repeated once again, this time in the presence of the solar gain and internal gain profiles applied to the 30 zone test network as in Section 3.6.1. As before, it is assumed that no measurement is available for these disturbances.

4.5.1 Identification without Disturbance Estimation

To assess the impact of the disturbances on the models without any disturbance estimation, $1^{st} - 4^{th}$ order models for each zone are first developed using only measured inputs (standard inputs - not PRBS) and outputs. In Fig. 4.11 and Fig. 4.12 the bode

plots associated with the first zone are shown (using the heat supply as the input in Fig. 4.11 and the external temperature as the input in Fig. 4.12 with the zone temperature used as the output for both). It can be seen that the dynamics are less accurately captured than in the undisturbed case, particularly for the higher order approximations.




Figure 4.11: Frequency responses of full-order system and data-driven low-order transfer functions between heat input and internal temperature for Zone 1 of the test network derived with disturbances present




Figure 4.12: Frequency responses of full-order system and data-driven low-order transfer functions between external and internal temperatures for Zone 1 of the test network derived with disturbances present

To better quantify the accuracy, the p -step ahead predicted output of the zone models is compared to the target output for the set of unseen, undisturbed validation data (again, using a 72-step ahead prediction to emphasise the mid-frequency range). This is shown in Fig. 4.13, while in Table 4.2, the R^2 values for these predicted outputs is given. It can be seen that the accuracy of the predicted output has been significantly affected for all model orders.




Figure 4.13: p -step (6 hours) ahead predicted output for $1^{st} - 4^{th}$ order models of Zone 1 (zone models derived in presence of disturbance)

Table 4.2: Coefficient of determination for p -step ahead predicted outputs of $1^{st} - 4^{th}$ order models of Zone 1 (zone models derived in presence of disturbance)

COEFFICIENT OF DETERMINATION (WITH UNMEASURED DISTURBANCES PRESENT)				
	1^{st} Order	2^{nd} Order	3^{rd} Order	4^{th} Order
R^2	-0.8887	0.5899	0.7924	0.6614

The reduction in accuracy seen from Table 4.1 to Table 4.2 indicate that the lower accuracies seen in Table 4.2 are a result of the disturbances, rather than the model order. Though the third-order results in the presence of disturbance are most accurate for this zone, this was not true of all zones. The undisturbed accuracy of the second-order models imply that there may not be a significant benefit to using higher-order approximations, particularly given the increased complexity that would be introduced. As such, second-order models are chosen to represent each of the individual zones.

4.5.2 Inclusion of the Spatio-Temporally Filtered Disturbance Estimates in the Identification Process

In Section 3.7, a method for refining model parameter estimates using disturbance estimates was developed. The same PCA-based spatial filtering technique is adapted for use here to adjust the parameters of the zone models, this time incorporated into the PEM identification formulation.

In place of the centralised Kalman filter used for the full platform, individual decentralised Kalman filters are first used to determine the disturbance estimates of each zone. As the disturbances are predominantly heat based, it is assumed that they affect the model states with the same dynamics as the radiator heat inputs. The structure of the discrete-time Kalman Filter for each zone j of the building is as follows, where the initial disturbance estimate associated with the zone is denoted \hat{x}_{d_j} :

$$\begin{bmatrix} \hat{\mathbf{x}}_j(k+1) \\ \hat{x}_{d_j}(k+1) \end{bmatrix} = \begin{bmatrix} A_j & B_j \\ 0_{1 \times 2} & 1 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}}_j(k) \\ \hat{x}_{d_j}(k) \end{bmatrix} + \begin{bmatrix} B_j & E_j \\ 0 & 0 \end{bmatrix} \begin{bmatrix} u_j(k) \\ T_{amb}(k) \end{bmatrix} + K_j (y_j(k) - \hat{y}_j(k)) \quad (4.30)$$

$$\hat{y}_j(k) = \begin{bmatrix} C_j & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}}_j(k) \\ \hat{x}_{d_j}(k) \end{bmatrix} \quad (4.31)$$

where $\hat{\mathbf{x}}_j \in \mathbb{R}^2$, u_j , and y_j denote the states, heat input and room temperature of the j^{th} zone respectively, while T_{amb} represents the external temperature (a measured disturbance). A_j , B_j , C_j and E_j represent the second-order model parameter matrices. The initial disturbance estimates from all the zones of the building are then grouped in north, south, east and west facing categories, as well as a central set of zones, as described in Chapter 3. Using the PCA process described in (3.29)-(3.36), a spatially filtered disturbance estimate for each zone, given by \hat{x}'_{d_j} is determined. The number of principal components is chosen to rebuild the disturbance estimates in each of the groups by observing the spectrum given by (3.40).

These disturbance estimates are included as inputs to the PEM identification procedure for each zone and the model parameters are updated. Where $\boldsymbol{\vartheta}_j$ is the set of model parameters, the structure of the model for the j^{th} zone can be expressed as:

$$\begin{aligned} \mathbf{x}_j(k+1) &= A(\hat{\boldsymbol{\vartheta}}_j)\mathbf{x}_j(k) + B(\hat{\boldsymbol{\vartheta}}_j) \left[u_j(k) + \hat{x}'_{d_j}(k) \right] + E(\hat{\boldsymbol{\vartheta}}_j)T_{amb}(k) \\ &\quad + \Gamma_{res}(\hat{\boldsymbol{\vartheta}}_j)\epsilon_j(k) \end{aligned} \quad (4.32)$$

$$y_j(k) = C(\hat{\boldsymbol{\vartheta}}_j)\mathbf{x}_j(k), \quad (4.33)$$

where ϵ_j represents the residual. The Kalman filter model of (4.30-4.31) is updated using the new model parameters and a new disturbance estimate is determined. The spatial filtering procedure is then repeated. The model parameters and disturbance estimates are iteratively refined until the model converges, by which point, no change occurs between updates. The full process of parameter identification and disturbance estimation is summarized in Fig. 4.14.

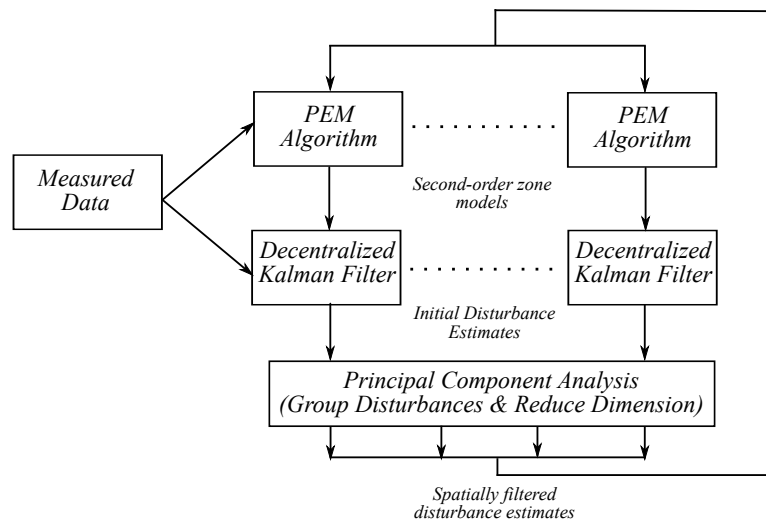


Figure 4.14: Iterative process of parameter identification and disturbance estimation

4.5.3 Models Derived Using Disturbed Test Network Data with Disturbance Estimation

Using the five day training data-set with unmeasured disturbances included, second order zone models are identified using the PEM formulation allied to the PCA disturbance estimation technique. The sum of the squared residuals is used as the cost function of the PEM formulation. To tune the Kalman filter, the variance of the disturbance-state error was assumed to be far greater than both the model state-error and the sensor noise. The off-diagonal terms of the model-error and sensor-error covariance matrices are set at 0. In the PCA process, the same zone groupings are used as in Section 3.7.4, with only the primary principal component used to rebuild the disturbance estimates in each group.

In Fig. 4.15 and Fig. 4.16, the bode plots associated with Zone 1 are shown for the full order system, the second-order system derived using spatially filtered disturbance estimates, the second-order system derived without disturbance estimation and, as a performance bound, the second-order system derived when no disturbances were present. As before, in the first plot, the heat supply is taken as the input while in the second, the external temperature is used. The internal zone temperature is used as the output for each. It can be seen that by including the disturbance estimates, the resulting model produces a better representation of the system in the low-mid frequency range, approaching the accuracy obtained when disturbance-free data was used for identification.

In Fig. 4.17, the 72-step ahead predicted output for the models derived with and with-




Figure 4.15: Frequency responses of full-order system and data-driven low-order transfer functions between heat input and internal temperature for Zone 1 of the test network derived with and without disturbance estimation




Figure 4.16: Frequency responses of full-order system and data-driven low-order transfer functions between external and internal temperatures for Zone 1 of the test network derived with and without disturbance estimation

out disturbance estimation are shown for an unseen validation data-set (simulating a nine day period). Table 4.3 shows the R^2 values obtained for this data-set using the undisturbed data, the disturbed data without estimation and the disturbed data with estimation. It can be seen that the model derived in the presence of unmeasured disturbances can achieve an accuracy approaching that of the undisturbed scenario if the disturbance estimation strategy is incorporated.

While a significant improvement can be achieved by estimating the unmeasured disturbance, the need for spatial filtering is clarified by observing the set of models derived without PCA. For certain zones, incorporating the disturbance estimates produced by the Kalman filters directly in the PEM formulation without using PCA can result in a




Figure 4.17: 72-step (6 hours) ahead predicted output for 2^{nd} order models of Zone 1 derived with and without disturbance estimation

Table 4.3: Coefficient of determination for 72-step ahead predicted outputs of 2^{nd} order models of Zone 1 derived without disturbances present, with disturbances present and with disturbances present and estimated

COEFFICIENT OF DETERMINATION (WITH & WITHOUT DISTURBANCE ESTIMATION)			
	Disturbance-Free	No Disturbance Estimate	PCA Disturbance Estimate
2^{nd} Order (R^2)	0.9790	0.5899	0.9325

high level of accuracy. This is not true for all zones however.

The initial Kalman filtered disturbance estimates will not only be a result of unmeasured disturbance, but also of output-error caused by model inaccuracy. As a large amount of the output-error can then appear to be accounted for by the disturbance estimate, the PEM formulation can only seek to minimise what remains of the output-error, which in certain zones, results in convergence to a biased model.

To highlight this, the 72-step ahead predicted output using the validation data-set for the models derived for Zone 4 with and without the use of PCA are shown in Fig. 4.18, as well as the model derived without any disturbance estimate. The R^2 values for these predicted outputs are also shown and compared to the value obtained when disturbance-free data is used.

It can be seen that to improve the identification process, the output-error portion of the disturbance estimate must first be separated from the remainder of the disturbance estimate and removed, in this case by using PCA.




Figure 4.18: 72-step (6 hours) ahead predicted output for 2^{nd} order models of Zone 4 derived with and without PCA

Table 4.4: Coefficient of determination for 72-step ahead predicted outputs of 2^{nd} order models of Zone 4

COEFFICIENT OF DETERMINATION (WITH & WITHOUT DISTURBANCE ESTIMATION)				
	Disturbance- Free	No Disturbance Estimate	KF only Estimate	KF-PCA Estimate
2^{nd} Order (R^2)	0.9387	0.6695	0.7578	0.9241

4.6 Deriving Zone Models Using Real Data - The Nimbus Centre

The techniques developed in Section 4.5 can be used to develop zone models appropriate for control purposes in a scenario in which a high order simulation platform is not available. Using the measured data taken from the Nimbus Centre (as outlined in Section 3.9.1), these methods are used to develop zone models which can represent the thermal dynamics of the zones in the Nimbus Centre.

4.6.1 Nimbus Centre Inputs, Outputs and Disturbances

The heating system of the Nimbus Centre must again be considered at this point, in this case taking into account which inputs are to be used as control variables. When using the test network, it is possible to use the heat supplied to each zone in *Watts* as the primary input. In the real heating system however, the heat flow into the zones is dependent on the flow-temperature of the water in the radiator (T_{fl}), the zone tem-

perature (T_z), the mass flow rate of the water through the radiator (\dot{m}_{fl}) and the heat transfer coefficient associated with the radiator (UA). Of these, T_{fl} and \dot{m}_{fl} could be regarded as the control variables. For the j^{th} zone using (3.46), if the heat transfer to surrounding zones is neglected, the zone model could be expressed as:

$$\dot{x}_j(t) = A_{c_j}x_j(t) + B_{c_j}\frac{2UA_j\dot{m}_{fl_j}(t)C_p}{2\dot{m}_{fl_j}(t)C_p + UA_j}(T_{fl}(t) - C_{c_j}x_j(t)) + E_{c_j}T_{amb}(t) \quad (4.34)$$

$$T_{z_j}(t) = C_{c_j}x_j(t) \quad (4.35)$$

where A_c , B_c , C_c and E_c are the continuous-time model parameter matrices. To maintain linearity, as discussed in Section 3.10.4, the mass flow rate and the heat transfer coefficient are taken to be constant. A linear discretized zone-model can then be formed using (4.34) as follows:

$$\dot{x}_j(k) = A_jx_j(k) + B_j\delta_j(T_{fl}(k) - C_jx_j(k)) + E_jT_{amb}(k) \quad (4.36)$$

in which δ_j is a constant. This equation can be rearranged as:

$$\dot{x}_j(k) = (A_j - B_j\delta_jC_j)x_j(k) + B_j\delta_jT_{fl}(k) + E_jT_{amb}(k), \quad (4.37)$$

which can then be defined as:

$$\dot{x}_j(k) = \tilde{A}_jx_j(k) + \tilde{B}_jT_{fl}(k) + E_jT_{amb}(k) \quad (4.38)$$

The control input to the zone model is now T_{fl} (which will be common to all zones on the same floor), T_{amb} is a measured disturbance (common to all zones) and the model matrices are given as \tilde{A}_j , \tilde{B}_j , C_j and E_j .

To estimate the unmeasured disturbances, the Kalman filter model of (4.30) - (4.31) is modified and given as:

$$\begin{bmatrix} \hat{\mathbf{x}}_j(k+1) \\ \hat{x}_{d_j}(k+1) \end{bmatrix} = \begin{bmatrix} \tilde{A}_j & \frac{\tilde{B}_j}{\delta_j} \\ 0_{1 \times 2} & 1 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}}_j(k) \\ \hat{x}_{d_j}(k) \end{bmatrix} + \begin{bmatrix} \tilde{B}_j & E_j \\ 0 & 0 \end{bmatrix} \begin{bmatrix} T_{fl}(k) \\ T_{amb}(k) \end{bmatrix} + K_j(y_j(k) - \hat{y}_j(k)) \quad (4.39)$$

$$\hat{y}_j(k) = \begin{bmatrix} C_j & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}}_j(k) \\ \hat{x}_{d_j}(k) \end{bmatrix} \quad (4.40)$$

The structure of the model used for identification is then given as:

$$\begin{aligned} \mathbf{x}_j(k+1) = & \tilde{A}(\hat{\boldsymbol{\vartheta}}_j)\mathbf{x}_j(k) + \tilde{B}(\hat{\boldsymbol{\vartheta}}_j) \left[T_{fl}(k) + \frac{\hat{x}'_{d_j}(k)}{\delta_j} \right] + E(\hat{\boldsymbol{\vartheta}}_j)T_{amb}(k) \\ & + \Gamma_{res}(\hat{\boldsymbol{\vartheta}}_j)\epsilon_j(k) \end{aligned} \quad (4.41)$$

$$y_j(k) = C(\hat{\boldsymbol{\vartheta}}_j)\mathbf{x}_j(k), \quad (4.42)$$

4.6.2 Zone Models Derived from Measured Nimbus Centre Data

With a training period of 15 days, input (T_{fl} & T_{amb}) and output data (T_z) taken from the Nimbus Centre was used to develop a set of 2^{nd} -order models for each of the zones of the building. The unmeasured disturbances were estimated and filtered using the same zone groupings and number of principal components as for the simulation platform development (Section 3.9.3). North, south, east and west facing zone groups were used as well as an additional central set. Using (3.40), one principal component was chosen to rebuild the disturbance estimates in each group.

Evaluating the resulting zone models is not as straightforward as for the test network, as only disturbed data can be retrieved from the building. A comparison between the zone model outputs and the actual outputs is not then informative - even if the output of the zone model with disturbance estimate included closely matches the measured output, an incorrect disturbance estimate may be concealing an inaccurate zone model. To determine the accuracy of the zone model dynamics without the distorting impact of the disturbances, the dynamics of the zone models, derived using measured data, were compared instead to the dynamics of the high-order simulation platform developed in Chapter 3.

It should be noted that this is an imperfect comparison used for illustrative purposes. The validity of the comparison relies on the assumed accuracy of the simulation platform, which itself cannot be directly compared to the real building without disturbances being present. In Fig. 4.19 a bode plot shows the dynamics between the input T_{fl} and the output T_z of a first-floor office taken from the simulation platform, the 2^{nd} -order zone model derived without disturbance estimation and the 2^{nd} -order model derived using a spatially filtered disturbance estimate. An improvement can be seen in the mid-frequency dynamics of the model derived using the disturbance estimation.

In Fig. 4.20, undisturbed data taken from the simulation platform is used to determine a 72-step (6 hour) ahead predicted output from each of the models for a period of 10 days. The R^2 values associated with each of the 2^{nd} -order models are shown in




Figure 4.19: Bode plot comparison of simulation platform and identified second-order control model (first-floor office)




Figure 4.20: p -step (6 hours) ahead predicted output for 2^{nd} order models of first-floor office

Table 4.5.

The importance of data quality was emphasised earlier in the chapter. Using measured data from the Nimbus Centre to derive the zone models, the importance is clarified further. In one south-facing zone of the building (the first floor seminar room), in which large amounts of glazing were present, a particularly low input excitation was observed for the training period. This appeared to be caused by solar gains pushing the room temperature above the comfort set-point, resulting in the thermostatic valve remaining closed for long periods. With the data obtained, it was not possible to obtain accurate models. Whereas for the simulation platform development, the initial physics-based estimate of the RC-network allowed for candidate solutions to be confined to a physically sensible search-space, the solution to a deterministic PEM problem is entirely dependent on the quality of measured data used.

Table 4.5: Coefficient of determination for p -step ahead predicted outputs of 2^{nd} order models of first-floor office derived using data from Nimbus Centre with and without disturbance estimation

COEFFICIENT OF DETERMINATION (THE NIMBUS CENTRE)		
	No Disturbance Estimate	PCA Disturbance Estimate
2^{nd} Order (R^2)	0.7811	0.9293

In Fig. 4.21, the bode plot for this poorly excited seminar room (with T_{fl} and T_z as the input and output respectively), derived using measured data shows the performance of the 2^{nd} -order model with and without disturbance estimation. Fig. 4.22 shows the 6-hour ahead predicted output for the models with the R^2 value given in Table 4.6. Allowing for the inevitable differences that will exist between the building from which data was obtained and the simulation platform to which the resulting model is being compared, this level of accuracy would clearly be of little use in a predictive control strategy.




Figure 4.21: Bode plot comparison of simulation platform and identified second-order control model (first floor seminar room)

As a comparison, in Fig. 4.23 the 6-hour ahead predicted output for a second order model derived using highly-excited data taken directly from the simulation platform is shown. This model has an R^2 value of 0.9757. While it is possible to obtain zone models from data without the availability of a high order simulation model using the techniques outlined in this chapter, to obtain accurate models, a great deal of care must be taken when using measured data for identification, particularly if functional tests are not carried out.




Figure 4.22: p -step (6 hours) ahead predicted output for 2^{nd} order models of first-floor seminar room (models derived using poorly excited data)

Table 4.6: Coefficient of determination for p -step ahead predicted outputs of 2^{nd} order models of first-floor seminar room derived using data from Nimbus Centre with and without disturbance estimation

COEFFICIENT OF DETERMINATION (THE NIMBUS CENTRE)		
	No Disturbance Estimate	PCA Disturbance Estimate
2^{nd} Order (R^2)	-10.5883	0.1827

4.7 Conclusion

This chapter analyses different approaches which can be used to obtain low-order building models for use within a numerical-optimisation based control formulation such as MPC. Broadly speaking, the approaches can be categorized as physics-based or data-driven.

The type of models suitable for use within an MPC framework were first considered. A compromise is required between complexity and accuracy to allow for reasonable computation times with useful predictions. With this in mind, the full building model was decomposed into a set of zone models which can be considered as thermodynamically isolated in a decentralised framework for simplicity.

To replicate a scenario whereby a high-fidelity simulation model is available, a methodology was outlined for developing reduced order zone models using Hankel singular value decomposition. As the parameters of the complex model will typically be in some way linked to physical building properties (as in Chapter 3), this identification




Figure 4.23: p -step (6 hours) ahead predicted output for 2^{nd} order models of first-floor seminar room derived using data from high order simulation platform

procedure is classed here as a physics-based approach. Reduced order zone models were developed to represent the thermal dynamics of a high order RC-network. The accuracy of the resulting models was compared for different levels of complexity. It can be seen that the $2^{nd} - 4^{th}$ order models produced frequency responses similar to those of the full-order model.

For scenarios in which high-order simulation models are not available, the model reduction techniques cannot be used. Data-driven approaches based on ARX and ARMAX model structures were outlined. Using input and output data taken from the test network, prediction-error methods were used to derive ARMAX structured zone models. Using disturbance-free data, accurate models were obtained, assessed in terms of the p -step ahead prediction accuracy. By applying unmeasured disturbances to the test network during data-acquisition, it was shown that the input/output relationship can become obfuscated in the data, resulting in inaccurate zone models. A methodology is introduced by which the disturbance estimation and PCA-based spatial filtering techniques developed in Chapter 3 can be incorporated in the PEM formulation to allow for the disturbances and model parameters to be simultaneously identified.

By these methods, 2^{nd} -order zone models were derived from input/output data corrupted by unmeasured disturbances which approach the prediction accuracy of the disturbance-free case. The PEM/PCA methodology was then applied to the real data taken from the Nimbus Centre. It was shown that in certain cases in which the input data was not sufficiently excited, the resulting zone model could not provide an accurate prediction of a zone's output.

References

- Afram, A. & Janabi-Sharifi, F. (2014), ‘Review of modeling methods for HVAC systems’, *Applied Thermal Engineering* **67**(1-2), 507–519.
- Agbi, C. & Krogh, B. (2014), ‘Decentralized identification of building models’, *2014 American Control Conference* pp. 1070–1075.
- Agbi, C., Song, Z. & Krogh, B. (2012), ‘Parameter identifiability for multi-zone building models’, *Proceedings of the IEEE Conference on Decision and Control* pp. 6951–6956.
- Bengea, S., Adetola, V., Kang, K., Liba, M. J., Vrabie, D., Bitmead, R. & Narayanan, S. (2011), Parameter estimation of a building system model and impact of estimation error on closed-loop performance, *in* ‘Proceedings of the IEEE Conference on Decision and Control’, Ieee, pp. 5137–5143.
- Berthou, T., Stabat, P., Salvazet, R. & Marchio, D. (2014), ‘Development and validation of a gray box model to predict thermal behavior of occupied office buildings’, *Energy and Buildings* **74**, 91–100.
- Camacho, E. & Alba, C. (2007), *Model predictive control*, second edn, Springer-Verlag, London.
- Chandan, V. & Alleyne, A. (2013), ‘Optimal partitioning for the decentralized thermal control of buildings’, *IEEE Transactions on Control Systems Technology* **21**(5), 1756–1770.
- Chandan, V. & Alleyne, A. G. (2011), Optimal control architecture selection for thermal control of buildings, *in* ‘IFAC Proceedings Volumes (IFAC-PapersOnline)’, Vol. 18, Ieee, pp. 3090–3095.
- Chandan, V. & Alleyne, A. G. (2012), ‘Decentralized architectures for thermal control of buildings’, *2012 American Control Conference (ACC)* **1**, 3657–3662.

- Chandan, V. & Alleyne, A. G. (2014), 'Decentralized predictive thermal control for buildings', *Journal of Process Control* **24**(6), 820–835.
- Chinde, V., Heylmun, J. C., Kohl, A., Jiang, Z., Sarkar, S. & Kelkar, A. (2015), Comparative Evaluation of Control Oriented Zone Temperature Prediction Modeling Strategies in Buildings, in 'Proceedings of the ASME 2015 Dynamic Systems and Control Conference'.
- Cigler, J. & Privara, S. (2010), 'Subspace identification and model predictive control for buildings', *2010 11th International Conference on Control Automation Robotics & Vision* (December), 750–755.
- Cole, W. J., Hale, E. T. & Edgar, T. F. (2013), Building energy model reduction for model predictive control using OpenStudio, in 'American Control Conference', pp. 449–454.
- Deng, K., Goyal, S., Barooah, P. & Mehta, P. G. (2014), 'Structure-preserving model reduction of nonlinear building thermal models', *Automatica* **50**(4), 1188–1195.
- Ferkl, L., Široký, J. & Prívara, S. (2010), Model predictive control of buildings: The efficient way of heating, in 'Proceedings of the IEEE International Conference on Control Applications', Yokohama, pp. 1922–1926.
- Freire, R. Z., Oliveira, G. H. C. & Mendes, N. (2008), 'Predictive controllers for thermal comfort optimization and energy savings', *Energy and Buildings* **40**(7), 1353–1365.
- Ghiaus, C. (2006), 'Experimental estimation of building energy performance by robust regression', *Energy and Buildings* **38**(6), 582–587.
- Gouda, M., Danaher, S. & Underwood, C. (2002), 'Building thermal model reduction using nonlinear constrained optimization', *Building and Environment* **37**, 1255–1265.
- Goyal, S., Liao, C. & Barooah, P. (2011), Identification of multi-zone building thermal interaction model from data, in 'Proceedings of the IEEE Conference on Decision and Control', Ieee, pp. 181–186.
- Hazyuk, I., Ghiaus, C. & Penhouet, D. (2012), 'Optimal temperature control of intermittently heated buildings using Model Predictive Control: Part I - Building modeling', *Building and Environment* **51**, 379–387.
- Kutner, M. H., Nachtsheim, C. J., Neter, J. & Li, W. (1996), *Applied Linear Statistical Models*, Vol. Fifth.

- Lazos, D., Sproul, A. B. & Kay, M. (2014), 'Optimisation of energy management in commercial buildings with weather forecasting inputs: A review', *Renewable and Sustainable Energy Reviews* **39**, 587–603.
- Li, X. & Wen, J. (2014), 'Review of building energy modeling for control and operation', *Renewable and Sustainable Energy Reviews* **37**, 517–537.
- Lin, Y., Middelkoop, T. & Barooah, P. (2012), Issues in identification of control-oriented thermal models of zones in multi-zone buildings, in 'Proceedings of the IEEE Conference on Decision and Control', Ieee, pp. 6932–6937.
- Ljung, L. (1999), *System Identification (2nd Ed.): Theory for the User*, second edn, Prentice-Hall, New Jersey.
- Ljung, L. (2002), 'Prediction error estimation methods', *Circuits, Systems, and Signal Processing* **21**(1), 11–21.
- Ma, J., Qin, S. & Salsbury, T. (2012), 'Model predictive control of building energy systems with balanced model reduction', *American Control Conference (ACC)*, 2012 pp. 3681–3686.
- Ma, Y., Matusko, J. & Borrelli, F. (2014), 'Stochastic Model Predictive Control for Building HVAC Systems: Complexity and Conservatism', *IEEE Transactions on Control Systems Technology* **23**(1), 1–1.
- Maasoumy, M., Razmara, M., Shahbakhti, M. & Vincentelli, a. S. (2014), 'Handling model uncertainty in model predictive control for energy efficient buildings', *Energy and Buildings* **77**, 377–392.
- Maciejowski, J. (2002), *Predictive Control: With Constraints*, Prentice-Hall, UK.
- Moroşan, P.-D., Bourdais, R., Dumur, D. & Buisson, J. (2010), 'Building temperature regulation using a distributed model predictive control', *Energy and Buildings* **42**(9), 1445–1452.
- Palomo Del Barrio, E. (2000), 'Using model size reduction techniques for thermal control applications in buildings', *Energy and Buildings* **33**(1), 1–14.
- Privara, S., Cigler, J., Vana, Z., Ferkl, L. & Sebek, M. (2010), 'Subspace Identification of Poorly Excited Industrial Systems', *49th IEEE Conference on Decision and Control* pp. 4405–4410.
- Radecki, P. & Hencey, B. (2013), Online thermal estimation, control, and self-

- excitation of buildings, in 'Proceedings of the IEEE Conference on Decision and Control', pp. 4802–4807.
- Ríos-Moreno, G. J., Trejo-Perea, M., Castañeda-Miranda, R., Hernández-Guzmán, V. M. & Herrera-Ruiz, G. (2007), 'Modelling temperature in intelligent buildings by means of autoregressive models', *Automation in Construction* **16**, 713–722.
- Rojas, C. R., Barenthin, M. & H. Hjalmarsson (2008), The cost of complexity in identification of {FIR} systems, in '17th IFAC World Congress', pp. 11451–11456.
- Touretzky, C. R. & Baldea, M. (2013), 'Model reduction and nonlinear MPC for energy management in buildings', *2013 American Control Conference* pp. 461–466.
- Wu, S. & Sun, J. Q. (2012), 'A physics-based linear parametric model of room temperature in office buildings', *Building and Environment* **50**, 1–9.

Chapter 5

Scalable Model Predictive Control Strategies for Building Heating Systems

5.1 Introduction

The focus of this chapter is on the design of suitable MPC formulations which can be set-up, adjusted and characterised in a way that can be scaled to different buildings without the need for significant commissioning effort or expert knowledge of advanced control.

Though not yet common in practice (Sturzenegger et al. 2015), many examples of MPC-type control of building energy systems can be found in literature showing an improved performance when compared to typical rule-based control (RBC) strategies (Shaikh et al. 2014, Lazos et al. 2014, Ma, Borrelli, Hencsey, Coffey, Bengesa & Haves 2012). The need for appropriate system models specific to each building in which the control is implemented however, allied to the need to tune the MPC problem in accordance with these system-specific models, is a significant disadvantage when compared to the simplicity of empirically adjusted RBC. Here, standard MPC approaches are examined, with new strategies then developed in which the improvement in energy efficiency and user comfort can be achieved without the need for extensive tuning and high optimisation complexity.

Objectives suited to the application of building energy systems are first outlined, highlighting issues that arise when attempting to objectively classify the subjective concept

of thermal comfort. A commonly used quadratic MPC formulation is then implemented in the simulation platform, without the presence of disturbances or uncertainties in the predictions. It is shown that, given the tendency for comfort and energy objectives to oppose one another (Yang & Wang 2012, O'Dwyer et al. 2015), defining a desired overall performance (and appropriately balancing the cost function to achieve this performance) is not straightforward.

Issues of scale and scalability are considered in terms of both the numerical optimisation complexity and the tuning problem complexity. Methods for simplifying the problem are introduced as well as guidelines for characterising the relationship between tuning parameters and tangible performance aspects.

Finally, a prioritised-objective approach is introduced which allows for the comfort and energy objectives to be solved separately in a way that precludes the need for tuning. The strategy is formulated to allow for nonlinear heating system elements to be set-apart from the linear constraints of the zone models, thus enabling the use of the strategy in larger systems. The prioritised-objective approach is shown to achieve improved energy efficiency while enabling a user, unfamiliar with advanced control, to adjust the comfort level obtained.

5.2 Defining the Cost - Balancing Energy Consumption with Occupant Comfort

5.2.1 Performance Criteria to be Optimised

The performance objectives of foremost importance to any building heating (or cooling) system control strategy could be separated into two main categories: reduction of energy consumption and satisfaction of the occupant's comfort demands (Castilla et al. 2011). The former objective is the more conceptually unambiguous, typically consisting of a cost in terms of units of energy (Liao & Dexter 2010, Sturzenegger et al. 2015) or units of currency (Oldewurtel et al. 2010, Ma, Qin, Salisbury & Xu 2012).

The comfort objective can be somewhat more abstract. Crucially, the notion of *comfort satisfaction* in a general sense is subjective to each individual occupant. A commonly used index for quantifying comfort is the Predicted Mean Vote (PMV), as outlined in Chapter 2, which is used for predictive control purposes in (Ferreira et al. 2012) and (Freire et al. 2008). Limitations associated with the use of PMV are outlined in

(Humphreys & Fergus Nicol 2002) in which it is noted that in surveys of individual buildings, the actual observed mean vote often does not correspond to calculated values. Furthermore, as the PMV model is nonlinear and quite complicated, it may be more suited to model-free approaches. For many strategies (Shaikh et al. 2014), particularly those for which humidity control is not available (as is the case for the hydronic heating system studied here), a comfort cost based on the deviation of the zone temperature from a given set-point is used. For the remainder of this chapter, *comfort* is defined by the proximity of a zone temperature to its set-point.

5.2.2 Standard Formulations - Quadratic Cost Functions

Standard model predictive control approaches seek to predict future inputs that will minimise a particular cost function (or performance index) over a receding horizon, with the first element of the calculated control sequence then applied to the plant. At the next sample the system states are measured (or estimated) and the updated optimisation problem is repeated. Typically, a state-space structure is used for the optimisation model, with a constrained numerical optimisation employed to determine the future control sequence (Maciejowski 2002, Camacho & Alba 2007). The cost function at the k^{th} sample is often of the form:

$$J(k) = \sum_{i=1}^H \|\hat{z}(k+i|k) - r(k+i)\|_{\mathcal{Q}}^2 + \sum_{i=0}^{H-1} \|\Delta \hat{u}(k+i|k)\|_{\mathcal{R}}^2 \quad (5.1)$$

where $\hat{z}(k+i|k)$ is the output predicted for i steps in the future, $r(k+i)$ is some desired reference, and $\hat{u}(k+i|k)$ is the predicted input, using a H -step prediction horizon. The tuning matrices \mathcal{Q} and \mathcal{R} can be adjusted to achieve different levels of control aggression. Larger values in the \mathcal{Q} matrix penalize deviations from the set-point more heavily, while larger \mathcal{R} values encourage smoother control trajectories by penalizing the accumulated input increment.

This objective is formulated so as to minimise the deviation between the plant and the reference, with the control increment included to introduce integral action to the formulation (Maciejowski 2002). The reasoning behind this standard cost function does not, however, naturally extend to the problem of building heating systems. Considering the performance criteria of Section 5.2.1, the control strategy should typically seek to minimise the deviation of the outputs (the zone temperatures) from a reference, while also minimising the sum of the squared inputs (to reduce energy supplied to the

building), as opposed to the sum of the input increments:

$$J(k) = \sum_{i=1}^H \|\hat{z}(k+i|k) - r(k+i)\|_{\mathcal{Q}}^2 + \sum_{i=0}^{H-1} \|\hat{u}(k+i|k)\|_{\mathcal{R}}^2 \quad (5.2)$$

Variations of this cost formulation can be seen in the building energy literature (Chandan & Alleyne 2013, 2014). By the nature of the problem, as energy is often required to improve comfort, the objectives of set-point tracking and input reduction will tend to oppose each-other. The result of this will be a cost function that attempts to strike a balance between comfort and energy, the bias of which will depend on how the function is weighted. This is a subjective problem which will vary with the preference of the user and the specifics of the underlying models. A strategy in which a non-trivial tuning procedure is required for each building in which the strategy is used is far from ideal.

A common strategy employed to avoid the inclusion of contradicting objectives in a single cost function is to use an economic MPC formulation where only the energy supplied is minimised (Oldewurtel et al. 2013, Hazyuk et al. 2012, Sturzenegger et al. 2015). Often the zone temperatures are included in the constraints instead of the cost function through the addition of a temperature band constraint (Široký et al. 2011). If hard comfort constraints are used, the optimisation problem may become infeasible (due to limits on the inputs). For robustness, to soften these constraints, slack variables are introduced (Cigler & Privara 2010) which, once again, must be included in the cost function with appropriately tuned weights.

5.2.3 Implementation of Standard Quadratic Objective in the Simulation Platform

To highlight the performance benefits of an MPC approach when compared to the more traditional empirical approaches, the standard strategy using a cost function of the form of (5.2) was compared to the standard weather-compensation strategy by implementing both strategies in the simulation platform, previously derived to represent the dynamics of the Nimbus Centre. This initial example is intended as a theoretical base-case in which unmeasured disturbances are not included, external temperature forecasts are assumed to be correct and nonlinear components associated with the heating system (for example the boiler efficiency curves) are disregarded.

For the j^{th} zone of the building (on the i^{th} -floor), the sequence of predicted zone

temperatures over a prediction horizon H , where the matrices $A_{i,j}$, $B_{i,j}$, $C_{i,j}$ and $E_{i,j}$ are the model matrices derived using the PEM/PCA methods outlined in Section 4.6, can be expressed as:

$$\begin{bmatrix} T_{z_{i,j}}(k+1) \\ T_{z_{i,j}}(k+2) \\ \vdots \\ T_{z_{i,j}}(k+H) \end{bmatrix} = \begin{bmatrix} C_{i,j}A_{i,j} \\ C_{i,j}A_{i,j}^2 \\ \vdots \\ C_{i,j}A_{i,j}^H \end{bmatrix} \mathbf{x}_{i,j}(k) + \begin{bmatrix} C_{i,j}B_{i,j} & \cdots & 0 \\ C_{i,j}A_{i,j}B_{i,j} & \ddots & 0 \\ \vdots & \ddots & \vdots \\ C_{i,j}A_{i,j}^{H-1}B_{i,j} & \cdots & C_{i,j}B_{i,j} \end{bmatrix} \mathbf{T}_{fl_{sp_i}}(k) + \begin{bmatrix} C_{i,j}E_{i,j} & \cdots & 0 \\ C_{i,j}A_{i,j}E_{i,j} & \ddots & 0 \\ \vdots & \ddots & \vdots \\ C_{i,j}A_{i,j}^{H-1}E_{i,j} & \cdots & C_{i,j}E_{i,j} \end{bmatrix} \mathbf{T}_e(k) \quad (5.3)$$

$$\mathbf{T}_{fl_{sp_i}}(k) = [T_{fl_{sp_i}}(k) \cdots T_{fl_{sp_i}}(k+H-1)]^T \quad (5.4)$$

$$\mathbf{T}_e(k) = [T_e(k) \cdots T_e(k+H-1)]^T \quad (5.5)$$

where $\mathbf{x}_{i,j}(k)$ denotes the vector of model states at time k .

The Nimbus Centre consists of two floors, with 11 zones on the ground-floor and 14 zones on the first-floor. Where M_i is the number of zones on the floor, projected over a horizon of H steps, the vector of current states of the zone models on the i^{th} floor is given as:

$$\mathbf{x}_i(k) = [\mathbf{x}_{i,1}^T(k), \cdots, \mathbf{x}_{i,M_i}^T(k)]^T \quad (5.6)$$

The vector of predicted temperature outputs for the i^{th} floor is given as:

$$\mathbf{T}_{z_i}(k) = [T_{z_{i_1}}(k+1), T_{z_{i_2}}(k+1) \cdots T_{z_{i_{M_i}}}(k+1) \cdots T_{z_{i_1}}(k+H), T_{z_{i_2}}(k+H) \cdots T_{z_{i_{M_i}}}(k+H)]^T \quad (5.7)$$

This can be determined in terms of the predicted sequence of input flow-temperatures ($\mathbf{T}_{fl_{sp_i}}(k)$) and the predicted sequence of external temperatures ($\mathbf{T}_e(k)$) as:

$$\mathbf{T}_{z_i}(k) = \Psi_i \mathbf{x}_i(k) + \Phi_i \mathbf{T}_{fl_{sp_i}}(k) + \varphi_i \mathbf{T}_e(k), \quad (5.8)$$

where, with \oplus denoting the direct sum operation:

$$\Psi_i = \begin{bmatrix} \bigoplus_{j=1}^{M_i} C_{i,j} A_{i,j} \\ \vdots \\ \bigoplus_{j=1}^{M_i} C_{i,j} A_{i,j}^H \end{bmatrix} \quad (5.9)$$

$$\Phi_{i(j-1)M_i+q,r} = \begin{cases} C_{i,q} A_{i,q}^{j-r} B_{i,q}, & \text{if } j-r \geq 0 \\ 0, & \text{otherwise} \end{cases} \quad (5.10)$$

$$\varphi_{i(j-1)M_i+q,r} = \begin{cases} C_{i,q} A_{i,q}^{j-r} E_{i,q}, & \text{if } j-r \geq 0 \\ 0, & \text{otherwise} \end{cases} \quad (5.11)$$

$$\forall j = 1 \cdots H \quad \forall r = 1 \cdots H \quad \forall q = 1 \cdots M_i$$

Denoting the zone temperature set-point as T_{sp} , the vector of predicted deviations from the comfort set-point in each i^{th} -floor zone is then given as:

$$\epsilon_i(k) = T_{sp} \mathbf{1}_{M_i H} - \mathbf{T}_{z_i}(k) \quad (5.12)$$

where $\mathbf{1}_{M_i H}$ denotes a column of ones of length $M_i H$. Representing the predicted free response of the i^{th} -floor zones as $\mathbf{y}_{f_i}(k) = \Psi_i \mathbf{x}_i(k) + \varphi_i \mathbf{T}_e(k)$, the vector of predicted set-point deviations for the floor can be given as:

$$\epsilon_i(k) = T_{sp} \mathbf{1}_{M_i H} - \mathbf{y}_{f_i}(k) - \Phi_i \mathbf{T}_{f_{sp_i}}(k) \quad (5.13)$$

Among the primary objectives of the strategy is to reduce energy consumption, which in this formulation, as the boiler is omitted for simplicity, is equivalent to reducing the accumulated heat flow into each of the zones of the building. Using (4.36), the flow to the j^{th} zone on floor i at time k can be given as:

$$Q_{i,j}(k) = \delta_{i,j} (T_{f_{l_{sp_i}}}(k) - T_{z_{i,j}}(k)) \quad (5.14)$$

where $\delta_{i,j}$ is related to the heat transfer coefficient of the radiator in the zone. Using (5.8) and (5.14), the total predicted heat flow to the i^{th} -floor (from the i^{th} mixing valve) over the prediction horizon, which is given by:

$$\mathbf{Q}_{MV_i}(k) = \left[\sum_{j=1}^{M_i} Q_{i,j}(k) \cdots \sum_{j=1}^{M_i} Q_{i,j}(k+H-1) \right] \quad (5.15)$$

can then be expressed as:

$$\mathbf{Q}_{MV_i}(k) = (I_H \otimes \boldsymbol{\delta}_i) \left((I_H \otimes \mathbf{1}_{M_i} - \Phi_i) \mathbf{T}_{fl_{sp_i}}(k) - \mathbf{y}_{f_i}(k) \right) \quad (5.16)$$

where \otimes is the Kronecker product, I_H denotes a $H \times H$ identity matrix and $\boldsymbol{\delta}_i = [\delta_{i,1}, \dots, \delta_{i,M_i}]$.

Denoting for brevity:

$$\begin{aligned} \Lambda_{1_i} &= (I_H \otimes \boldsymbol{\delta}_i) \\ \Lambda_{2_i} &= (I_H \otimes \mathbf{1}_{M_i} - \Phi_i) \\ \Lambda_{3_i} &= (\Lambda_{1_i} \Lambda_{2_i})^{-1} \end{aligned}$$

and rearranging (5.16), a relationship between $\boldsymbol{\epsilon}_i$ and \mathbf{Q}_{MV_i} can be formed (by substituting for $\mathbf{T}_{fl_{sp_i}}$ in (5.13)) as:

$$\boldsymbol{\epsilon}_i(k) = T_{sp} \mathbf{1}_{M_i H} - (\Phi_i \Lambda_{3_i} \Lambda_{1_i} + I_{M_i H}) \mathbf{y}_{f_i}(k) - \Phi_i \Lambda_{3_i} \mathbf{Q}_{MV_i}(k) \quad (5.17)$$

For the simulation platform, with 25 zones on 2 floors, the vector of decision variables to be found by numerical optimisation are then given as $\boldsymbol{\theta} = [\boldsymbol{\epsilon}_1^T, \boldsymbol{\epsilon}_2^T, \mathbf{Q}_{MV_1}^T, \mathbf{Q}_{MV_2}^T]^T$. The quadratic cost function for the problem can now be expressed in terms of these variables as:

$$J = \min_{\boldsymbol{\theta}} \left(\boldsymbol{\theta}^T \begin{bmatrix} \mathcal{Q} & \mathbf{0}_{25H \times 2H} \\ \mathbf{0}_{2H \times 25H} & \mathcal{R} \end{bmatrix} \boldsymbol{\theta} \right) \quad (5.18)$$

where $\mathcal{Q} \in \mathbb{R}^{25H \times 25H}$ and $\mathcal{R} \in \mathbb{R}^{2H \times 2H}$ are weighting matrices. A set of linear equality constraints can also be constructed using (5.17). For the full building (ground-floor and first-floor) these can be expressed as:

$$\begin{aligned} \begin{bmatrix} I_{11H} & \mathbf{0}_{11H \times 14H} & \Phi_1 \Lambda_{3_1} & \mathbf{0}_{11H \times H} \\ \mathbf{0}_{14H \times 11H} & I_{14H} & \mathbf{0}_{14H \times H} & \Phi_2 \Lambda_{3_2} \end{bmatrix} \boldsymbol{\theta} \\ = \begin{bmatrix} T_{sp} \mathbf{1}_{11H} - (\Phi_1 \Lambda_{3_1} \Lambda_{1_1} + I_{11H}) \mathbf{y}_{f_1} \\ T_{sp} \mathbf{1}_{14H} - (\Phi_2 \Lambda_{3_2} \Lambda_{1_2} + I_{14H}) \mathbf{y}_{f_2} \end{bmatrix} \end{aligned} \quad (5.19)$$

Additional inequality constraints can then be used to ensure that the flow-temperature of the water in the radiators remains within maximum and minimum limits, denoted here as $T_{fl_{sp}}^+$ and $T_{fl_{sp}}^-$ respectively. In reality, these bounds are functions of the amount of heat available from the heating system. In this simplified example however, as the

boiler is omitted, they are assigned as constant values. Using (5.16), the inequality constraints can be formed in terms of θ as:

$$\begin{bmatrix} 0_{4H \times 25H} & \vdots & \frac{I_{2H}}{I_{2H}} \\ & & \vdots & \frac{I_{2H}}{I_{2H}} \end{bmatrix} \theta \leq \begin{bmatrix} \Lambda_{1_1} (\Lambda_{2_1} T_{fl_{sp}}^+ \mathbf{1}_H - \mathbf{y}_{f_1}) \\ \Lambda_{1_2} (\Lambda_{2_2} T_{fl_{sp}}^+ \mathbf{1}_H - \mathbf{y}_{f_2}) \\ -\Lambda_{1_1} (\Lambda_{2_1} T_{fl_{sp}}^- \mathbf{1}_H - \mathbf{y}_{f_1}) \\ -\Lambda_{1_2} (\Lambda_{2_2} T_{fl_{sp}}^- \mathbf{1}_H - \mathbf{y}_{f_2}) \end{bmatrix} \quad (5.20)$$

The outlined strategy was implemented in the simulation platform without the inclusion of disturbances and with the assumption of perfect external temperature forecasts. Second-order models were used, derived with data taken from the simulation platform using the PEM parameter identification methodologies and PCA-based spatio-temporally filtered disturbances estimation techniques developed in Chapter 4. As no nonlinear elements were included at this point, the optimisation was carried out using an interior-point algorithm inbuilt in the *quadprog* function of Matlab (used for solving constrained optimisation problems with quadratic costs and linear constraints).

For the application of building energy, different approaches are taken in the literature with regard to horizon choice. In (Mendoza-Serrano & Chmielewski 2012) for example, an infinite-horizon economic MPC formulation is outlined. Using an infinite horizon can provide stability guarantees and with sufficient prediction accuracy, a more desirable behaviour may be obtained. Given the unpredictable nature of the disturbances affecting the building however, as well as the complexity of the underlying thermal processes, longer prediction horizons may not be beneficial, as noted in (Brooks et al. 2015). In (Lefort et al. 2013), the impact of using different finite horizons is more thoroughly investigated in order to select an approach which achieves a performance that is sufficiently close to optimal, without excessive complexity.

Using a sample time of 600 seconds, a period of 9 days was simulated for 3 different prediction horizons - 5 samples (50 minutes), 10 samples (100 minutes) and 20 samples (200 minutes). The \mathcal{Q} and \mathcal{R} matrices were chosen to be diagonal in structure, tuned by trial-and-error to provide a high level of thermal comfort without excessive aggression. To simulate the weather-compensation strategy, $T_{fl_{sp_1}}$ and $T_{fl_{sp_2}}$ were determined based on the external temperature using the relationship:

$$T_{fl_{sp_1}}(k) = T_{fl_{sp_2}}(k) = 80 - 2.5T_e(k) \quad (5.21)$$

In all strategies, the zone-level valves were operated using a comfort band of $\pm 1^\circ\text{C}$. Fig. 5.1 shows the simulated temperature obtained for the ground-floor meeting room

using each of the horizon values, while the optimal input calculated for the ground floor ($T_{fl_{sp1}}$) is shown in Fig. 5.2.




Figure 5.1: Zone temperature using MPC with different prediction horizons for the ground-floor meeting room (2 day period)




Figure 5.2: Input ($T_{fl_{sp1}}$) determined using MPC with different prediction horizons (ground-floor, 2 day period)

It can be seen that the predictive strategies anticipated the commencement of the occupancy period before it occurred, reducing the set-point deviation at the start of each simulated day. A longer prediction horizon allows for predictable changes (such as the set-point change during occupied hours) to be accounted for further in advance. In the second day, a certain amount of overheating can be seen in the MPC strategies. As only one input is available for several zones ($T_{fl_{sp_i}}$), this may be necessary to avoid under-heating in another zone. The tuning of the cost function also resulted in an aggressively changing input signal.

A second set of results for the ground-floor meeting room zone temperatures is shown in Fig 5.3 using a different set of tuning parameters, in this case, balanced more in

favour of energy reduction (each element of the updated Q matrix was 10 times smaller than before). The comfort deviation was significantly larger than in the previous example, though the energy consumed was less. A less aggressive control action can be seen in the input profile ($T_{fl_{sp_i}}$) shown in Fig 5.4. It can also be seen that the horizon length has a noticeable effect on the tuning of the comfort/energy balance.




Figure 5.3: Zone temperature using re-tuned MPC with different prediction horizons (ground-floor meeting room - 2 day period)




Figure 5.4: Input ($T_{fl_{sp_1}}$) determined using re-tuned MPC with different prediction horizons (ground-floor - 2 day period)

The total accumulated set-point deviation for all zones of the building obtained using prediction horizons of 5 samples and 20 samples with each set of tuning parameters can be seen in Fig. 5.5. This was measured by finding the average set-point error in each zone for every hour and summing across the full building. The accumulation over time of this set-point deviation is then given in $^{\circ}C.h$. The energy consumption of the building can be seen in Fig. 5.6, given in kWh . Unsurprisingly, the problem appears to be sensitive to how the cost is tuned but, perhaps more significantly, the performance

obtained from a particular set of tuning parameters is highly sensitive to the specifics of the formulation (such as the prediction horizon selected).




Figure 5.5: Accumulated set-point deviation across full building for different prediction horizons and tuning parameters




Figure 5.6: Accumulated energy consumption of building for different prediction horizons and tuning parameters

It should be noted at this point that the schedule for the weather-compensation approach could be empirically adjusted to achieve an improved comfort performance over the full occupancy period. It is also worth noting that by switching on the heating system in advance of the occupancy period, more energy will be expended.

As such, comparing the overall performance of different strategies in this sense is not straightforward. For any conditions, the best possible energy performance would be achieved by never switching the heating system on (most likely resulting in an unacceptable comfort performance). Conversely, to achieve an improved comfort performance at the expense of additional energy may not be a desirable outcome. Once again, the subjectivity involved in assigning an appropriate balance to the problem is clear.

Without suitable tuning, the optimal solution of an MPC strategy may not align in any way with what a user would consider to be *optimal*. A more suitable classification of comfort performance may require the definition of a generally acceptable level of temperature set-point deviation (preferably user-defined). The energy performance could then be taken as a measure of how efficiently this comfort level is achieved. Two significant obstacles prohibit the use of a standard QP-type approach to solve the problem using these criteria however. Determining a set of tuning parameters that can guarantee a specific comfort level (if possible within the limitations of the constraints) poses a significant challenge. Additionally, to achieve the most energy efficient trajectory, the boiler efficiency must be taken into account, thus introducing nonlinearities to the optimisation problem. Formulation of the problem in a more suitable structure to achieve these goals is considered later in the chapter.

5.3 Scale and Scalability

Maintaining a practicable level of simplicity is critical to the design of an MPC strategy for the application of building energy systems. Given that buildings can vary greatly in shape, size and construction, the commissioning effort required to adapt a control formulation to suit each new building must be considered as an important factor in the strategy-design process. The link between energy savings and additional information (and as such, additional complexity) is noted in (Goyal et al. 2013).

As outlined in Chapter 2, the predominant building heating control strategies follow empirically chosen rule based formulations, with no information of the building's thermodynamics incorporated. The resulting simplicity allows for building users to modify the strategies to achieve a more desirable level of thermal comfort. To retain this adaptability in an energy efficient strategy poses certain difficulties. As noted in (Peeters et al. 2008), though users tend to opt for energy efficient heating system appliances, they are often not then used properly (boiler settings, for example, are arbitrarily selected without consideration of the specific situation at hand).

Formulating a strategy which can achieve an energy efficient control performance by including knowledge of the thermodynamic properties of each specific building in which it is used will inevitably introduce a higher level of complexity (particularly for larger buildings). In this section, possible approaches for handling this complexity in a way that allows for the strategy to be scaled easily for use in a range of different buildings without an impractical commissioning requirement are considered.

5.3.1 Handling the Growing Complexity in Larger Buildings

5.3.1.1 Distributing the Optimisation Problem

In Chapter 4 the choice between centralised, decentralised and distributed zone models was discussed. By omitting the heat transfer between different zones through internal walls from the calculations, the dynamics of each zone can be regarded as independent from other zones. As common inputs are chosen in the optimisation process to be applied to multiple zones however (in the case of the Nimbus Centre, all zones on the same floor receive the same radiator flow-temperature), the individual zones models must still be combined in the same numerical optimisation. For smaller buildings, this is not problematic, provided the optimisation problem is reasonably well conditioned, as was the case for the Nimbus Centre example (2 floors, 25 zones, convex quadratic optimisation with linear constraints). If larger buildings with additional floors and many more zones are considered, calculation of the optimal constrained control solution may become prohibitively complex.

Distributed MPC (dMPC) provides a natural framework for solving optimisation problems in which a globally satisfactory solution is desired for a large system with many objectives (Camponogara et al. 2002, Ma et al. 2011), particularly when the interaction between subsystems is strong (Stewart et al. 2010). Examples in which dMPC is formulated for the application of building energy can be found in (Moroşan et al. 2010a, Ma et al. 2011, Moroşan et al. 2010b).

Broadly speaking, the distributed approach decomposes a large problem into a set of smaller interconnected subproblems. The local solution to each subproblem is found in parallel and communicated to neighbouring subproblems. Additionally, a cooperative element can be introduced whereby each of the local optimisations is iteratively repeated with its local solution communicated to neighbouring problems between iterations. If formulated correctly, this can allow for convergence towards a global solution, similar to what would be achieved using a centralised approach, but with far less computational effort (Camponogara et al. 2002).

In Fig. 5.7 the problem separation is shown (where the superscript p denotes a predicted output). In this example, the i^{th} local subproblem determines optimal local inputs u_i^* and the associated local output predictions y_i^p , using local measurements from the building y_i and the predicted outputs calculated by the neighbouring subproblems.

Due to the relatively small scale of the Nimbus Centre, as well as the difficulty in characterising the interaction between zones using data driven methods (as seen in

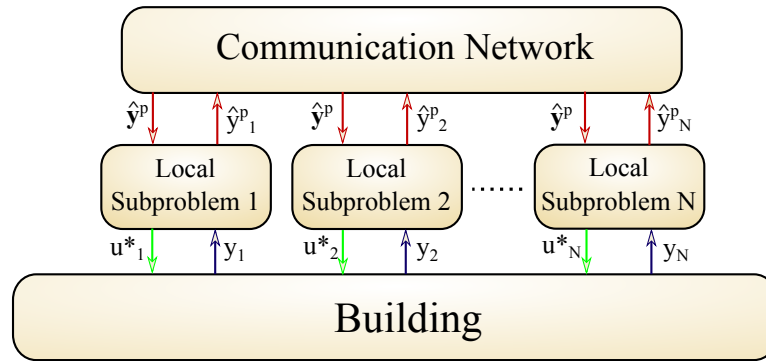


Figure 5.7: Distributed MPC - Problem is divided into manageable subproblems

Chapter 4), the optimisation problem was handled in this chapter without the use of these distributed approaches. Nevertheless, at a larger scale, such frameworks may become preferable.

5.3.1.2 A Hierarchical Approach - Reducing the Problem

An alternative approach, developed in (O'Dwyer et al. 2014), determines the optimal control for two prototype zones, representing the hot and cold extremes of the building. Zones of the building were classified as *hot* or *cold* based on historical measurements and assigned the same control input as the associated prototype. In this way, the large number of varying objectives and low number of controller options is acknowledged.

The middle-layer of the control hierarchy - the mixing valve layer - was used to optimise the comfort and energy balance in the cold prototype zone by selecting an appropriate radiator flow-temperature set-point. Control of the lowest layer of the hierarchy - the zone-level valves - is then optimised to prevent overheating in the hot prototype zone by switching off at appropriate times. In this strategy, the zone-level on-off valves follow a pulse-width modulation (PWM) signal, which, when averaged over one sample period, are open for a percentage of time equal to the duty cycle (d_{cyc}). This type of strategy is useful as it allows for a valve that is binary in nature to be represented by a continuous variable, thus avoiding the need for mixed integer optimisation (in reality, the operation of mechanical zone-level valves in this manner may not be realistic however, while the resolution would be limited by the speed of movement of the valves).

Using this strategy, the mixing valves prevent under-heating in *cold* zones while the zone-level valves prevent overheating in *hot* zones. Regardless of the number of zones in the building, the numerical optimisation need only consider the dynamics of two

zones, though the validity of the approach is dependent on how well the zones conform to the *hot* and *cold* designations. In Fig. 5.8, the control hierarchy of this representative zone approach is shown.




Figure 5.8: Control hierarchy of approach using hot and cold representative zones to simplify optimisation

In (O'Dwyer et al. 2014) this hierarchical approach was implemented in the Nimbus Centre simulation platform for a 7 day period using measured external temperature data from April 2013. With a horizon of 10 samples and a sample time of 15 minutes, the mixing-valve control was optimised for the coldest zone of the building, while the zone-level valve control was optimised for the hottest zone. The hot and cold zones were chosen based on data from a three month period. By including the zone-level valves in the optimisation process (by way of the PWM duty cycle d_{cyc}), the zone models can no longer be considered linear in form, with the heat supplied to each zone j on the i^{th} floor can now be given as:

$$Q_{i,j}(k) = d_{cyc_{i,j}}(k) \delta_{i,j} \left(T_{flsp_i}(k) - T_{z_{i,j}}(k) \right) \quad (5.22)$$

To account for this, an interior-point algorithm which can handle nonlinear constraints provided the problem is twice continuously differentiable is used (IPOPT - (Wächter & Biegler 2006)). In Fig. 5.10 and Fig. 5.9, the accumulated set-point deviation and energy consumption obtained using the hierarchical MPC approach and the standard weather-compensation approach is shown for the simulated period. Using a quadratic cost function once again (tuned by trial-and-error), the hierarchical formulation consumed approximately 14% less energy with a 15% reduction in set-point deviation.




Figure 5.9: Accumulated set-point deviation across the Nimbus Centre model using hierarchical MPC formulation and a more standard weather compensated formulation




Figure 5.10: Accumulated energy consumption of the Nimbus Centre model using hierarchical MPC formulation and a more standard weather compensated formulation

5.3.2 Tuning Considerations - Setting a Desired Level of Performance

Many examples can be found in literature highlighting an improved building heating/cooling system control performance by the use of an MPC or optimal control based strategy. A review of such strategies can be found in (Shaikh et al. 2014), while a brief review of strategies that have been fully implemented in buildings is given in (Sturzenegger et al. 2015). While the applicability of the strategy to buildings of larger scales must be considered, a truly scalable MPC strategy must also be adaptable in terms of the objective balance. As discussed, many zones, each with an individual comfort objective must compete with an overall building energy consumption objective in a way that is dependent on how the cost function is tuned. This can be seen in the simulated examples of Section 5.2.3. Despite the clear benefits, strategies are

unlikely to be adopted on a large scale if the tuning process for every building requires the presence of an MPC specialist.

In (O'Dwyer et al. 2015), a tuning methodology was proposed which provides a direct link between each tuning parameter and a tangible performance metric. More specifically, the parameters are chosen based on a notional steady-state set-point deviation in each zone taken from the analytic solution of the unconstrained optimisation problem. To allow for a practical determination of the analytic solution associated with the problem, several assumptions were made. A PWM-type control strategy was assumed for the zone-level valves, which allowed for the heat supply to each zone to be chosen independently of all other zones. First-order zone models and diagonal tuning matrices were also assumed.

If the inequality constraints associated with the heating system are omitted, a relationship between the tuning parameters and the set-point deviation can be found by observing the KKT conditions associated with the problem. While this simplification of the problem is useful for tuning parameter selection, the inequality constraints can be reinserted for the actual implementation of the MPC strategy.

Where ϵ_j is the sequence of future deviations from the comfort set-point of the temperature in zone j , and \mathbf{u}_j is the sequence of future heat inputs to the zone, the Karush-Kuhn-Tucker (KKT) conditions for the problem can be expressed using Lagrange multipliers λ and μ as:

$$\nabla J(\epsilon_j, \mathbf{u}_j) = - \sum_{q=1}^{2H} \mu_q \nabla g_q(\mathbf{u}_j) - \sum_{p=1}^H \lambda_p \nabla h_p(\epsilon_j, \mathbf{u}_j) \quad (5.23)$$

$$g_q(\mathbf{u}_j) \leq 0 \quad (5.24)$$

$$h_p(\epsilon_j, \mathbf{u}_j) = 0 \quad (5.25)$$

$$\mu_q \geq 0 \quad (5.26)$$

where the tuning matrices are given by $\mathcal{Q}_j = q_j I_H$ and $\mathcal{R}_j = I_H$ and the model

parameters given by a_j , b_j and e_j :

$$J(\epsilon_j, \mathbf{u}_j) = \sum_{i=1}^H \left(\epsilon_j^T \mathcal{Q}_j \epsilon_j + \mathbf{u}_j^T \mathcal{R}_j \mathbf{u}_j \right) \quad (5.27)$$

$$g_i = -u_j(k + i - 1|k) \quad (5.28)$$

$$g_{H+i} = u_j(k + i - 1|k) - u_{MAX} \quad (5.29)$$

$$h_i = T_{sp} - \epsilon_j(k + i|k) - a_j^i T_{z_j}(k) - \begin{bmatrix} a_j^{i-1} b_j & a_j^{i-2} b_j & \cdots & b_j & 0 & \cdots & 0 \end{bmatrix} \mathbf{u}_j \\ - \begin{bmatrix} a_j^{i-1} e_j & a_j^{i-2} e_j & \cdots & e_j & 0 & \cdots & 0 \end{bmatrix} \mathbf{T}_e \quad (5.30)$$

$$\forall i = 1, 2 \dots H$$

By omitting the inequality constraints, the equations of (5.23) and (5.25) can be solved in terms of the weight q_j . For a prediction horizon of $H = 2$ for example, the system of equations could be represented as:

$$\begin{bmatrix} 2q_j & 0 & 0 & 0 & -1 & 0 \\ 0 & 2q_j & 0 & 0 & 0 & -1 \\ 0 & 0 & 2 & 0 & -b_j & -a_j b_j \\ 0 & 0 & 0 & 2 & 0 & -b_j \\ 1 & 0 & b_j & 0 & 0 & 0 \\ 0 & 1 & a_j b_j & b_j & 0 & 0 \end{bmatrix} \begin{bmatrix} \epsilon_j^*(k+1|k) \\ \epsilon_j^*(k+2|k) \\ u_j(k|k) \\ u_j(k+1|k) \\ \lambda_1 \\ \lambda_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ T_{sp} - y_{f_j}(k+1|k) \\ T_{sp} - y_{f_j}(k+2|k) \end{bmatrix} \quad (5.31)$$

where y_f is the free response of the system. From this, the future optimal set-point deviation $\epsilon_j^*(k+1|k)$ can be found in terms of the current set-point deviation, $\epsilon_j(k|k) = T_{sp} - T_{z_j}(k)$, the tuning parameter q_j and the external temperature forecast. Here, where α_1 , α_2 and α_3 are H^{th} -order polynomials in terms of q_j , the following relationship can be formed:

$$\epsilon_j^*(k+1|k) = \alpha_1(q_j) \epsilon_j(k|k) + \alpha_2(q_j) T_{sp} + \alpha_3(q_j) \mathbf{T}_e \quad (5.32)$$

Once again for simplicity, if a perfect plant model is assumed and the external temperature is held constant (\bar{T}_e), an equation for the steady-state set-point deviation ϵ_{jss} can be formed as:

$$\epsilon_{jss} = \frac{\alpha_2(q_j) T_{sp} + \alpha_3(q_j) \bar{T}_e}{1 - \alpha_1(q_j)} \quad (5.33)$$

Using (5.33), a tuning parameter can be selected for each zone, j , to obtain a specified value for ϵ_{jss} . Though many simplifications are required - most significantly the use of first-order zone models and the assumption that the zone-level valves can be operated

in a pulse-width-modulated fashion - taking this approach offers the ability to characterise a tuning parameter for each zone in terms of a tangible performance metric (the steady-state set-point deviation). While it is possible to get solutions in this manner using higher order zone models, the complexity of the relationship between ϵ_{jss} and q_j quickly becomes unmanageable, particularly as the horizon length is increased.

In (O'Dwyer et al. 2015), using a three day simulation period, this tuning and control strategy was implemented in the simulation platform with a prediction horizon of 10 samples (with a sample time of 15 minutes) and an external temperature profile taken from measured data from April 2013. The IPOPT interior-point algorithm was used once again due to the bilinear relationship between d_{cyc} and T_{flsp} .

The tuning parameters were selected to obtain a steady-state temperature deviation of 1°C below a set-point of 20°C . As an example, the relationship between the tuning parameter associated with the ground-floor meeting room and the steady-state set-point deviation for an external temperature of 10°C , is shown in Fig. 5.11. For this zone, to obtain a set-point deviation of 1°C , a value of 155 is chosen for the tuning weight.




Figure 5.11: Relationship between the steady-state set-point deviation and the tuning weight for ground floor meeting room

In Fig. 5.12, the output temperatures obtained for six of the ground-floor zones are shown. Disturbances in the form of solar and internal gains were included in the simulation platform, but, as they were considered unmeasured, they were not included in the tuning procedure. These disturbances resulted in slight over-heating at times.




Figure 5.12: 6 ground-floor zone temperatures for three day period using a tuning strategy based on the steady-state set-point deviation

5.4 A Prioritised-Objectives MPC Approach for Building Heating Systems

In this section, using a methodology developed in (O'Dwyer et al. 2016), an MPC formulation is constructed to handle the scalability issues inherent in the building energy control problem, resulting from the large number of objectives and tuning parameters. By separating the problem in a suitable manner, the costs and constraints are organized to allow for a more straightforward commissioning and readjustment process without the need for tuning weights. To improve energy efficiency, control at a higher level of the heating system - specifically the boiler - is incorporated in the strategy.

5.4.1 Multiple Objectives and Lexicographic MPC Formulation

For systems with a large number of controllable outputs, expressing all objectives in a single, appropriately weighted cost function can pose challenges due to the number of decision variables present. For certain systems, decomposing the problem into several objectives which are solved individually can provide a more conceptually intuitive framework for strategy design. A typical multi-objective formulation seeks a Pareto-optimal solution to each of the objectives, with the overall solution then being a Pareto-front, a solution set in which decreasing the cost of one objective cannot be achieved without increasing the cost of at least one other objective. Specifically, for a set of n objectives, $\mathbf{V}(\boldsymbol{\theta}) := [V_1(\boldsymbol{\theta}) \cdots V_n(\boldsymbol{\theta})]^T$, a Pareto optimal minimiser, $\boldsymbol{\theta}^* \in \Theta$, is one for which no point $\boldsymbol{\theta} \in \Theta$ exists that satisfies $V_i(\boldsymbol{\theta}) < V_i(\boldsymbol{\theta}^*)$ and $\mathbf{V}(\boldsymbol{\theta}) \leq \mathbf{V}(\boldsymbol{\theta}^*)$ $\forall i \in \{1, \dots, n\}$, where Θ is the set of feasible solutions. An overview of multi-

objective optimal control is given in (Gambier & Badreddin 2007).

For illustrative purposes, the set of feasible solutions for two objectives (J_1 and J_2) is shown in Fig. 5.13. For a particular subset, there are no alternative feasible solutions which result in a reduced cost for both objectives. A solution from this set is said to be a weak-Pareto-optimal solution. A multi-objective formulation seeks to find a minimum solution to a weighted combination of the objectives, which will fall within this Pareto front. The relative weighting of the objectives determines the point at which the overall solution is found. It should be noted that this is a two-dimensional version of the tuning problem faced when using the quadratic cost function of (5.2), in which the single objective could be equivalently considered as a set of multiple objectives corresponding to comfort in each zone and overall energy consumption.




Figure 5.13: Combined Pareto optimal solution for two objectives

In (Kerrigan & Maciejowski 2002), a lexicographic strategy in which priorities are assigned to the objectives was outlined. Arranging the objectives in order of descending priority from $1 \dots n$, the argument $\theta^* \in \Theta$ is a lexicographic minimiser of the overall problem if and only if:

$$\theta^* \in \{\theta \in \Theta | V_j(\theta) \leq V_j^*, j = 1, \dots, n\} \quad (5.34)$$

where, $\forall i \in \{2, \dots, n\}$:

$$V_1^* = \min_{\theta \in \Theta} V_1(\theta) \quad (5.35)$$

$$V_i^* = \min_{\theta \in \Theta} \{V_i(\theta) \mid V_j(\theta) \leq V_j^*, j = 1, \dots, i-1\} \quad (5.36)$$

Approaching the problem in this manner ensures that lower priority objectives are only improved if doing so does not have a detrimental effect on higher priority objectives. As a result of this ordering, weighting the objectives is not necessary. Furthermore, individual objectives can be removed or reordered without rebalancing the problem.

Using once again the two-objective example of Fig. 5.13, the Lexicographic solution is now shown in Fig. 5.14 where J_1 is taken as the primary objective. It can be seen that the solution is unambiguous.

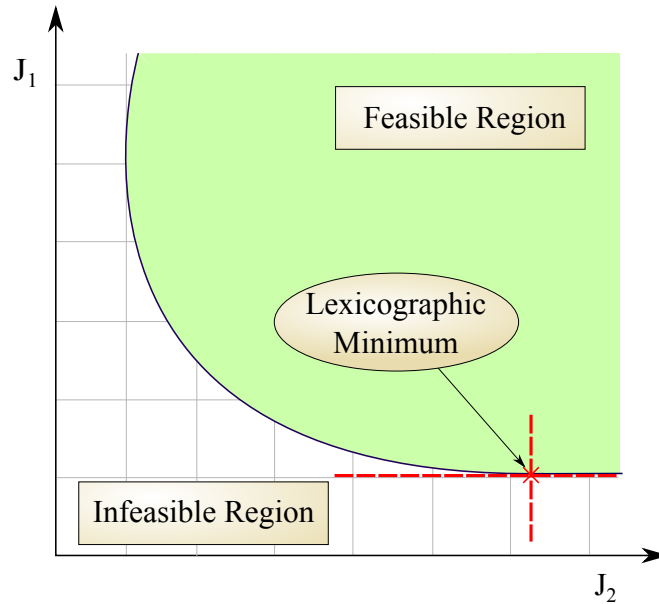


Figure 5.14: Lexicographic optimal solution for two prioritised-objectives

While the advantages are clear, the problem must be formulated in a suitable framework. Crucially, the first $n - 1$ objectives should not be strictly convex, as a unique solution would render lower priority objectives unnecessary.

In (O'Dwyer et al. 2015), a lexicographic approach was proposed for the building heating system domain. Once again, the main objectives to be considered are those of minimising temperature set-point deviation in all zones of the building, while also minimising energy consumption. It should be noted once again that in this lexicographic form, the problem is formed as a set of separable objectives, solved sequentially in a specified order. In this way it differs from standard multi-objective approaches which

seek to find some desirable Pareto-optimal solution to a weighted combination of objectives.

Achieving the first objective of the strategy requires driving all zone temperatures, \mathbf{T}_z , as close as possible to some comfort set-point, \mathbf{T}_{sp} , over some prediction horizon H when the building is in use. The second objective seeks to reduce the accumulated energy input to the heating system over the horizon, $\int_t^{t+H\Delta t} P_{in}(t)dt$. To accommodate these aims into a lexicographic framework, comfort satisfaction is taken as the primary objective while the energy consumption is minimised subject to the constraint that comfort-level determined by the primary objective is not impacted.

5.4.2 Development of the Cost Functions

To ensure a generally unique solution is not obtained for the primary objective, the set-point is replaced by a temperature band $\underline{\mathbf{T}}_{sp} < \mathbf{T}_{sp} < \bar{\mathbf{T}}_{sp}$. The deviation in zone temperature away from this band is penalised by the introduction of variables $\bar{\epsilon}$ and $\underline{\epsilon}$. These are vectors comprised of the zone temperature deviations from the set-point in each zone over the prediction horizon. Denoting the set of optimisation variables as $\boldsymbol{\theta} = [\bar{\epsilon}^T, \underline{\epsilon}^T, \mathbf{P}_{in}^T]^T$, the primary optimisation at time k can be formed (at this point ignoring the heating system constraints) as:

$$V_1^* = \min_{\boldsymbol{\theta}} \Phi_1 \boldsymbol{\theta} \quad (5.37)$$

s.t.

$$\bar{\epsilon}(k+i) \geq \mathbf{T}_z(k+i) - \bar{\mathbf{T}}_{sp} \quad (5.38)$$

$$\underline{\epsilon}(k+i) \geq \underline{\mathbf{T}}_{sp} - \mathbf{T}_z(k+i) \quad (5.39)$$

$$\bar{\epsilon}, \underline{\epsilon} \geq 0 \quad (5.40)$$

$$\forall i = 1 \dots H$$

For a building of N zones, the vector $\Phi_1 \in \Re^{H(2N+1)}$ is given by:

$$\Phi_1 = \begin{bmatrix} \mathbf{1}_{NH}^T & \mathbf{1}_{NH}^T & \mathbf{0}_{1 \times H} \end{bmatrix}, \quad (5.41)$$

where $\mathbf{1}_{NH}$ is a column of NH ones. The linear cost is convex, but not strictly convex and so the solution is not generally unique.

The secondary objective can then be expressed as:

$$V_2^* = \min_{\theta} \theta^T \Phi_2 \theta \quad (5.42)$$

s.t.

$$V_1(\theta) \leq V_1^*, \quad (5.43)$$

where in this case:

$$\Phi_2 = \begin{bmatrix} 0_{2NH \times 2NH} & 0_{2NH \times H} \\ 0_{H \times 2NH} & I_{H \times H} \end{bmatrix} \quad (5.44)$$

On first inspection, the cost function does not appear to be strictly convex as Φ_2 is semi-positive definite. As this is the final objective, a unique solution is more desirable. Considering however that the minimum of $V_1(\theta)$ has been found in (5.37), the constraint given by (5.43) ensures that $V_1(\theta) = V_1^*$. As this is a constant, it can be added to the cost function without affecting the solution:

$$\theta^* = \arg \min_{\theta} (\theta^T \Phi_2 \theta + V_1(\theta)) \quad (5.45)$$

The cost function can now be rewritten as:

$$V(\theta) = P_{in}^T P_{in} + [1, \dots, 1] \begin{bmatrix} \bar{\epsilon} \\ \epsilon \end{bmatrix}, \quad (5.46)$$

which can be shown to be a strictly convex function by the following analysis. Two possible solutions, denoted θ_1 and θ_2 are first defined as:

$$\theta_1 = [P_{in_1}, \bar{\epsilon}_1, \epsilon_1]^T \quad (5.47)$$

$$\theta_2 = [P_{in_2}, \bar{\epsilon}_2, \epsilon_2]^T \quad (5.48)$$

To prove strict convexity, where $\theta_1 \neq \theta_2$ the following inequality can be shown:

$$\lambda V(\theta_1) + (1 - \lambda)V(\theta_2) > V(\lambda\theta_1 + (1 - \lambda)\theta_2) \quad (5.49)$$

$$\forall \lambda \in (0, 1)$$

Substituting (5.47) and (5.48) into (5.49), the inequality becomes:

$$\lambda \left(\mathbf{P}_{in_1}^T \mathbf{P}_{in_1} + [1, \dots, 1] \begin{bmatrix} \bar{\epsilon}_1 \\ \underline{\epsilon}_1 \end{bmatrix} \right) + (1 - \lambda) \left(\mathbf{P}_{in_2}^T \mathbf{P}_{in_2} + [1, \dots, 1] \begin{bmatrix} \bar{\epsilon}_1 \\ \underline{\epsilon}_2 \end{bmatrix} \right) > \\ (\lambda \mathbf{P}_{in_1} + (1 - \lambda) \mathbf{P}_{in_2})^T (\lambda \mathbf{P}_{in_1} + (1 - \lambda) \mathbf{P}_{in_2}) + [1, \dots, 1] \begin{bmatrix} \lambda \bar{\epsilon}_2 + (1 - \lambda) \bar{\epsilon}_1 \\ \lambda \underline{\epsilon}_1 + (1 - \lambda) \underline{\epsilon}_2 \end{bmatrix} \quad (5.50)$$

This simplifies to:

$$\lambda \mathbf{P}_{in_1}^T \mathbf{P}_{in_1} + (1 - \lambda) \mathbf{P}_{in_2}^T \mathbf{P}_{in_2} > \lambda^2 \mathbf{P}_{in_1}^T \mathbf{P}_{in_1} \\ + \lambda (1 - \lambda) (\mathbf{P}_{in_1}^T \mathbf{P}_{in_2} + \mathbf{P}_{in_2}^T \mathbf{P}_{in_1}) + (1 - \lambda)^2 \mathbf{P}_{in_2}^T \mathbf{P}_{in_2} \quad (5.51)$$

By rearranging this, the following inequality is found:

$$[\mathbf{P}_{in_1} - \mathbf{P}_{in_2}]^T [\mathbf{P}_{in_1} - \mathbf{P}_{in_2}] > 0 \quad (5.52)$$

which is true for $\mathbf{P}_{in_1} \neq \mathbf{P}_{in_2}$.

From the perspective of the end-user, the tuning and adjustment of the MPC strategy to achieve some desired performance should now be more accessible. An acceptable level of zone temperature set-point deviation is set by assigning a width to the comfort band around the set-point. The MPC strategy will attempt to push the zone temperatures into, or as close as possible to, this comfort band without violating the heating system constraints. The formulation will then seek to achieve this level of comfort in the most energy efficient way possible.

5.4.3 Supply and Demand - Nonlinear Heating System Components

The previous strategies outlined in this chapter sought to minimise the energy required by the building to minimise discomfort. This could be seen as the demand-side problem. The use of MPC is advantageous for this as, if correctly configured, the need to empirically determine a suitable heating schedule is removed and additionally, the overheating of zones should be less likely. To truly achieve a greater level of energy efficiency, however, the generation of the heat energy should also be optimised (this could be thought of as the supply-side problem). Here, the heat is assumed to be generated by a boiler. From this point onwards the energy consumed by the boiler, as

opposed to the energy required by the building, is taken as the most significant energy term to be minimised.

In order to incorporate the efficiencies and constraints associated with the boiler in the prioritised framework, the heating system must again be considered. In this case, a three-layer hydronic heating system for a generic building of F floors is examined. At the highest level, enough heat energy must be produced (by the boiler) to heat the building as a whole. The heat must then be distributed from the system header to all parts of the building at a suitable temperature. Finally, thermostatic valves allow radiators in each zone to switch off if the zone temperature exceeds a set-point. As before, the temperature of the water supplied to all radiators on each floor is dictated by PI-controlled mixing valves while the zone-level thermostatic valves follow a hysteresis loop. In Fig. 5.15, the three layers of the heating system hierarchy are shown for the building.




Figure 5.15: Three-layer heating system control hierarchy

Maintaining the autonomy of the lowest layer (the zone-level thermostatic-valves) is desirable for robust operation (Maciejowski 2002, O'Dwyer et al. 2015). The MPC formulation is then focused on the mixing valves and the boiler, while the zone-level valves operate as per the hysteresis loop of the standard strategy. The inputs to the mixing valves are the set-points of the PI-controllers which control the temperature of the water flowing to the radiators on each floor, $(T_{fl_{sp_i}})$, while the input to the boiler is the input-power (P_{in}) .

5.4.3.1 Heating System Considerations and Nonlinearities

As in the quadratic formulation developed for the Nimbus Centre shown in (5.19) and (5.20), the heating system equations must be included in the new control strategy in the form of constraints. In addition to the zone model and mixing valve dynamics, the

delivery of heat from the boilers to the mixing valves through the header must also be considered. Generalizing once again to a building of F floors, in which the i^{th} floor consists of M_i zones, the heat supplied to the i^{th} floor from its associated mixing valve at time k is denoted $Q_{MV_i}(k)$. Using (3.47), the header temperature dynamics can be described by the difference equation (5.53), driven by the total heat flow into the header from the boiler Q_{Bo} and the total heat out of the header to each of the mixing valves. The constant β_h , represents the thermal capacity of the header.

$$T_h(k+1) = T_h(k) + \frac{1}{\beta_h} \left(Q_{Bo}(k) - \sum_{i=1}^F Q_{MV_i}(k) \right) \quad (5.53)$$

The relationship between the boiler input power P_{in} and the output from the boiler Q_{Bo} , where η denotes the boiler efficiency is given as:

$$P_{in}(k)\eta(k) = Q_{Bo}(k) \quad (5.54)$$

The boiler efficiency η , is often given by the manufacturers efficiency curves as a function of the power input and the return water temperature (which in this case is the header temperature T_h). In Fig. 5.16 for example, the relationship between Q_{Bo} and P_{in} is shown for two different header temperatures determined using efficiency curves given for the *Viessmann Virocrossal 200* series boiler.




Figure 5.16: Input and output power for different values of T_h for *Viessmann Virocrossal 200* series boiler

With this information, it is then possible to fit a nonlinear efficiency surface to data-points taken from these curves that relates Q_{Bo} to T_h and the input power P_{in} :

$$Q_{Bo}(k) = f(T_h(k), P_{in}(k)) \quad (5.55)$$

In Fig. 5.17 the derived efficiency surface showing the relationship between $P_{in}(t)$,

$Q_{Bo}(t)$ and $T_h(t)$ (taken from the efficiency curves of Fig. 5.16) can be seen.




Figure 5.17: Example boiler output as a function of boiler input & header temperature for Viessmann Virocrossal 200 series boiler

The heating system equations must be included in the control strategy in the form of constraints. An important consideration here however, is that the relationship between the boiler input and output power is not linear (5.55). Furthermore, the operating region of operation for a boiler is often discontinuous. Though it may be switched off, operating the boiler at less than a specific input power may not be permissible. The operating range is then given by the discontinuous interval $P_{in}(t) \in \{0 \cup [P_{min}, P_{max}]\}$. Approaches for solving nonlinear optimisation problems have been widely studied (Floudas 1995, Nocedal & Wright 1999), however, in many cases both the objective and constraints are assumed to be not only continuous, but twice-continuously differentiable (Forsgren et al. 2002, Wächter & Biegler 2006). Though methods for nonlinear mixed integer (MINLP) and combinatorial programming exist (Wolsey 1998), such problems are generally NP-hard (D'Ambrosio & Lodi 2013), becoming prohibitively complex as the number of variables and constraints increases.

To avoid discontinuity, the variable Q_{Bo} is approximated by the continuous function Q'_{Bo} , given by:

$$Q'_{Bo}(k) = \frac{1}{2} \left(1 + \tanh \left[\rho \left(\frac{P_{in}(k)}{P_{max}} - P_{min} \right) \right] \right) Q_{Bo}(k) \quad (5.56)$$

Though solutions within the region $(0, P_{min})$ are now possible, the tanh function renders such solutions less desirable. The constant ρ determines the *steepness* of the tanh function. For $\rho = 10$, the efficiency surface from Fig. 5.17 is now modified as shown in Fig. 5.18.




Figure 5.18: Boiler output as a function of boiler input & header temperature

5.4.3.2 Separating the Supply and Demand Problems

As the number of zones and the prediction horizon length increases, the optimisation problem can become quite large. Each zone has $2H$ variables (associated with the zone temperature deviation above and below the set-point) while there are H input variables associated with each floor of the building. Though the problem is now continuous, including (5.56) in the constraints of the two-level prioritised framework results in a pair of nonlinear problems that are not generally convex, with $H(2N + F)$ variables (where N is the total number of zones). While methods for solving such minimisation problems are available (Wächter & Biegler 2006), it is desirable first to reduce the number of constraints and variables involved if possible.

Here, the overall formulation is separated into a demand-side problem and a supply-side problem. The demand-side problem uses the lexicographic approach as described, with the heat input to each floor of the building (from the header) Q_{MV_i} used in place of P_{in} . The set-points $T_{fl_{sp_i}}$ to be sent to the mixing valves are deduced from the optimal choice $Q_{MV_i}^*$. The boiler is only considered in terms of the limits it puts on the available heat and so only the maximum output power Q_{Bo}^+ is required. The nonlinearities of (5.56) can then be omitted from the formulation.

Once a unique solution for the heat to be supplied to all floors of the building, $Q_{MV}^* = \sum_1^F Q_{MV_i}^*$, has been established, it is possible to ignore all constraints associated with the zones and just determine the minimum boiler input power P_{in} needed to supply Q_{MV}^* . This supply-side problem contains the nonlinear boiler dynamics, but only comprises H variables. The complete framework now contains a linear optimisation to minimise set-point deviation ($H(2N + F)$ variables), a quadratic optimisation to minimise demand-side energy ($H(2N + F)$ variables) and a nonlinear optimisation to min-

imise supply-side energy (H variables). The control hierarchy is shown in Fig. 5.19.

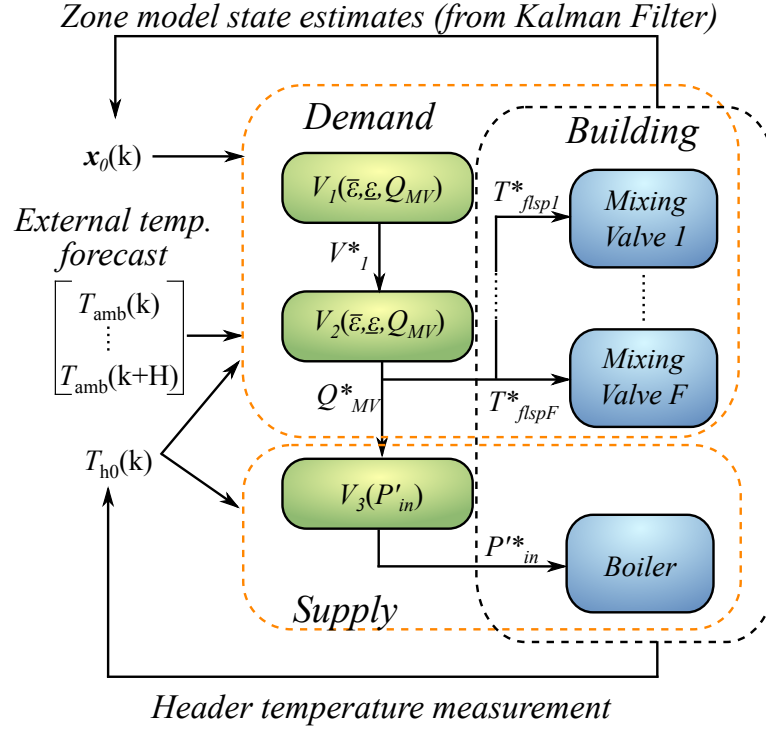


Figure 5.19: Three-layer control hierarchy

5.4.4 Formulation of the Prioritised-Objective MPC Strategy

This section outlines the full set of constraints and objectives for each layer of the hierarchy for a building of N zones and F floors, using the notation developed in Section 5.2.3.

5.4.4.1 Demand-Side Constraints

Using (5.19), the inequality constraints of (5.38) and (5.39) can be expressed in terms of the optimisation variables $\theta = [\bar{\epsilon}^T, \epsilon^T, Q^T_{MV_1}, \dots, Q^T_{MV_F}]^T$ as:

$$\begin{bmatrix} -I_{2NH} & \vdots & \bigoplus_{i=1}^F \Phi_i \Lambda_{3_i} \\ \vdots & \ddots & \vdots \\ \vdots & \vdots & -\bigoplus_{i=1}^F \Phi_i \Lambda_{3_i} \end{bmatrix} \theta \leq \begin{bmatrix} \bar{\mathbf{T}}_{sp} - \mathbf{y}_{f_1} - \Phi_1 \Lambda_{3_1} \Lambda_{1_1} \mathbf{y}_{f_1} \\ \vdots \\ \bar{\mathbf{T}}_{sp} - \mathbf{y}_{f_F} - \Phi_F \Lambda_{3_F} \Lambda_{1_F} \mathbf{y}_{f_F} \\ \mathbf{y}_{f_1} + \Phi_1 \Lambda_{3_1} \Lambda_{1_1} \mathbf{y}_{f_1} - \underline{\mathbf{T}}_{sp} \\ \vdots \\ \mathbf{y}_{f_F} + \Phi_F \Lambda_{3_F} \Lambda_{1_F} \mathbf{y}_{f_F} - \underline{\mathbf{T}}_{sp} \end{bmatrix} \quad (5.57)$$

As the radiator flow-temperature is controlled by mixing water from the header with return water from the radiators, the set-points, $\mathbf{T}_{fl_{sp_i}}$, sent to the mixing valves must be greater than or equal to the zone temperatures \mathbf{T}_z and less than or equal to the header temperature \mathbf{T}_h . The former inequality constraint can be found in terms of θ using (5.8) and (5.16) as:

$$\begin{bmatrix} 0_{NH \times 2NH} & -\bigoplus_{i=1}^F \Lambda_{2_i} \Lambda_{3_i} \end{bmatrix} \theta \leq \begin{bmatrix} \Lambda_{2_1} \Lambda_{3_1} \Lambda_{1_1} \mathbf{y}_{f_1} - \mathbf{y}_{f_1} \\ \vdots \\ \Lambda_{2_F} \Lambda_{3_F} \Lambda_{1_F} \mathbf{y}_{f_F} - \mathbf{y}_{f_F} \end{bmatrix} \quad (5.58)$$

To ensure that enough power can be provided by the boiler when operating at maximum power Q_{Bo}^+ to keep \mathbf{T}_h higher than $\mathbf{T}_{fl_{sp}}$ and \mathbf{T}_h^- (the minimum permissible header temperature) over the horizon, the following two constraints are introduced, using 5.53 and (5.16) where T_{h0} is the measured header temperature at time k :

$$\begin{bmatrix} 0_{H(F+1) \times 2NH} & \vdots & \mathbf{1}_F \mathbf{1}_F^T \otimes \Gamma + \bigoplus_{i=1}^F \Lambda_{3_i} \\ \vdots & \ddots & \vdots \\ \vdots & \vdots & \mathbf{1}_F^T \otimes \Gamma \end{bmatrix} \theta \leq \begin{bmatrix} T_{h0} \mathbf{1}_H + \Gamma Q_{Bo}^+ \mathbf{1}_H - \Lambda_{3_1} \Lambda_{1_1} \mathbf{y}_{f_1} \\ \vdots \\ T_{h0} \mathbf{1}_H + \Gamma Q_{Bo}^+ \mathbf{1}_H - \Lambda_{3_F} \Lambda_{1_F} \mathbf{y}_{f_F} \\ \hline (T_{h0} - T_h^-) \mathbf{1}_H + \Gamma Q_{Bo}^+ \mathbf{1}_H \end{bmatrix} \quad (5.59)$$

where $\Gamma \in \mathbb{R}^{H \times H}$ is given by:

$$\Gamma_{i,j} = \begin{cases} \frac{1}{\beta_h}, & \text{for } i = j \\ 0, & \text{otherwise} \end{cases}$$

Finally, $T_{fl_{sp_i}}$ must be chosen to be lower than T_{fl}^+ , the maximum permissible header

temperature. Using (5.53), this constraint can be represented as:

$$\left[\begin{array}{c|c} 0_{FH \times 2NH} & \bigoplus_{i=1}^F \Lambda_{3_i} \end{array} \right] \theta \leq \begin{bmatrix} T_h^+ \mathbf{1}_H - \Lambda_{3_1} \Lambda_{1_1} \mathbf{y}_{f_1} \\ \vdots \\ T_h^+ \mathbf{1}_H - \Lambda_{3_F} \Lambda_{1_F} \mathbf{y}_{f_F} \end{bmatrix} \quad (5.60)$$

The linear cost function and full set of constraints for the primary objective can then be expressed as:

$$V_1^* = \min_{\theta \in \Theta} \Phi_1 \theta \quad (5.61)$$

s.t.

$$\Omega \theta \leq \omega \quad (5.62)$$

$$\theta \geq 0, \quad (5.63)$$

where the inequality in (5.62) is taken from the set of inequalities (5.57)-(5.60) and $\Phi_1 = [\mathbf{1}_{NH}^T \quad \mathbf{1}_{NH}^T \quad 0_{1 \times FH}]$. The secondary (quadratic) formulation is then:

$$V_2^* = \min_{\theta \in \Theta} \theta^T \Phi_2 \theta \quad (5.64)$$

s.t.

$$\begin{bmatrix} \Phi_1 \\ \Omega \end{bmatrix} \theta \leq \begin{bmatrix} V_1^* \\ \omega \end{bmatrix} \quad (5.65)$$

$$\theta \geq 0, \quad (5.66)$$

where

$$\Phi_2 = \left[\begin{array}{c|c} 0_{2NH \times 2NH} & 0_{2NH \times FH} \\ \hline 0_{FH \times 2NH} & I_{FH \times FH} \end{array} \right]$$

The flow-temperature set-point $T_{fl_{sp_i}}$ to be sent to the mixing valve PI controller for the i^{th} floor can then be found using (5.16) from the optimal solution for $Q_{MV_i}^*$ by:

$$T_{fl_{sp_i}} = \Lambda_{3_i} (Q_{MV_i}^* + \Lambda_{1_i} \mathbf{y}_{f_i}) \quad (5.67)$$

$$\forall i = 1, 2 \dots, F$$

5.4.4.2 Supply-Side Constraints

Having determined a unique solution for the heat energy leaving the header $\mathbf{Q}_{MV}^* = \sum_{i=1}^F \mathbf{Q}_{MV_i}$, the constraints associated with zones can be removed when considering the supply-side problem of minimising the accumulated boiler input power, given by:

$$(\mathbf{P}_{in}(k) = [P_{in}(k) \cdots P_{in}(k + H - 1)]) \quad (5.68)$$

The variables to be considered are now P_{in} , \mathbf{Q}'_{Bo} and \mathbf{T}_h , the latter of which can be expressed in terms of \mathbf{Q}'_{Bo} using (5.53) as:

$$\mathbf{T}_h = T_{h0} \mathbf{1}_H + \Gamma (\mathbf{Q}'_{Bo} - \mathbf{Q}_{MV}^*) \quad (5.69)$$

In this problem, \mathbf{T}_h must remain within the bounds of T_h^+ and T_h^- and above \mathbf{T}_{flsp} . The optimisation problem can be expressed as:

$$\mathbf{P}_{in}^* = \arg \min_{\mathbf{P}_{in}, \mathbf{Q}'_{Bo}} \sum_{k=1}^H P_{in}(k) \quad (5.70)$$

subject to the set of linear constraints is given by:

$$\begin{bmatrix} \Gamma \\ -\Gamma \\ -(\mathbf{1}_F \otimes \Gamma) \end{bmatrix} \mathbf{Q}'_{Bo} \leq \begin{bmatrix} (T_h^+ - T_{h0}) \mathbf{1}_H + \Gamma \mathbf{Q}_{MV}^* \\ (T_{h0} - T_h^-) \mathbf{1}_H - \Gamma \mathbf{Q}_{MV}^* \\ T_{h0} \mathbf{1}_H - \Gamma \mathbf{Q}_{MV}^* - \mathbf{T}_{flsp1} \\ \vdots \\ T_{h0} \mathbf{1}_H - \Gamma \mathbf{Q}_{MV}^* - \mathbf{T}_{flspF}^* \end{bmatrix} \quad (5.71)$$

$$0 \leq \mathbf{P}_{in} \leq P_{in}^+ \mathbf{1}_H, \quad (5.72)$$

An additional nonlinear equality constraint is taken from (5.55)-(5.56) and (5.69):

$$\mathbf{Q}'_{Bo} = f(\mathbf{P}_{in}, T_{h0}, \mathbf{Q}_{MV}^*) \quad (5.73)$$

5.4.5 Performance of Prioritised Objective Strategy in Simulation Platform

To assess the performance and usability, the prioritised-objective strategy was implemented with the validated model of the Nimbus Building and compared with the standard weather-compensation in terms of comfort and energy. The primary optimisation

problem was solved using the *linprog* function in Matlab, the secondary was solved using the *quadprog* function, while the supply optimisation problem was solved using the IPOPT algorithm.

Simulating a period of 10 days, a sample time of 600 seconds was used with a prediction horizon of 10 samples. The prediction horizon was chosen to reflect the observation made when using the simulation platform that, in typical conditions when the heating system was turned on, 100 minutes was an adequate time-scale to achieve comfort requirements in most zones. Longer horizons only served to increase complexity. A comfort band of $\pm 1^\circ\text{C}$ was used for the demand-side optimisation problem. An external temperature profile was taken from measured and recorded data from the Nimbus Centre from February 2015, while unmeasured disturbances were also generated and included in the platform to represent typical solar gain, internal occupancy and equipment gain profiles.

Fig. 5.20 shows the accumulated energy input to the boiler for both strategies, while Fig. 5.21 shows the accumulated absolute deviation of the zone temperatures from the comfort set-point band in units of $^\circ\text{C}.h$. Using the prioritised MPC, energy consumption was reduced by 7.1% while deviation from the set-point band was reduced by 29.2%. This performance is dependent on the set-point band used for the MPC approach - additional energy savings could be made with a possible compromise of set-point deviation. These results were chosen to show that both energy and comfort can be simultaneously improved by using the predictive strategy.




Figure 5.20: Energy consumption for the Nimbus Centre simulation platform using prioritised MPC vs. Weather Compensation (10 days, Feb 2015)




Figure 5.21: Temperature set-point deviation in the Nimbus Centre simulation platform using prioritised MPC vs. Weather Compensation (10 days, Feb 2015)

5.4.6 Load Flexibility - Choosing the Acceptable Level of Set-point Deviation

In the example shown, an improved performance was achieved when compared to standard approaches. It should be noted again that such an improvement is not uncommon, with many examples in literature highlighting the benefits of an optimal control strategy for the application of building energy (a review can be found in (Shaikh et al. 2014)). The more significant outcome is that the performance was achieved without the need for a large-scale tuning parameter selection process, the solution to which may seem almost arbitrary. In this strategy, a user can subjectively ascribe an acceptable level of discomfort to the strategy by adjusting the zone temperature set-point band. The strategy will then attempt to satisfy this comfort requirement with as little energy as possible. This is a far more tangible adjustment, as it only concerns room temperatures (in $^{\circ}\text{C}$).

While increasing the set-point band will have the effect of reducing energy consumption, the exact relationship (knowing how much energy will be saved by increasing the band by 1°C for example) will be dependent on measured and unmeasured disturbances. The potential to make energy savings by adjusting the set-point band is clear however. To illustrate this, the test scenario of the previous example was repeated using the same MPC formulation, but with different set-point bands.

In Fig. 5.22, the energy consumption is plotted for comfort bands of 0.5°C , 1°C , 2°C , 3°C and 4°C . It can be seen for example, that over the 10 day period, a further saving of 537kWh (or 9.5%) would be achieved if the acceptable temperature deviation was increased from 1°C to 3°C . The internal temperatures that would result for each setting

are shown for the first-floor administration office in Fig. 5.23.




Figure 5.22: Energy consumption change as acceptable comfort range is increased using prioritised model predictive control strategy




Figure 5.23: Simulated zone temperature obtained using weather-compensation and prioritised MPC for 10 day period (first-floor administration office of Nimbus Centre)

To assist in the set-point selection, the predicted load profile can also be observed. Using the full prediction horizon of the optimisation problem, a projection of the energy use can be obtained for a given period of time. By projecting the load profile for different levels of acceptable set-point deviation, an estimate can be made of the possible energy savings.

An example is shown in Fig. 5.24, in which a one day projection of energy use is shown for different comfort band settings. Each of the projections was obtained by increasing the prediction horizon of the control problem to 12 hours and observing the full sequence of predicted inputs, calculated at the first time-step.




Figure 5.24: Projected energy consumption over period of 11.5 hours for different acceptable comfort ranges

In this example, it can be seen that a 21% energy reduction was predicted for the day by increasing the acceptable temperature deviation from 1°C to 3°C . Furthermore, little advantage in terms of energy consumption was predicted to be achievable by increasing the comfort band from 1°C to 2°C . This results from the fact that when using a wider comfort band, both over-heating and under-heating in the zones may be increased - the former acts to increase the energy consumption, while the latter acts to reduce it. A direct relationship between energy use and comfort satisfaction cannot then be assumed.

The validity of the projected energy use is dependent on the prediction accuracy. As the zone prediction models are selected to provide accuracy over a shorter prediction horizon, it is likely that the true daily energy use will deviate from the projected value. Additionally, disturbances will affect the true values. Nonetheless, the comparison between the different set-points can be instructive when selecting an appropriate set-point band.

5.5 Conclusion

In this chapter, different predictive control formulations, applicable to building heating systems were analysed. Suitable objectives associated with this application were outlined, with the PMV and PPD indices given as a possible method of quantifying comfort satisfaction. A quadratic cost function, typically employed to allow for a reduction in both energy consumption and zone temperature set-point deviation was implemented using the simulation platform developed in Chapter 3 to represent the Nimbus Centre. The zone models developed using methods from Chapter 4 were in-

corporated in the formulation to predict the future zone temperature outputs. It was shown that the performance of the formulation in terms of comfort satisfaction and energy reduction is dependent on the weights used in the cost function. Selection of these weights to achieve some desired performance is not a straightforward process.

To handle complexity, both in terms of the optimisation problem and the tuning procedure, different MPC formulations were introduced. Using a hierarchical approach in which only the hot and cold extremities of the problem were considered in the optimisation, the comfort and energy performance was shown to be better than a traditional weather-compensation based strategy. A guideline for tuning quadratic cost functions was also introduced by observing the unconstrained solution to the KKT conditions of the optimisation problem.

A prioritised-objective formulation was then developed whereby an acceptable temperature band for comfort satisfaction was defined for all zones. A linear optimisation determines the minimum cumulative deviation outside this band over a prediction horizon. The solution to this is used as a constraint in a second, quadratic, optimisation problem, which seeks to determine the minimum energy required to achieve this level of comfort. The cost functions of the optimisation problems do not require tuning parameters.

A methodology was also introduced whereby the full optimisation problem is further broken up into demand and supply problems. The demand problem consists of a large number of linear constraints (and variables), solved using a lexicographic approach, while the supply problem consists of a low number of nonlinear constraints. The performance of the prioritised MPC approach compared favourably to weather-compensation in terms of energy consumption and set-point deviation, without the need for tuning parameter selection. The ability to adjust the strategy using tangible metrics was also outlined by using projections of the predicted energy requirement needed to achieve an acceptable level of comfort over a set period of time.

References

- Brooks, J., Kumar, S., Goyal, S., Subramany, R. & Barooah, P. (2015), ‘Energy-efficient control of under-actuated HVAC zones in commercial buildings’, *Energy and Buildings* **93**, 160–168.
- Camacho, E. & Alba, C. (2007), *Model predictive control*, second edn, Springer-Verlag, London.
- Camponogara, E., Jia, D., Krogh, B. H. & Talukdar, S. (2002), ‘Distributed model predictive control’, *Control Systems, IEEE* **22**(February), 44–52.
- Castilla, M., Álvarez, J., Berenguel, M., Rodríguez, F., Guzmán, J. & Pérez, M. (2011), ‘A comparison of thermal comfort predictive control strategies’, *Energy and Buildings* **43**(10), 2737–2746.
- Chandan, V. & Alleyne, A. (2013), ‘Optimal partitioning for the decentralized thermal control of buildings’, *IEEE Transactions on Control Systems Technology* **21**(5), 1756–1770.
- Chandan, V. & Alleyne, A. G. (2014), ‘Decentralized predictive thermal control for buildings’, *Journal of Process Control* **24**(6), 820–835.
- Cigler, J. & Privara, S. (2010), ‘Subspace identification and model predictive control for buildings’, *2010 11th International Conference on Control Automation Robotics & Vision* (December), 750–755.
- D’Ambrosio, C. & Lodi, A. (2013), ‘Mixed integer nonlinear programming tools: an updated practical overview’, *Annals of Operations Research* **204**(1), 301–320.
- Ferreira, P. M., Ruano, a. E., Silva, S. & Conceição, E. Z. E. (2012), ‘Neural networks based predictive control for thermal comfort and energy savings in public buildings’, *Energy and Buildings* **55**, 238–251.
- Floudas, C. a. (1995), ‘Nonlinear and Mixed-Integer Optimization’, *Handbook of Applied Optimization* p. 462.

- Forsgren, A., Gill, P. & Wright, M. (2002), 'Interior methods for nonlinear optimization', *SIAM review* **44**(4), 525–597.
- Freire, R. Z., Oliveira, G. H. C. & Mendes, N. (2008), 'Predictive controllers for thermal comfort optimization and energy savings', *Energy and Buildings* **40**(7), 1353–1365.
- Gambier, A. & Badreddin, E. (2007), Multi-Objective Optimal Control: An Overview, in 'IEEE International Conference on Control Applications', number October, pp. 170–175.
- Goyal, S., Ingley, H. a. & Barooah, P. (2013), 'Occupancy-based zone-climate control for energy-efficient buildings: Complexity vs. performance', *Applied Energy* **106**, 209–221.
- Hazyuk, I., Ghiaus, C. & Penhouet, D. (2012), 'Optimal temperature control of intermittently heated buildings using Model Predictive Control: Part II - Control algorithm', *Building and Environment* **51**, 388–394.
- Humphreys, M. A. & Fergus Nicol, J. (2002), 'The validity of ISO-PMV for predicting comfort votes in every-day thermal environments', *Energy and Buildings* **34**(6), 667–684.
- Kerrigan, E. & Maciejowski, J. (2002), 'Designing model predictive controllers with prioritised constraints and objectives', *IEEE International Symposium on Computer Aided Control System Design* pp. 33–38.
- Lazos, D., Sproul, A. B. & Kay, M. (2014), 'Optimisation of energy management in commercial buildings with weather forecasting inputs: A review', *Renewable and Sustainable Energy Reviews* **39**, 587–603.
- Lefort, A., Bourdais, R., Ansanay-Alex, G. & Guéguen, H. (2013), 'Hierarchical control method applied to energy management of a residential house', *Energy and Buildings* **64**, 53–61.
- Liao, Z. & Dexter, A. (2010), 'An Inferential Model-Based Predictive Control Scheme for Optimizing the Operation of Boilers in Building Space-Heating Systems', *IEEE Transactions on Control Systems Technology* **18**(5), 1092–1102.
- Ma, J., Qin, J., Salsbury, T. & Xu, P. (2012), 'Demand reduction in building energy systems based on economic model predictive control', *Chemical Engineering Science* **67**(1), 92–100.
- Ma, Y., Anderson, G. & Borrelli, F. (2011), A Distributed Predictive Control Approach

- to Building Temperature Regulation, in 'American Control Conference', pp. 2089–2094.
- Ma, Y., Borrelli, F., Hancey, B., Coffey, B., Bengea, S. & Haves, P. (2012), 'Model Predictive Control for the Operation of Building Cooling Systems', *IEEE Transactions on Control Systems Technology* **20**(3), 796–803.
- Maciejowski, J. (2002), *Predictive Control: With Constraints*, Prentice-Hall, UK.
- Mendoza-Serrano, D. I. & Chmielewski, D. J. (2012), HVAC control using infinite-horizon economic MPC, in 'Proceedings of the IEEE Conference on Decision and Control', Vol. 60616, pp. 6963–6968.
- Moroşan, P.-D., Bourdais, R., Dumur, D. & Buisson, J. (2010a), 'Building temperature regulation using a distributed model predictive control', *Energy and Buildings* **42**(9), 1445–1452.
- Moroşan, P. D., Bourdais, R., Dumur, D. & Buisson, J. (2010b), Distributed model predictive control based on Benders' decomposition applied to multisource multizone building temperature regulation, in 'Proceedings of the IEEE Conference on Decision and Control', IETR-SUPELEC, Avenue de la Boulaie - B.P. 81127, F-35511 Cesson-S; Cedex, France, Ieee, pp. 3914–3919.
- Nocedal, J. & Wright, S. (1999), *Numerical Optimization*, 2 edn, Springer-Verlag.
- O'Dwyer, E., Cychowski, M., Kouramas, K. & Lightbody, G. (2014), A hierarchical Model-based Predictive Control strategy for building heating Systems, in '25th IET Irish Signals & Systems Conference 2014', Limerick, pp. 298–303.
- O'Dwyer, E., Cychowski, M., Kouramas, K., Tommasi, L. D. & Lightbody, G. (2015), Scalable , Reconfigurable Model Predictive Control for Building Heating Systems, in 'European Control Conference 2015 (ECC)', pp. 2253–2258.
- O'Dwyer, E., Tommasi, L. D., Kouramas, K., Cychowski, M. & Lightbody, G. (2016), 'Prioritised Objectives for Model Predictive Control of Building Heating Systems', *IEEE Transactions on Control Systems Technology* p. (Under Review).
- Oldewurtel, F., Sturzenegger, D. & Morari, M. (2013), 'Importance of occupancy information for building climate control', *Applied Energy* **101**, 521–532.
- Oldewurtel, F., Ulbig, A., Parisio, A., Andersson, G. & Morari, M. (2010), 'Reducing peak electricity demand in building climate control using real-time pricing and model predictive control', *49th IEEE Conference on Decision and Control (CDC)* pp. 1927–1932.

- Peeters, L., Van der Veken, J., Hens, H., Helsen, L. & D'haeseleer, W. (2008), 'Control of heating systems in residential buildings: Current practice', *Energy and Buildings* **40**, 1446–1455.
- Shaikh, P. H., Nor, N. B. M., Nallagownden, P., Elamvazuthi, I. & Ibrahim, T. (2014), 'A review on optimized control systems for building energy and comfort management of smart sustainable buildings', *Renewable and Sustainable Energy Reviews* **34**, 409–429.
- Široký, J., Oldewurtel, F., Cigler, J. & Prívara, S. (2011), 'Experimental analysis of model predictive control for an energy efficient building heating system', *Applied Energy* **88**(9), 3079–3087.
- Stewart, B. T., Venkat, A. N., Rawlings, J. B., Wright, S. J. & Pannocchia, G. (2010), 'Cooperative distributed model predictive control', *Systems and Control Letters* **59**, 460–469.
- Sturzenegger, D., Gyalistras, D., Morari, M. & Smith, R. S. (2015), 'Model Predictive Climate Control of a Swiss Office Building : Implementation , Results , and Cost – Benefit Analysis', *IEEE Transactions on Control Systems Technology* pp. 1–12.
- Wächter, A. & Biegler, L. (2006), 'On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming', *Mathematical programming* **106**(1), 25–57.
- Wolsey, L. A. (1998), *Integer Programming*, Vol. 98, John Wiley & Sons, Inc., New York.
- Yang, R. & Wang, L. (2012), 'Multi-objective optimization for decision-making of energy and comfort management in building automation and control', *Sustainable Cities and Society* **2**(1), 1–7.

Chapter 6

Inaccurate Predictions and System Faults in the MPC Formulation

6.1 Introduction

Previously, the performance of the MPC formulations were considered under conditions in which the controlled building was assumed to follow expected operating patterns, with events unfolding in a predictable fashion. In this chapter the impact of future uncertainties caused by inaccurate models and component failures in a building energy MPC formulation is considered. The effect of modelling inaccuracy is first examined as well as the effect of incorrect weather forecasts. Different methods for predicting unmeasured disturbances based on previously estimated values are then compared. Finally, faults and unexpected system changes are studied, both in terms of characterisation of the fault and reconfiguration of the control to achieve a desired performance level post-fault.

While the performance of predictive control strategies can exceed that of more traditional empirical approaches in the application of building energy systems, the performance of any predictive approach is dependent on the validity of the predictions. The impact of disturbances and uncertainties in the strategy is analysed here by observing the effect of modelling error as well as disturbance prediction error. The quadratic MPC formulation of Chapter 5 is implemented using lower-accuracy zone models to determine the effect on the optimal solution derived for different prediction horizon lengths. Uncertainties are then included in the external temperature forecast followed by unmeasured disturbances. Different methods of handling the uncertainties are introduced and evaluated by comparing the similarity of the solutions to the solution

obtained without prediction uncertainties present.

Given the large number of sensors and actuators present in the building, the occurrence of component failures and system changes must also be considered. If the effects of such failures are not incorporated in the predictions, the performance may be greatly reduced. Ideally, if a fault occurs, the plant should be able to operate to some specified level of performance. To achieve fault tolerance, the fault or system change must first be detected and identified. A model-free approach using principal component analysis (PCA) for sensor-fault detection is initially analysed using the simulation platform. Following this, a model-based methodology using Kalman filtering and simple machine-learning techniques is implemented to detect and identify mixing valve and zone-level valve faults.

The reconfiguration of the control strategy after a fault has been identified is then analysed. It is shown that the prioritised-objective MPC formulation developed in Chapter 5 can provide a simpler means of reconfiguration when compared to a more standard quadratic formulation. This is due to the lack of tuning parameters as well as the greater ease with which the pre-fault and post-fault performance can be characterised.

6.2 Analysing the Impact of Inaccurate Predictions

When considering an MPC-based control strategy for a building heating system, different types of uncertainty must be taken into account (Maasoumy et al. 2014, Maasoumy & Sangiovanni-Vincentelli 2012). Low-order zone models can only approximate the thermal properties of the building, while disturbances will also have an effect on the validity of any predictions made. Many of these disturbances may not be easily predictable. Those for which forecasts are available will inevitably still contain a certain amount of inaccuracy.

If the uncertainty can be in some way characterised, methods can be incorporated in the MPC strategy to handle it. In (Oldewurtel et al. 2012, 2014) for example, prediction uncertainty is acknowledged and accounted for by the use of stochastic MPC (SMPC) with chance constraints. In place of a strict comfort constraint, temperatures in the building are only required to satisfy constraints with a given probability. In this section, the impact of these uncertainties is analysed.

6.2.1 Effect of Zone Model Inaccuracies on Optimisation

In Section 5.2.3, the standard QP-type formulation was implemented in the simulation platform using different prediction horizons. It was shown that by increasing the horizon length, set-point changes were anticipated earlier. The longer horizon results in a greater optimisation complexity however, with the presence of more variables and constraints. An additional aspect to be considered when selecting the horizon length is the prediction accuracy. Having an estimate of the future plant states is only useful if the estimate is valid. If the predicted plant behaviour drifts significantly from the reality, a longer prediction horizon may negatively impact the control performance. This is noted in (Negenborn et al. 2008) for example, in which a strategy is implemented with a sample time that increases further into the future. The lower resolution reflects the reduced confidence in the prediction.

In Chapter 4, different methodologies were considered for the development of low-order models for use in a predictive control strategy, compared in terms of prediction accuracy. The impact of inaccuracies resulting from these zone models is now investigated. To assess the impact of zone-model uncertainty in a standard MPC strategy, the quadratic formulation of Section 5.2.3 in Chapter 5 in which no disturbances were present is considered once again.

The performance aspect considered here is the optimal input trajectory. The strategy is implemented using a set of high-accuracy models, while a second optimisation in which an alternative sequence of inputs is calculated in parallel (but not applied) at each sample using a set of lower accuracy zone models. Running the alternate optimisation in open loop allows for the the optimisation problem posed at each sample to be the same for the high-accuracy and low-accuracy models. The optimal input trajectory calculated with the high-accuracy models is used as the target to which the trajectories obtained using the low-accuracy set of models are compared.

The first set of zone models used was derived from disturbance-free data taken from the simulation platform. The average coefficient of determination (R^2) of the 72-step ahead prediction (the accuracy metric used in the previous chapter) determined for the full set of zone models was 0.9546. The solution obtained using these high-accuracy models is considered here to be the target solution to the problem. A second set of models was derived using data from the simulation platform in which a high level of disturbance was present and no disturbance estimation was included in the identification process. The average R^2 value found for the 72-step ahead prediction using this set of zone models was 0.7263.

Control of the system was carried out as before using the high-accuracy zone model-set with prediction horizons of 5, 10 and 20 samples. The control sequence determined using the low-accuracy zone model-set within the same formulation was calculated in parallel using measurements from the plant (but not applied). At each sample, the control trajectories determined using the low-accuracy model set are compared to those determined using the high-accuracy set for different prediction horizons.

In Fig. 6.1 the first-floor radiator flow-temperature input, $T_{fl_{sp_{ff}}}$, applied to the system (determined using the high-accuracy models with a horizon of 20 steps) is shown, as well as the input that would have been applied at each sample if the low-accuracy zone models were used $T'_{fl_{sp_{ff}}}$. It can be seen that the heating requirement was underestimated when the low-accuracy zone models were used.




Figure 6.1: Optimal input determined using zone models of different accuracy

At sample k , the difference between the projected inputs for each step in the horizon, determined with low-accuracy and high-accuracy zone models, was calculated. The rms value for each point in the horizon, calculated for the full set of N_{sim} samples was then found for both floors (with the subscript gf representing the ground-floor and the subscript ff representing the first-floor):

$$e_{T_{fl_{gf}}}(k+i|k) = T_{fl_{sp_{gf}}}(k+i|k) - T'_{fl_{sp_{gf}}}(k+i|k) \quad (6.1)$$

$$e_{T_{fl_{ff}}}(k+i|k) = T_{fl_{sp_{ff}}}(k+i|k) - T'_{fl_{sp_{ff}}}(k+i|k) \quad (6.2)$$

$$e_{rms}(i) = \sqrt{\frac{1}{2N_{sim}} \sum_{k=1}^{N_{sim}} (e_{T_{fl_{gf}}}(k+i|k)^2 + e_{T_{fl_{ff}}}(k+i|k)^2)} \quad (6.3)$$

$$\forall i = 1, 2, \dots, H$$

In Fig. 6.2, the calculated value of e_{rms} is shown for each point along the horizon,

for strategies with prediction horizon lengths of 5, 10 and 20 samples. When a longer prediction horizon is used, a larger deviation from the target can be seen, particularly for predictions at the start of the horizon. This is to be expected. Over shorter horizons, there is simply less time for the inaccurate predictions to deviate from the target. While longer prediction horizons may improve performance if model accuracy is high, the tendency for inaccuracies to be accentuated must be considered when selecting an appropriate horizon length.




Figure 6.2: Difference between the inputs calculated with low-accuracy and with high-accuracy models at each point of the horizon

As the first element of the optimal input sequence is of primary importance here (as this is the only element applied to the system), a more in-depth analysis of the impact on model accuracy on the first element is next carried out. In this case, the analytic solution to the QP problem with input constraints removed was determined for the set of ground floor zones (the QP problem once again seeks to find an optimal sequence of mixing valve flow temperature set-points). The solution (as outlined in (Maciejowski 2002)) for these inputs is given by:

$$\mathbf{T}_{fl_{sp1}}^* = \left(\Phi_1^T \mathbf{Q} \Phi_1 + \mathcal{R} \right)^{-1} \Phi_1^T \mathbf{Q} (\Psi_1 \mathbf{x}_{0_1} + \varphi_1 \mathbf{T}_e) \quad (6.4)$$

where the zone models are given by the relationship outlined using the notation of Chapter 5 as:

$$\mathbf{T}_{z_1}(k) = \Psi_1 \mathbf{x}_1(k) + \Phi_1 \mathbf{T}_{fl_{sp1}}(k) + \varphi_1 \mathbf{T}_e \quad (6.5)$$

For this analysis, the predicted vector of external temperatures \mathbf{T}_e is set as a constant temperature of 15°C . The first element of the solution is determined using the

high-accuracy models and the low-accuracy models for horizon lengths ranging from 1 sample (5 minutes) to 100 samples (500 minutes). In Fig. 6.3, these solutions are plotted, with the horizon length on the x-axis and the first element of the optimal sequence of flow temperatures shown on the y-axis. It can be seen that the solutions begin to deviate for horizon lengths of greater than 15 minutes, with a significant difference noticeable above 100 minutes.




Figure 6.3: First element of optimal solution to unconstrained QP problem for different horizons with high-accuracy and low-accuracy models

To illustrate the effect of the inaccurate models and the resulting inputs on the output of the system, the low-accuracy optimisation was implemented in the simulation platform in closed loop with a prediction horizon of 20 samples. The simulated zone temperature of the ground-floor meeting room is shown in Fig. 6.4, compared to the output obtained when the higher accuracy models are used. It can be seen that, though the inputs are different in each case, the resulting output does not deviate significantly. This highlights the importance of the receding-horizon aspect of MPC. Though the predictions may deviate from the reality, the models are prevented from drifting as they are updated with measurements from the plant at each sample.

6.2.2 Forecasting Disturbances in the Formulation

In the previous sections, the control strategy was implemented without the inclusion of unknown disturbances. Furthermore, the forecast of the measured disturbance that was included - the external temperature - was assumed to be correct. In this section, the impact of events unaccounted for at the time of prediction is considered.




Figure 6.4: Difference between ground-floor meeting room output determined using low-accuracy and high-accuracy zone models within MPC formulation

6.2.2.1 Inaccuracies in Measured Disturbance Forecast

The ability to exploit knowledge of the future ambient temperature values is an important benefit to a building energy predictive control strategy. A predicted drop in external temperature for example, can be accounted for in advance of its occurrence as opposed to traditional approaches, which can only react to changes that have already happened. Given the slow thermal dynamics of a building, a reactive strategy may result in extended periods of time for which thermal comfort is affected or energy is unnecessarily consumed.

Including the external temperature forecast may not be beneficial however, if it is inaccurate. The Irish meteorological service (Met Éireann) provide external weather forecasts for a 24 hour period of approximately 88% average accuracy for the maximum predicted temperature and 81% average accuracy for the minimum predicted temperature whereby an individual temperature prediction is deemed to be accurate if it is within $\pm 2^{\circ}\text{C}$ of the observed temperature (*Forecast Accuracy - Met Éireann - The Irish Meteorological Service Online* 2016). This uncertainty may need to be considered in the strategy.

To do so, the external temperature forecast used for prediction was changed so as to be different to the actual external temperature profile applied to the simulation platform. The forecast profile T_{ef} was generated by adding an integrated Gaussian noise term as

follows (for a set of N_{sim} samples):

$$T_{ef}(k) = T_e(k) + w_e(k) \quad (6.6)$$

$$w_e(k) = w_e(k-1) + \xi_e(k) \quad (6.7)$$

$$\forall k = 1, \dots, N_{sim}; \quad w_e(0) = 0; \quad \xi_e \sim \mathcal{N}(0, 0.01)$$

At each sample, the forecast was then corrected using the measured value. The predicted value i steps into the future at time k is given as:

$$T'_{ef}(k+i) = T_{ef}(k+i) - (T_{ef}(k) - T_e(k)) \quad (6.8)$$

The external temperature profile, T_e and the initial temperature forecast, T_{ef} are shown in Fig 6.5, as well as the corrected forecast, T'_{ef} , that would be obtained at time $k = 100$.




Figure 6.5: External temperature measurement and forecast for 66 hour period, as well as the corrected forecast obtained at time $k = 100$

The simulated experiment was again repeated for prediction horizon lengths of 5, 10 and 20 samples, this time calculating, in parallel, a set of optimal inputs using an incorrect external temperature forecast.

The rms of the projected difference in the input sequences calculated with perfect and imperfect external temperature predictions is shown in Fig. 6.6. Once again, the impact was more pronounced when a longer horizon was used. The deviations are relatively small however, when compared to those observed when inaccurate zone models were used. For the time-scales considered here, inaccuracies in the forecast do not appear to have a significant effect.




Figure 6.6: Difference between the inputs calculated with correct and with incorrect weather forecast information at each point of the horizon

6.2.2.2 Prediction of the Unmeasured Disturbances for Improved Control Performance

The impact of unmeasured disturbances on the control performance is next considered. In Chapter 3, a set of disturbance profiles were generated to replicate the impact of solar gains, equipment gains and occupancy gains on the building. These are now applied to the simulation platform.

At each sample k , a disturbance state, $\hat{x}_{d_j}(k)$, associated with the j^{th} zone was estimated by the Kalman filtering process outlined in Section 4.5. Though at each sample, estimates of the current and past disturbances are known, it must be assumed that no knowledge of the future disturbances is available. Different ways are introduced here for including the disturbance estimate in the predictive strategy, though other approaches are possible (Lazos et al. 2014, Pawlowski et al. 2010).

Option 1:

Most simplistically, the disturbance state could be assumed to be zero for the prediction horizon:

$$\hat{x}_{d_j}(k+i) = 0 \quad \forall i = 1, \dots, H \quad (6.9)$$

Option 2:

The disturbance term could be assumed to be constant for the horizon, equal to the

current estimate:

$$\hat{x}_{d_j}(k+i) = \hat{x}_{d_j}(k) \quad \forall i = 1, \dots, H \quad (6.10)$$

Option 3:

Time-series approaches to forecasting solar radiation are evaluated in (Pawlowski et al. 2010). One such approach is to project ahead an exponentially weighted moving average (EWMA) of previous disturbance estimates (Chilin et al. 2012). Where N_d is the sample window size, the i^{th} step ahead prediction is given as:

$$\hat{x}_{d_j}(k+i) = \frac{\hat{x}_{d_j}(k+i-1) + \dots + (1-\beta_d)^{N_d} \hat{x}_{d_j}(k+i-N_d-1)}{1 + (1-\beta_d) + \dots + (1-\beta_d)^{N_d}} \quad (6.11)$$

where $\beta_d \in (0, 1)$ is a parameter chosen to increase or decrease the impact of older estimates on the forecast, depending on whether the capture of slow or fast moving trends is preferable.

Option 4:

It may be possible to empirically approximate disturbance profiles based on historical trends. As the primary sources of disturbance are assumed here to be solar, occupancy and equipment gains which, to a certain extent, follow a daily pattern, a sensible approach may be to apply the estimated disturbance profile of the previous day.

While the assumption that the same disturbances will occur on successive days is unlikely to be correct, it may allow for a practical approximation to be made. At each sample, as the current disturbance is estimated, the forecast can be updated in the same manner as (6.8). Where N_{day} is the number of samples in a day, the predicted disturbance is given as:

$$\hat{x}_{d_j}(k+i) = \hat{x}_{d_j}(k+i-N_{day}) - (\hat{x}_{d_j}(k-N_{day}) - \hat{x}_{d_j}(k)) \quad (6.12)$$

To determine a *best-case-scenario* for the formulation in the presence of disturbance, the strategy was implemented with the assumption of perfect knowledge of the future disturbance estimate. In parallel to this, an alternative set of input sequences was calculated at every sample using each of the disturbance prediction options outlined above.

The rms values of the projected deviations in the input sequences from the *best-case-scenario*, calculated at each sample using (6.3), are shown for each option in Fig. 6.7-

Fig. 6.10 for prediction horizons of 5, 10 and 20 steps. Using Option 1, with no disturbance assumed, the calculated input was affected for all horizons, particularly at the beginning of the prediction horizon (which is relevant considering that only the first element of the derived input sequence is implemented).




Figure 6.7: Difference between inputs calculated with perfect prediction of disturbance estimate and prediction using Option 1 (future disturbance = 0)

Using Option 2 (Fig. 6.8), in which the current disturbance estimate was held constant across the full horizon, a significant reduction in deviation from the *best-case-scenario* can be seen. For longer horizons, it can be seen that the assumption of a constant disturbance becomes less valid further into the future.




Figure 6.8: Difference between inputs calculated with perfect prediction of disturbance estimate and prediction using Option 2 (future disturbance = current estimated disturbance)

In Fig. 6.9, the input deviation obtained using an EWMA prediction of disturbance

values is shown. The EWMA prediction was obtained using the previous 20 disturbance estimates with $\beta_d = 0.01$ to allow for historical values to have a large influence on the predicted future values. A larger input deviation than the constant disturbance case was observed however.




Figure 6.9: Difference between inputs calculated with perfect prediction of disturbance estimate and prediction using Option 3 (future disturbance = EWMA of previous 20 samples)

Using Option 4, the disturbance profile of the previous day was used for prediction. It can be seen in Fig. 6.10 that the resulting performance does not offer an improvement relative to the more simple approach of Option 2.




Figure 6.10: Difference between inputs calculated with perfect prediction of disturbance estimate and prediction using Option 4 (future disturbance = profile from previous day)

In Table 6.1 the rms deviation found for the first step of the horizon is shown for all

cases.

Table 6.1: Difference between first element of derived input sequence calculated with perfect prediction of disturbance estimate and those calculated with different disturbance prediction methods

RMS INPUT DEVIATION USING DISTURBANCE PREDICTION			
	$H = 5$	$H = 10$	$H = 20$
Option 1 (e_{rms})	1.241	1.415	1.489
Option 1 (e_{rms})	0.101	0.133	0.180
Option 3 (e_{rms})	0.357	0.356	0.405
Option 4 (e_{rms})	0.111	0.159	0.238

The results indicate that for the horizons considered here, the disturbances can be assumed to be constant.

6.3 Faults and System Changes - Detection and Identification

In large systems, with many actuators, the occurrence of faults and system changes may not be immediately apparent, yet, if appropriate measures are not taken, significant performance degradation can result (Zhang & Jiang 2008). The handling of such scenarios is often (appropriately) referred to as abnormal event management (AEM) (Venkatasubramanian, Rengaswamy, Yin & Kavuri 2003). Whereas in Section 6.2, the impact of uncertainties in standard operating conditions was analysed, the operation of the system in non-standard conditions is now considered.

The fault-tolerant properties of MPC have been widely studied (Camacho et al. 2010, Boskovic & Mehra 2002, Maciejowski 1999a). In (Maciejowski 1998) for example, the implicit 'daisy-chaining' property of MPC is outlined in the presence of actuator failures, whereby in certain circumstances, knowledge of the fault may not be required provided redundant actuators are available and the control increment is included in the cost function. More generally, if knowledge of the fault is available, the optimisation constraints can be updated to incorporate the post-fault dynamics of the system to allow for some performance criteria to be maintained. A comprehensive review on

fault-tolerant control (FTC) can be found in (Zhang & Jiang 2008). The application of these methods to the building heating control strategy is considered here.

6.3.1 Fault Detection and Identification using Measured Data

When a fault occurs, a fault-tolerant strategy will typically include stages for fault detection (determining if there is a fault), fault isolation (determining the location of the fault), fault identification (determining the magnitude of the fault) and reconfiguration (updating the control formulation to allow for an acceptable performance to be obtained) (Hwang et al. 2010). Without effective methods of fault detection and identification (FDI), the reconfiguration possibilities within the strategy will be severely limited.

Some possible approaches, suitable for the application of building heating systems, are explored here using the simulation platform. This analysis is not exhaustive - FDI in itself is an expansive research topic with many possible alternatives (a detailed three-part review of different approaches taken is provided in (Venkatasubramanian, Rengaswamy, Kavuri & Yin 2003, Venkatasubramanian, Rengaswamy & Kavuri 2003, Venkatasubramanian, Rengaswamy, Yin & Kavuri 2003)).

Faults and changes that can affect a building are broadly categorised here in one of three groups: actuator faults, system faults and sensor faults. A fault in which a mixing valve experiences a reduced range of possible positions or a reduced rate of change of position for example, would be classed as an actuator fault. Similarly, control of the zone-level thermostatic valves may be lost. User interaction with the building by opening doors or windows could be classed as a system fault, as the underlying model representing the dynamic interactions between different parts of the building and the external environment changes. Sensor faults could come in the form of offsets, additional noise or complete loss of communication with a sensor. FDI techniques can be split into model-based techniques and data-based techniques (Isermann 2005). A data-based methodology is first implemented to detect the presence of sensor faults in the system, while a model-based approach is then used to detect the presence of actuator faults.

6.3.2 Detection and Identification of Sensor Faults

Given the large amount of sensor data present, in some circumstances, automatic detection of a fault occurrence may not be trivial. A single zone temperature for example,

when viewed in isolation, may not appear to be unusual until compared in some way to the temperatures of the surrounding zones. In (Wang & Xiao 2004), a technique is outlined whereby sensor faults are detected by observing the relationship between measurements, rather than observing the measurements themselves. This methodology is implemented here using the simulation platform by re-visiting the concepts of principal component analysis (PCA) that were outlined in Chapter 3 for the purposes of disturbance estimation.

For a training period in which no fault is present, a batch of measured data from each sensor is collected in the matrix Z_{train} , where the columns of Z_{train} correspond to m different sensors and the rows correspond to n measurements. An m -dimensional variable then exists for each sample. This data is normalized so that each dimension is zero-mean and unit variance. For the i^{th} row and j^{th} column:

$$\tilde{Z}_{train,i,j} = \frac{(Z_{train,i,j} - \bar{Z}_{train,j})}{\sigma_{Z_j}} \quad (6.13)$$

where $\bar{Z}_{train,j}$ is the sample mean and σ_{Z_j} is the standard deviation of the data measured by the j^{th} sensor, given by:

$$\sigma_{Z_j} = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (Z_{train,i,j} - \bar{Z}_{train,j})^2} \quad (6.14)$$

As in (3.32), the covariance matrix Σ_{pca} of the normalized data is found. The eigenvalues (Γ_{pca}) and eigenvectors (Λ_{pca}) associated with this matrix are then determined. An appropriate number of principal components are selected by observing the relative magnitudes of the elements of (Λ_{pca}). Using this number of eigenvalues, the normalized data-set is projected to a lower-dimensional subspace. Where Γ_p denotes the set of eigenvectors corresponding to the p largest eigenvalues, the set of estimated data points is then given by:

$$\tilde{Z}'_{train} = \Gamma_p \Gamma_p^T \tilde{Z}_{train} \quad (6.15)$$

The sum of the squared differences between this PCA-filtered estimate and the normalized data-set (often referred to as the Q-statistic) is the found as:

$$Q_{st_{train,i}} = \sum_{j=1}^m (\tilde{Z}_{i,j} - \tilde{Z}'_{i,j})^2 \quad (6.16)$$

$$\forall i = 1, 2, \dots, n$$

The Q-statistic determined for the unfaulted training set can act to provide an indication of the typical level of variance between the sensor measurements. As new data is measured (Z_{new}), it is normalized and projected to the subspace determined in the training period by the transformation:

$$\tilde{Z}'_{new} = \Gamma_p \Gamma_p^T \tilde{Z}_{new} \quad (6.17)$$

The Q-statistic for the new data is then determined and compared to a threshold, chosen to reflect the typical values observed during the training period. If the threshold is exceeded, a fault is assumed to have occurred.

As previously noted, detection of the fault is the first stage of the process - the location and nature of the fault is still unknown. A possible approach for locating the sensor in which the fault has occurred is to consider the contribution of each of the sensors to the overall variance indicated by the Q-statistic, or more specifically, the *change* in contribution of each sensor. This is referred to as Q-contribution analysis. The proportion of the Q-statistic accounted for by each sensor can be determined at each sample and plotted in a Q-contribution plot. A schematic outlining the fault detection and isolation strategy is shown in Fig. 6.11.

To demonstrate this methodology, two scenarios in which sensor faults occur were replicated using the simulation platform, with weather compensation used for control of the heating system. Initially, disturbance-free data was used to demonstrate the concept. The impact of disturbances on the strategy was then considered. A four day set of unfaulted training-data was used to develop a suitable principal component subspace, using measurements of all the ground-floor zone temperatures. Using a different four day simulation period, three different scenarios were investigated. In the first scenario, no fault was applied so as to observe suitable threshold ranges to avoid false alarms. In the second scenario, a bias of $+2^\circ C$ was applied to the zone temperature sensor of an office on the east side of the building after the first day of the simulation. For the final scenario, a bias of $-2^\circ C$ was applied after the second day to the zone temperature sensor of a different office, this time on the west side of the building.

In Fig. 6.12, the Q-statistics determined for each scenario, as well as the training period, are shown. The occurrence of each fault is clear, while it is also evident that the Q-statistic determined with no fault present was far higher than that determined for the training set. This implies that to avoid false alarms, a different set of data should be used to determine a fault threshold than was used to develop the PCA-based subspace

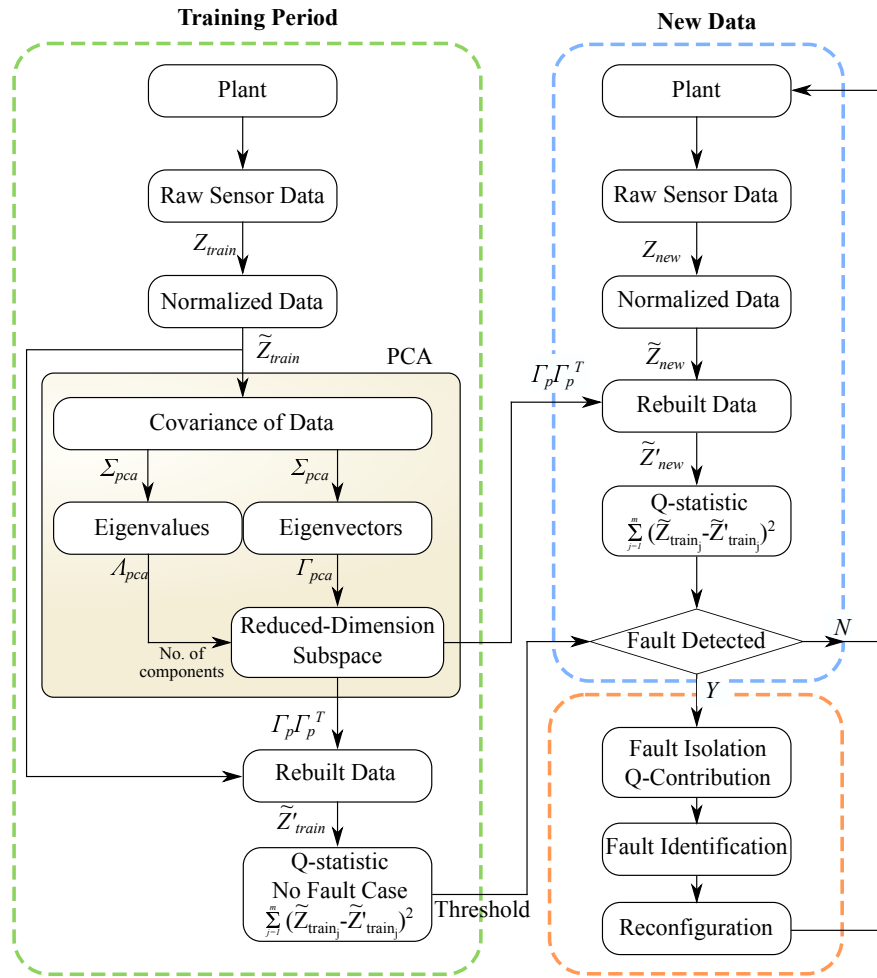


Figure 6.11: Diagram of sensor fault detection and isolation process using PCA-based techniques with measured data

transform.

To locate the faults at the time of detection, the Q-contribution plots are shown for each scenario in Fig. 6.13-Fig. 6.15, where different colours correspond to different zone sensor contributions. For a period of standard operation, the relative contribution of the variance associated each specific sensor to the overall variance of the Q-statistic undergoes a certain amount of natural variation. This can be seen in Fig. 6.13 which shows the unfaulted scenario.

If a large increase in the overall Q-statistic is observed due to an increase in variance of one of the sensors, the contribution of that sensor to the overall variance should greatly increase. At the time when a fault occurrence was detected using the Q-statistic plot of Fig. 6.12, the Q-contributions of the faulted sensors in scenario 2 and scenario 3 underwent a significant increase, dominating the variance of the data-set. This can be seen in Fig. 6.14 which shows the Q-contribution for scenario 2 in which a fault




Figure 6.12: Q-statistic showing detection of sensor faults in two zones




Figure 6.13: Q-contribution plot showing the proportion of total data variance contributed by each sensor in the unfaulted scenario

occurs in Sensor 9 (the east office) after one day and in Fig. 6.15, which shows the Q-contribution for scenario 3 in which a fault occurs in Sensor 3 (the west office) after two days.

The simulation was repeated with the addition of disturbances in the form of solar and internal gains. In the presence of a changing unmeasured disturbance, the detection of faults becomes less straightforward. The Q-statistic for each scenario is shown in Fig. 6.16 along with the same threshold used in the undisturbed case. It can be seen that the sensor faults are detected in scenario 2 and scenario 3 as before. In this case however, the fault threshold is exceeded on two occasions in the unfaulted scenario. While a higher threshold might then be advisable to allow for this, it can also be seen that the faults have the same impact on the variance as in the unfaulted case. The difference between the Q-statistic observed when a fault occurs and that which occurs during normal operation is far less than in the disturbance-free example, making faults more difficult to distinguish.




Figure 6.14: Q-contribution plot showing the proportion of total data variance contributed by each sensor with fault in Sensor 9




Figure 6.15: Q-contribution plot showing the proportion of total data variance contributed by each sensor with fault in Sensor 3




Figure 6.16: Q-statistic showing detection of sensor faults in two zones when disturbed data is used

In Fig. 6.17 the Q-contribution plot for the unfaulted data-set is shown. At the time when the first false alarm occurs, the contribution of sensor 8 appears to grow, while at the time of the second false alarm, the largest increases can be seen in sensor 2 and sensor 10.




Figure 6.17: Q-contribution plot determined with disturbances present in the data - unfaulted scenario

These increases in contribution are far less in each individual sensor than those observed when an actual fault occurs however. This can be seen in the Q-contribution plots of scenarios 2 and 3 in Fig. 6.18 Fig. 6.19, in which the distinction between fault and false alarm is clear. The disturbances applied to the simulation platform tend to impact several zones at once. As such, though the overall Q-statistic may increase beyond the threshold due to a particularly large disturbance, the increase in variance tends to be spread across several zones. Observing both the Q-statistic and the Q-contribution may allow for a more confident decision on fault presence to be made.




Figure 6.18: Q-contribution plot determined with disturbances present in the data - fault in Sensor 9

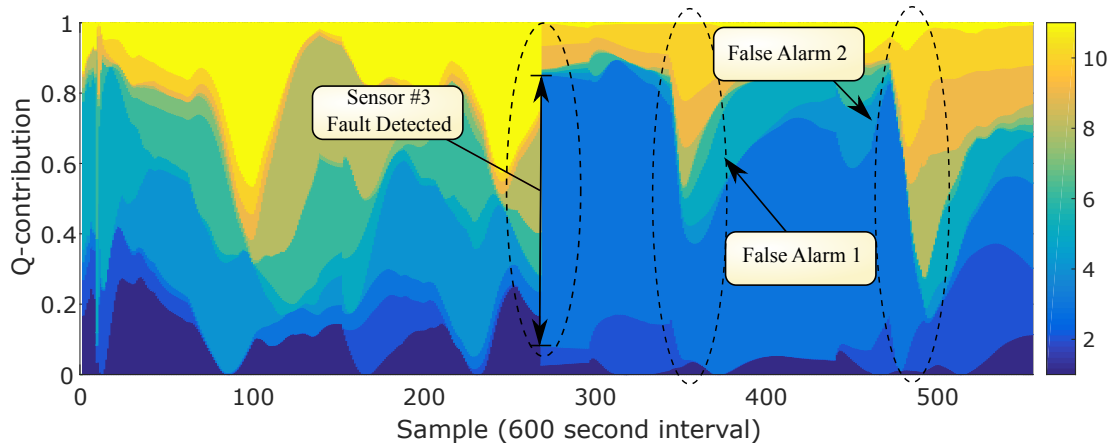


Figure 6.19: Q-contribution plot determined with disturbances present in the data - fault in Sensor 3

6.3.3 Detection and Identification of Actuator Faults

For detection of sensor faults, model-free approaches can be suitable. A sensor offset, such as that simulated in the previous section for example, may not result in a change in the underlying dynamics of the model. For actuator and system faults however, the model dynamics will undergo a change. As a fault to one of the components may affect the measurements obtained in many sensors, the fault may be less obvious if only the sensor data is considered. Observing a change in the relationship between the inputs that are applied and the outputs that are measured may offer greater clarity. As such, a model-based fault identification methodology is developed in this section to detect and isolate faults in actuator components. In this case, training data-sets are initially obtained for each of fault scenarios considered. These training sets are grouped in clusters, based in some way on the similarity of the each data point to the value predicted by an underlying model. As each new measurement is obtained, classification techniques are used to assign the new variable to the most suitable training data cluster.

In the methodology outlined here, the decentralised Kalman filtering techniques and zone models outlined in (4.30)-(4.31) are used to determine a disturbance state associated with each measurement given for the j^{th} sensor as \hat{x}_{d_j} . The use of these disturbance states is predicated on the idea that in a faulted scenario, a larger-than-normal deviation of the system from the modelled dynamics should occur. This should be captured in the disturbance state of the augmented Kalman Filter - a fault is then considered to resemble an unmeasured disturbance.

For the q^{th} scenario (faulted or unfaulted), the full batch of n training period samples from m sensors is given by \hat{X}_{dq} . The data from the j^{th} sensor from the q^{th} data-set is

normalized as:

$$\tilde{X}_{dqj} = \frac{|\hat{X}_{dqj}|}{\max\{|\bar{X}_{d1j}|, \dots, |\bar{X}_{dmj}|\}} \quad (6.18)$$

where $|\bar{X}_{dqj}|$ is the mean absolute value associated with the j^{th} column of $|\hat{X}_{dq}|$. Each of the normalized training data-sets is then considered to be a separate cluster. As new data is measured, it is normalized in the same way and classified as belonging to one of the training clusters.

In (House et al. 1999), different classification techniques are outlined for identifying faults in an air-handling unit. Techniques considered include k-Nearest Neighbour (k-NN) classification, Bayes classification, which seeks to minimise classification errors, artificial neural network (ANN) classifiers, which map input patterns to output patterns and rule-based classifiers which exploit knowledge of specific patterns that are known to be associated with specific scenarios using pre-defined IF-THEN rules. In all cases, fuzzy membership functions can also be used which allow for one data-point to be classified in several clusters with different levels of membership strength. In this section k-NN classification is used.

Using some distance metric (typically the Euclidean distance), for each new data-point, a k-NN algorithm finds the k 'closest' training data-points. In its simplest form, it assigns the new data-point to the class that is most common among these k neighbours. A variation of this which reduces the dimension of the problem, referred to as k-nearest prototype classification, only considers the proximity of the new data-point to prototype points associated with each of the training classes (usually, the centroid of the training data). In Fig. 6.20, an example in which the classification of a new two-dimensional data-point into one of two classes is shown, first using the 9 nearest neighbours, then using the nearest prototype. In both cases, the new data-point would be assigned membership to Class 1.

The full fault isolation process where two fault classes are determined is shown in Fig. 6.21.

Using the simulation platform, three seven day training periods were simulated. In the first, no fault was present. The second training data-set simulated the operation of the plant whereby the ground floor mixing valve gets stuck at a position in which 10% of supply water is mixed with 90% return water. The final simulation period represented a scenario in which the zone-level valves of the ground floor stairwells were stuck in the closed position.

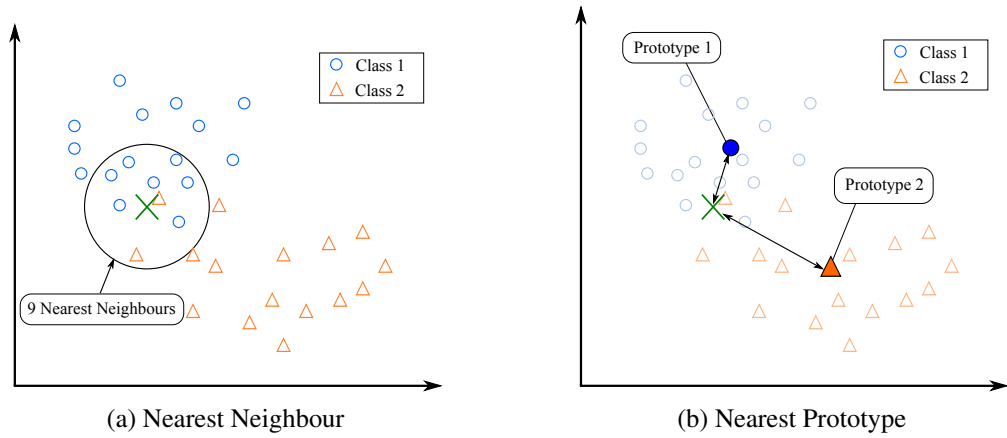


Figure 6.20: k-Nearest Neighbour and k-Nearest Prototype classification methods for classifying new data-point




Figure 6.21: Diagram of actuator fault detection and identification process using Kalman filtering and machine learning

In each scenario, the augmented Kalman filter model structure of (4.30)-(4.31) was used to determine sequences of estimated disturbances associated with each zone, which were then normalized using (6.18). For each sample, the ℓ largest normalized disturbance estimates were taken to represent an ℓ -dimensional data-point for the sample. The data-points from each of the three scenarios are shown in Fig. 6.22, with $\ell = 3$.




Figure 6.22: Training clusters developed for classification of actuator faults - $\ell = 3$

The separation between each of the scenarios becomes more clear by observing in Fig. 6.23-Fig. 6.25 the relationships between each of the dimensions.




Figure 6.23: Training clusters - largest deviation from model estimate vs. second largest

Three new seven-day sets of validation data associated with each scenario were generated and assigned to one of the three classes using the k-NN methodology, where




Figure 6.24: Training clusters - largest deviation from model estimate vs. third largest




Figure 6.25: Training clusters - second largest deviation from model estimate vs. third largest

the Euclidean distance was used to determine the K nearest neighbours. Results were determined for $\ell = 3$ and $\ell = 5$, using $k = 5$ and $k = 25$ neighbours.

To determine the classification accuracy, the probability of the data being correctly assigned was calculated for each data-set, given for the q^{th} set of data-points as:

$$P(q|Sc_q) = \frac{class_q}{n_q} \quad (6.19)$$

where n_q is the number of data-points in the q^{th} data-set while $class_q$ is the number of data-points in the q^{th} data-set that are correctly classified. To allow for false alarms to be taken into account, Bayes theorem was then used to determine the probability of a particular assignment being correct. The probability, $P(Sc_q|q)$, of an assignment to scenario q being correct is given by:

$$P(Sc_q|q) = \frac{P(q|Sc_q)}{Tot_q} \frac{n_q}{\sum_{ii=1}^3 n_{ii}} \quad (6.20)$$

where Tot_q is the proportion of validation data from all data-sets that was classified as being from scenario q . More concisely, $P(q|Sc_q)$ gives the probability of a data-point, known to be from the scenario q , being correctly classified, whereas $P(Sc_q|q)$ gives the probability of a data-point of unknown scenario, which has been classified as q , being from scenario q . As the true scenario of the data-point will not in reality be known, the latter probability is the more relevant measure in terms of fault detection.

In Table 6.2, the accuracy with which the data is classified is shown for each of the scenarios. The most accurate results were obtained using the 5 nearest neighbours with the data dimension $\ell = 5$.

Table 6.2: Accuracy of fault identification for unfaulted scenario, zone-valve fault scenario and mixing-valve fault scenario

CLASSIFICATION ACCURACY				
	$\ell = 3$		$\ell = 5$	
	$k = 5$	$k = 25$	$k = 5$	$k = 25$
$P(1 Sc_1)$	0.9963	0.9926	0.9926	0.9963
$P(2 Sc_2)$	0.7593	0.4963	0.8444	0.8704
$P(3 Sc_3)$	0.4222	0.3704	0.8630	0.8370
$P(Sc_1 1)$	0.8078	0.8072	0.8072	0.8078
$P(Sc_2 2)$	0.6212	0.4891	0.9744	0.9476
$P(Sc_3 3)$	0.7755	0.4902	0.9549	0.9869

Though this technique can allow for a distinction to be made between many faults with a high-accuracy, training data is required for each of the considered faults. Given the number of actuators present in a large building, as well as the range of different faults by which each can be compromised, it may only be possible to characterise a low number the most significant possible faults.

6.4 Reconfiguration of the MPC Formulation when a Fault Occurs

If knowledge of a fault or a general system change is available, it may be necessary to reformulate the MPC problem in order to obtain some desired level of performance. As

previously mentioned, a certain level of fault tolerance is available if the cost function used penalizes a combination of the set-point deviation and input increment. Where at time k , if $\hat{z}(k)$ is the model output, $\hat{r}(k)$ is the set-point reference and $\hat{u}(k)$ is the input, such a form could be given for a prediction horizon H by:

$$J(k) = \sum_{i=1}^H \|\hat{z}(k+i|k) - r(k+i)\|_Q^2 + \sum_{i=0}^{H-1} \|\Delta\hat{u}(k+i|k)\|_{\mathcal{R}}^2 \quad (6.21)$$

where Q and \mathcal{R} are weighting matrices.

Including the input increment in the objective introduces a kind of integral action to the strategy, removing (if possible) the steady-state offset between the output and the desired reference (Maciejowski 2002). This could be seen as *passive* FTC. For the formulations outlined here however, the input increment is not included. If a form is used such as:

$$J(k) = \sum_{i=1}^H \|\hat{z}(k+i|k) - r(k+i)\|_Q^2 + \sum_{i=0}^{H-1} \|\hat{u}(k+i|k)\|_{\mathcal{R}}^2 \quad (6.22)$$

an output tracking error is inevitable in unfaulted conditions, the magnitude of which depends on how the objective is weighted. If the system changes, this offset will change. *Active* FTC is then required.

Depending on the magnitude and nature of the fault, a decision as to whether operation can continue must be made. For certain applications (particularly safety-critical applications), shut-down of the plant may not be a viable option and some kind of safe-mode of operation is required (Maciejowski 1999b). For building energy applications however, a more relevant consideration is whether the plant can operate with a performance level similar to the pre-fault conditions, or whether some reduced level of performance in certain zones should be specified. In both cases, updating the underlying optimisation model (and hence, the associated optimisation problem constraints) and updating the objective(s) may be required.

6.4.1 Updating the Model and the Optimisation Constraints

Different types of faults will tend to necessitate different reconfiguration approaches. With reliable FDI assumed, in the case of an equipment (e.g. boiler) or actuator fault, constraints associated with the affected elements can be directly updated to incorporate new limitations imposed by the fault. Restricted movement of the mixing valve for example, can be included as a constraint. Sensor biases can also be easily compensated

for, while the occurrence of excess sensor noise can be handled by re-tuning the local Kalman filters (which are already present in the strategy to provide state estimates).

For system faults (such as open windows), automatically updating the underlying optimisation models to capture the new faulted plant dynamics may be required. In many applications, particularly those in which high-fidelity modelling techniques are used, a bank of parallel pre-defined models can be developed in advance, each incorporating the properties associated with a specific fault or set of faults. When a fault is identified, the optimisation constraints are updated accordingly using the appropriate new system models (Hwang et al. 2010). This multiple-model approach is often advocated for re-configurable flight control strategies where faults must be immediately accounted for to ensure safe operation (Zhang & Jiang 1999, Boskovic & Mehra 2002, 1998).

An effective multiple-model based strategy must consider all possible failure modes including faults on each actuator as well as combinations of these faults. Given the large number of actuators and sensors present in a typical office building, for the application of building energy systems, this may not be reasonable. Another alternative, which is possibly more suited to an application in which data-driven models are used, is an adaptive control approach whereby model parameters are updated based on newly measured data (Maasoumy et al. 2014). If a fault occurs, the change in system dynamics may be captured in the new data and new models can be derived using system-identification techniques (though the data-quality may be compromised). Using such a strategy will result in a delayed reaction to the fault as the parameters are re-tuned. For the non-safety-critical application of building energy, this may be acceptable.

An additional alternative, given the approximately decentralised nature of the building energy system, is to remove faulted elements from the optimisation. If a fault occurs in a single zone valve for example, removing the affected zone from the problem may be the best solution. Diminished performance will be experienced in the faulted zone, but the remaining zones can still be considered unfaulted.

In Table 6.3, some of the possible faults which could affect the heating system control strategy are listed, as well as the possible actions which may need to be taken if such a fault occurred.

6.4.2 Reformulating the Objective Function in Faulted Conditions

Though the optimisation constraints can be updated to incorporate knowledge of a fault, if the underlying models are changed, the problem objective may also need to

Table 6.3: Possible actions to be taken when the system is affected by a particular fault

RECONFIGURATION APPROACH				
	Update Constraints	Adaptive Identification	Re-tune Kalman filter	Remove Faulted Zone
Actuator Fault (Zone Valves)	-	-	-	✓
Actuator Fault (Mixing Valves)	✓	-	-	-
Actuator Fault (Equipment)	✓	-	-	-
Sensor Fault (Noise)	-	-	✓	-
Sensor Fault (Loss)	-	-	-	✓
System Fault	✓	✓	-	-

be reconfigured. The issues associated with this reconfiguration problem are examined here (and in (O'Dwyer et al. 2016)) by analysing an example in which a zone-level valve fault enforces the removal of the zone from the formulation. The standard approach quadratic formulation as well as the prioritised objective approach of Chapter 5 are compared in terms of the ease with which a desired post-fault performance can be achieved.

As previously referred to, when using a standard single cost function formulation such as (6.22), a large number of weighting parameters are required to balance the comfort objectives in each zone with the energy minimisation objective. Typically, these weights will be designed for a disturbance-free, fault-free scenario. In the presence of disturbances and faults however, the balance between comfort and energy can become somewhat arbitrary (Kerrigan & Maciejowski 2002). Even if the changes to the system are accounted for in the constraints, the cost function must be recomputed in order to obtain some desired performance (O'Dwyer et al. 2015).

The need for retuning is illustrated in the following example where the temperatures of two zones are controlled by a single input using an MPC strategy with the cost function given by 6.22.

The sequence of zone temperatures (T_{z1} and T_{z2}) over the prediction horizon H at time k is expressed in terms of the sequence of inputs $U(k)$ and current zone states

$\mathbf{x}_1(k)$ and $\mathbf{x}_2(k)$ as:

$$\begin{bmatrix} \mathbf{T}_{z1} \\ \mathbf{T}_{z2} \end{bmatrix} = \begin{bmatrix} \Psi_1 & 0 \\ 0 & \Psi_2 \end{bmatrix} \begin{bmatrix} \mathbf{x}_1(k) \\ \mathbf{x}_2(k) \end{bmatrix} + \begin{bmatrix} \Phi_1 \\ \Phi_2 \end{bmatrix} \mathbf{U}(k) \quad (6.23)$$

The optimal unconstrained solution of (6.23) for the input sequence is then (Maciejowski 2002):

$$\mathbf{U}^*(k) = \left(\begin{bmatrix} \Phi_1^T & \Phi_2^T \end{bmatrix} \mathcal{Q} \begin{bmatrix} \Phi_1 \\ \Phi_2 \end{bmatrix} + \mathcal{R} \right)^{-1} \begin{bmatrix} \Phi_1^T & \Phi_2^T \end{bmatrix} \mathcal{Q} \begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix}_k \quad (6.24)$$

where \mathcal{Q} and \mathcal{R} are the tuning matrices and ξ_1 and ξ_2 are the predicted sequence of free response set-point deviations over the prediction horizon at time k . This relationship represents the balance between energy consumption (the accumulated input) and comfort satisfaction (the set-point deviation) in each of the zones.

A scenario is now considered whereby control of the second zone is lost (a stuck valve, for example, may block the input), as shown in Fig. 6.26. The best approach may be




Figure 6.26: Controlling 2 zones - Reconfiguration in Fault Scenario

to remove the faulted zone from the optimisation problem without affecting the energy/comfort balance of the unfaulted zone. Given knowledge of the fault, the formulation could be reconfigured to simulate the presence of the fault (and remove the faulted zone set-point deviation from the cost function) by setting $\Phi_2 = 0$. The new optimal solution of the unconstrained problem is:

$$\mathbf{U}_{fault}^*(k) = \left(\Phi_1^T \tilde{\mathcal{Q}} \Phi_1 + \mathcal{R} \right)^{-1} \Phi_1^T \tilde{\mathcal{Q}} \xi_1(k), \quad (6.25)$$

where $\mathcal{Q} \in \mathbb{R}^{H \times H}$ is given by $\tilde{\mathcal{Q}}_{i,j} = \mathcal{Q}_{i,j}$, for $i = 1, 2, \dots, H$ and $j = 1, 2, \dots, H$.

Though the set-point deviation term of the faulted zone has been eliminated from the cost, the balance between the input and the unfaulted zones set-point deviation has been altered. To achieve the same balance as the pre-faulted scenario, the tuning matrices must be recalculated. In the case of a building with a large number of zones

and actuators, pre-defined weightings for all possible combinations of fault scenarios quickly becomes intractable.

Using a lexicographic approach however, as no tuning is required, the post-fault behaviour of the system is pre-defined (Miksch & Gambier 2011) and a fault of this class in one zone does not propagate across the building. The primary and secondary objectives remain unchanged before and after the fault, provided the fault has been accounted for in the model (in this example, by setting $\Phi_2 = 0$).

To illustrate the highlighted reconfigurability issues using the simulation platform, a scenario was simulated whereby the radiator valve of the ground floor corridor was stuck in the closed position and control to the zone was lost. Knowledge of the fault was incorporated into the prediction models by setting the gain of the input to the faulted zone to zero.

The problem was first simulated using the standard quadratic cost function formulation. The weighting matrices Q and R were chosen to provide a (subjectively) sensible balance comfort satisfaction and energy (in this case, the energy and comfort performance of the unfaulted system is not significant, only the change in performance observed when the fault occurs).

The operation of the control strategy was simulated with and without the presence of the fault for the same period. In Fig. 6.27, the simulated output of the ground floor meeting room can be seen for both scenarios. This room should be unaffected by the fault, as the fault was in another zone and the change of dynamics was accounted for in the constraints, but as can be seen (and as shown in (6.24)-(6.25)), updating the constraints of the problem without recomputing the weighting matrices can impact the performance in the unfaulted zones.

In Fig. 6.28, the output of the same zone is shown for the faulted and unfaulted scenarios, this time using the prioritised-objective approach. If the fault is accounted for in the constraints, no adjustment of the cost function is required. Though the solution may be different (as the system is different), the comfort and energy objectives of the unfaulted zones remains the same.

6.5 Conclusion

To observe the effect of inaccurate zone models in an MPC formulation, the quadratic MPC strategy of Chapter 5 was repeated, calculating two different sets of optimal in-




Figure 6.27: Simulated temperature measurement of ground floor meeting room with and without the presence of a zone-level fault in corridor - Standard quadratic MPC approach




Figure 6.28: Simulated temperature measurement of ground floor meeting room with and without the presence of a zone-level fault in corridor - Prioritised MPC approach

puts. The first set of inputs was calculated using high-accuracy zone models derived from undisturbed data taken from the simulation platform and applied to the plant. The second set of inputs, calculated in parallel but not applied, was determined using lower accuracy zone models derived from data in which disturbances were present without disturbance estimation included. It was shown that with a longer horizon, the difference between the high-accuracy and low-accuracy calculated input became significantly more pronounced. An additional example in which an external temperature profile forecast was used was also simulated. Again, a longer horizon amplifies the impact of an incorrect prediction, though the overall effect was shown to be small.

Unmeasured disturbances in the form of solar, equipment and occupancy gains were applied to the system, with different methods for including disturbance prediction in the formulation attempted. The best results (i.e. closest to a hypothetical *best-case-scenario* in which the future disturbances were known) were obtained when the current

disturbance in each zone, estimated by Kalman filtering, was assumed to be constant for the full prediction horizon.

The impact of faults on the MPC strategy was next considered. A model-free strategy for locating sensor faults using principal component analysis was implemented in the simulation platform. The performance of the detection and identification strategy was shown to reduce in the presence of disturbances by way of an increased number of false alarms. A model-based fault identification strategy using Kalman filtered disturbance estimates and k-Nearest Neighbour classification was then introduced to detect and identify occurrences of mixing valve and zone-level valve faults.

Approaches for reconfiguring the MPC strategy in terms of both the objectives and constraints were outlined. The robust qualities of the prioritised-objective MPC formulation (developed in Chapter 5) were highlighted by simulating a scenario in which control to one zone is lost. It was shown that, when using a standard single-objective quadratic formulation, a fault in one zone can propagate to unfaulted zones even if it has been accounted for in the prediction models, unless the weighting matrices are recomputed. The prioritised formulation outlined does not have this recomputational requirement.

References

- Boskovic, J. & Mehra, R. (1998), A multiple model-based reconfigurable flight control system design, *in* ‘Proceedings of the 37th IEEE Conference on Decision and Control (Cat. No.98CH36171)’, Vol. 4, pp. 4503–4508.
- Boskovic, J. & Mehra, R. (2002), ‘Fault accommodation using model predictive methods’, *Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301)* **6**, 4–9.
- Camacho, E., Alamo, T. & de la Pena, D. (2010), ‘Fault-tolerant model predictive control’, *Emerging Technologies and Factory Automation (ETFA)* .
- Chilin, D., Liu, J., Davis, J. F. & Christofides, P. D. (2012), ‘Data-based monitoring and reconfiguration of a distributed model predictive control system’, *International Journal of Robust and Nonlinear Control* **22**, 68–88.
- Forecast Accuracy - Met Éireann - The Irish Meteorological Service Online* (2016).
- House, J., Lee, W. & Shin, D. (1999), ‘Classification Techniques for Fault Detection and Diagnosis of an Air-Handling Unit’, *ASHRAE Transactions* .
- Hwang, I., Kim, S., Kim, Y. & Seah, C. (2010), ‘A survey of fault detection, isolation, and reconfiguration methods’, *IEEE Transactions on Control Systems Technology* **18**(3), 636–653.
- Isermann, R. (2005), ‘Model-based fault-detection and diagnosis - Status and applications’, *Annual Reviews in Control* **29**(1), 71–85.
- Kerrigan, E. & Maciejowski, J. (2002), ‘Designing model predictive controllers with prioritised constraints and objectives’, *IEEE International Symposium on Computer Aided Control System Design* pp. 33–38.
- Lazos, D., Sproul, A. B. & Kay, M. (2014), ‘Optimisation of energy management in commercial buildings with weather forecasting inputs: A review’, *Renewable and Sustainable Energy Reviews* **39**, 587–603.

- Maasoumy, M., Razmara, M., Shahbakhti, M. & Vincentelli, a. S. (2014), 'Handling model uncertainty in model predictive control for energy efficient buildings', *Energy and Buildings* **77**, 377–392.
- Maasoumy, M. & Sangiovanni-Vincentelli, A. (2012), 'Optimal control of building HVAC systems in the presence of imperfect predictions', *Proceedings of the Dynamic System Control Conference* pp. 1–10.
- Maciejowski, J. (1998), 'The implicit daisy-chaining property of constrained predictive control', *Applied Mathematics and Computer Science* **8**(4), 101–117.
- Maciejowski, J. (1999a), 'Fault-tolerant aspects of MPC', *Control: Techniques and Applications-Day 2* (1), 1–4.
- Maciejowski, J. (1999b), 'Modelling and predictive control: Enabling technologies for reconfiguration', *Annual Reviews in Control* **23**, 13–23.
- Maciejowski, J. (2002), *Predictive Control: With Constraints*, Prentice-Hall, UK.
- Miksch, T. & Gambier, A. (2011), 'Fault-tolerant control by using lexicographic multi-objective optimization', *8th Asian Control Conference (ASCC)* pp. 1078–1083.
- Negenborn, R. R., Houwing, M., De Schutter, B. & Hellendoorn, H. (2008), 'Adaptive prediction model accuracy in the control of residential energy resources', *2008 IEEE International Conference on Control Applications* pp. 311–316.
- O'Dwyer, E., Cychowski, M., Kouramas, K., Tommasi, L. D. & Lightbody, G. (2015), Scalable , Reconfigurable Model Predictive Control for Building Heating Systems, in 'European Control Conference 2015 (ECC)', pp. 2253–2258.
- O'Dwyer, E., Tommasi, L. D., Kouramas, K., Cychowski, M. & Lightbody, G. (2016), 'Prioritised Objectives for Model Predictive Control of Buidling Heating Systems', *IEEE Transactions on Control Systems Technology* p. (Under Review).
- Oldewurtel, F., Jones, C. N., Parisio, A. & Morari, M. (2014), 'Stochastic Model Predictive Control for Building Climate Control', *IEEE Transactions on Control Systems Technology* **22**(3), 1198–1205.
- Oldewurtel, F., Parisio, A., Jones, C. N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B. & Morari, M. (2012), 'Use of model predictive control and weather forecasts for energy efficient building climate control', *Energy and Buildings* **45**, 15–27.
- Pawlowski, A., Guzman, J. L., Rodríguez, F., Berenguel, M. & Sanchez, J. (2010),

- Application of time-series methods to disturbance estimation in predictive control problems, in 'IEEE International Symposium on Industrial Electronics', pp. 409–414.
- Sun, J., Fang, W., Wu, X., Palade, V. & Xu, W. (2012), 'Quantum-Behaved Particle Swarm Optimization: Analysis of Individual Particle Behavior and Parameter Selection', *Evolutionary Computation* **20**(3), 349–393.
- Sun, J., Feng, B. & Xu, W. (2004), 'Particle swarm optimization with particles having quantum behavior', *Proceedings of the 2004 Congress on Evolutionary Computation (IEEE Cat. No.04TH8753)* **1**, 325–331.
- Venkatasubramanian, V., Rengaswamy, R. & Kavuri, S. N. (2003), 'A review of process fault detection and diagnosis: Part II: Qualitative models and search strategies', *Computers & Chemical Engineering* **27**(3), 313–326.
- Venkatasubramanian, V., Rengaswamy, R., Kavuri, S. N. & Yin, K. (2003), 'A review of process fault detection and diagnosis part III: Process history based methods'.
- Venkatasubramanian, V., Rengaswamy, R., Yin, K. & Kavuri, S. N. (2003), 'A review of process fault detection and diagnosis: Part I: Quantitative model-based methods', *Computers & Chemical Engineering* **27**(3), 293–311.
- Wang, S. & Xiao, F. (2004), 'AHU sensor fault diagnosis using principal component analysis method', *Energy and Buildings* **36**(2), 147–160.
- Zhang, Y. & Jiang, J. (1999), 'Design of integrated fault detection, diagnosis and reconfigurable control systems', *Proceedings of the 38th Conference on Decision and Control* (December), 3587–3592.
- Zhang, Y. & Jiang, J. (2008), 'Bibliographical review on reconfigurable fault-tolerant control systems', *Annual Reviews in Control* **32**(2), 229–252.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

This thesis introduced scalable model predictive control (MPC) techniques suitable for the application of building heating systems. Techniques were developed for deriving detailed high-order building simulation models from data corrupted by disturbance. These techniques were expanded upon to allow for their use in deriving low-order zone models, designed to be incorporated in an MPC-based control strategy. MPC strategies were then considered, focussing particularly on scalability in terms of both the problem complexity and the tuning requirement. A prioritised-objective approach was derived to achieve this. Finally, the ability of MPC to handle inaccurate predictions and faults was considered, with strategies developed for the detection of sensor and actuator faults. The reconfigurability of the prioritised-objective approach was demonstrated.

An overview is now provided of the main conclusions drawn from each chapter.

Chapter 2

In Chapter 2, a background was given to the overall problem of building energy system control, underlining the opportunities now available due to the accessibility of an increased amount of sensor data. Traditional methods currently in use, as well as methods proposed in literature were outlined. Intelligent approaches based on fuzzy logic control, neural networks and multi-agent systems were covered, which can use learning techniques to implicitly incorporate the complex nature of a building into the control decision. This can allow for an improved performance over the more traditional classical approaches, which assume that the system operates within an unchanging lin-

ear operating region. MPC was introduced as an alternative approach which can use predictions of the future building load requirements to determine an appropriate control input. The primary limitation of MPC approaches for implementation on a wider scale was given as the mismatch that exists between the complexity required for performance and the simplicity required for competitive commissioning effort and efficient on-line operation.

Chapter 3

To analyse and compare the control strategies developed in this work, a simulation platform was required which could accurately represent the complex thermodynamics of a building. In Chapter 3 methods for developing such a model were introduced. Using input and output data taken from an RC-network representation of a notional 30 zone building, a model was trained to replicate the network using metaheuristic search algorithms. A chemotaxis algorithm was initially used, though it was found that such an algorithm was not suitable due to a tendency for convergence to local minima. Particle swarm optimisation (PSO) and quantum-behaved particle swarm optimisation (Q-PSO) approaches which are based on swarm intelligence were implemented in place of the chemotaxis algorithm to prevent this. The Q-PSO approach provided the best performance of the algorithms tested.

Incorporating the impact of typical unmeasured disturbances on the data was shown to be necessary to avoid significant bias, particularly given the over-parameterised nature of the problem. A spatio-temporal filtering process was developed to allow for these disturbances to be derived using measured data. Kalman filtering was first used to estimate a disturbance state associated with each zone of the building. Principal component analysis (PCA) was then used to accentuate the commonalities present in disturbance estimates across multiple zones. The spatio-temporally filtered disturbance estimates were then included as an additional input to the system in the identification process. Iteratively, using Q-PSO, Kalman filtering and PCA, the model parameters and disturbance estimates were refined until convergence was achieved. The strategy was shown to work well with the notional 30 zone building model.

The techniques developed were implemented using measured data taken from the Nimbus Centre, a two-storey office building with a hydronic heating system. The model of the building and a representation of the heating system were combined in Simulink to form a simulation platform which can replicate the operation of the building.

Chapter 4

Techniques for developing low-order models for use within a control strategy were de-

veloped in Chapter 4. Differing from the detailed simulation modelling of the previous chapter, these models favour simplicity over complexity, at the expense of a certain amount of accuracy. The objective of the chapter was to develop methods that are straightforward to implement, to derive models which can be used to obtain meaningful predictions of the future state of the plant.

Different frameworks in which the full building can be represented as a set of smaller zone models were first considered. Model reduction techniques were then analysed, in which low-order zone models were extracted from high-order simulation models using Hankel singular value decomposition (HSVD). While such techniques were shown to be effective, the assumption that an accurate high-order model of the building is initially available may not be universally applicable.

Data-driven approaches which use system-identification techniques to directly derive zone models from measured data were advocated as a more scalable alternative. Methods were introduced by which an updated version of the spatio-temporal filtering process (developed in Chapter 3) was combined with prediction-error estimation methods (PEM) to develop ARMAX models that can account for the presence of disturbances in the data. The results obtained indicate that second-order models are adequate, with the provision that the disturbance estimation techniques are used. Once again, the methods were implemented with data measurements from the Nimbus Centre.

Chapter 5

In Chapter 5, scalable MPC formulations were considered. Using the simulation platform to represent the Nimbus Centre (as developed in Chapter 3), standard quadratic MPC formulations were first simulated. The objective was to minimise a combination of energy consumption and internal temperature set-point deviation over a prediction horizon for each zone (where the zone models used were as developed in Chapter 4). It was shown that due to the conflicting nature of energy performance and thermal comfort satisfaction, when such a formulations are used, characterising a desired set of performance criteria, as well as determining the tuning parameters which can produce this desired performance, can result in a problem that is both subjective and abstract in nature. By considering the large number of zones and associated tuning weights, the need to approach the problem differently was shown to become clear.

Different methodologies were introduced to accommodate the scaling issues of the problem. A hierarchical approach was developed in which the control strategy for the full building was determined by only considering two zones, the first of which represents the cold extremes of the building, while the second represents the hotter

extremes. This strategy allowed for the cost function weighting problem complexity to be greatly reduced. A method for determining cost function weights for each zone by considering the explicit solution to an unconstrained version of the optimisation problem was also derived.

A prioritised-objective approach was then formulated which allowed for the tuning weights to be completely removed, while maintaining feasibility of the optimisation using a lexicographic framework. An acceptable *comfortable* temperature band was first defined for all zones. Setting up the problem in an LP structure, the minimum possible deviation from this band was found. The minimum demand-side energy consumption possible to achieve this level of comfort was then found using a QP-type optimisation. To allow for the heating system energy to be minimised, an additional supply-side layer was added to the strategy, incorporating the efficiency curves of the boiler. The formulation was shown to outperform the standard weather-compensated control approach used in the Nimbus Centre in terms of both comfort and energy consumption. The positive results obtained imply that with such an approach, MPC can be viewed as a scalable alternative to traditional strategies, without the need for a significant tuning effort and a knowledge of advanced control.

Chapter 6

The performance of MPC formulations in the face of a typical range of inaccurate predictions and system changes/failures to be expected when controlling a building heating system was investigated in Chapter 6. The optimal input sequences calculated using low-accuracy and high-accuracy zone models were first compared. It was shown that with a lower accuracy, the use of a shorter prediction horizon may become necessary.

Further results demonstrated that when the external temperature forecast was corrected by the current measurement at each sample, the impact of an inaccurate external temperature forecast was minimal for the prediction horizons tested. Methods for predicting additional unmeasured disturbances were then introduced. Comparing the results of the different methods, it was found that the performance that most closely resembled a hypothetical *best-case* scenario (in which perfect disturbance predictions were available) was obtained when the current disturbance estimate was held constant for the duration of the prediction horizon.

The impact of sensor faults and actuator faults in the system was then considered. A model-free approach for detecting and isolating sensor faults using PCA was applied successfully in the simulation platform, however it was shown that with the inclusion of

disturbances in the measurements, a certain degree of clarity was lost. A model-based approach for detecting actuator faults was then implemented using Kalman filtered disturbance estimates and simple machine learning techniques.

A final analysis was carried out indicating that the prioritised-objective strategy developed in Chapter 5 can be reconfigured more easily in the event of a fault to one of the zones than would be the case for a more standard quadratic approach. It was shown that if a quadratic approach is implemented, a change to one of the zone models due to a fault must be accompanied by an update of the cost function to avoid a scenario in which the control of all unfaulted zones is compromised. This was not the case for the prioritised objective formulation, in which the objectives were unaffected by the underlying models.

7.2 Future Work

The potential research opportunities arising from the work carried out in this thesis are outlined in this section, split into four categories which broadly correspond to each of the chapters from Chapter 3-Chapter 6.

Simulation Modelling

With regard to the techniques developed in Chapter 3 for deriving detailed simulation models from data, an area which could be further exploited is in the design of more sophisticated training data profiles that could be applied to the building. This would involve the establishment of a set of inputs which excite the system across as wide a frequency range as possible, while remaining within the feasible operating range of the building at hand.

Following this, a formalisation of the identifiability issues involved would be useful. When implementing the metaheuristic search algorithms, a certain amount of trial and error was required to determine the set of adjustable parameters as well as the appropriate levels of search aggression used. A more strictly defined search-space, based on the excitation of the available data may simplify the design process. Furthermore, as only chemotaxis and PSO-type metaheuristic search algorithms were investigated, the implementation of an evolutionary algorithm approach may provide additional insight.

In the spatio-temporally filtered disturbance estimation process, the derivation of a stricter set of guidelines for the selection of a suitable number of principal components

should be carried out, as well as a more rigorous methodology for grouping the zones, to avoid trial and error based approaches in the design procedure.

Optimisation Modelling

While the benefits of using data-based approaches for the development of zone models was clear, the system identification methods used could be improved once again by considering a suitable range of inputs to sufficiently excite the system. Implementation of functional tests in a real system would then allow for models to be derived which could be compared with those derived using data measurements taken from standard closed-loop scenarios. An analysis of the potential accuracy improvements available by including thermal interaction between zones could be made. If identification of these interactions was possible from data measurements, a distributed control strategy could be considered.

If additional measurements were taken, such as solar gain and occupancy, a comparison between the unmeasured disturbance estimates (developed using the methods of Chapter 4) and the real disturbances could be made. Better methods for prediction of these disturbances could also be incorporated into the control itself.

Control Formulation

Implementation of the prioritised objective control strategy in a real building over an extended period of time could offer a better insight into the usability of the formulation. For the testing to be valid, the building would have to be occupied, which may not be straightforward. Furthermore a comparison with other intelligent approaches (such as fuzzy logic control and multi-agent systems) would be useful. In this work, the MPC algorithms are only compared to more traditional (though modern) approaches, over which an advantage may be more likely.

An extension of the formulation to include additional types of equipment and different system architectures would greatly improve the prospects of the technology. In particular, the adaptation of the techniques to suit a system with forced cooling (such as a HVAC system) or with additional actuators (such as controlled shading) could move the research in a valuable direction. Other factors could also be integrated into the strategy, such as fuel prices and lighting levels. With the inclusion of a CHP in the strategy, the electrical load demand could also offer significant benefits, while posing interesting control challenges (particularly in terms of the problem complexity). Co-operation of distributed energy systems is becoming an evermore expansive field of research.

Fault Tolerance

To achieve a more fault tolerant control strategy, the fault detection and isolation techniques introduced in Chapter 6, as well as associated reconfiguration techniques should be incorporated into the MPC formulation. For an FDD strategy to be effective, a characterisation of all possible fault classes which could affect the system should be carried out. Given the large number of faults which could occur in the system, an investigation into the faults which have the biggest impact on the control would be useful to determine the set of faults for which action should be taken.

In terms of reconfiguration post-fault, an investigation into adaptive methods may be suitable. As data-driven approaches are used to derive the initial zone models, updating the zone models using measured data after a fault has occurred may be sensible in certain scenarios. Fault tolerant and reconfigurable control is a wide ranging field. As such, many alternative techniques could also be applied to the MPC strategies developed.

Appendix A

Boiler Efficiency Curves

The boiler used in the Nimbus Centre is a *Viesmann Vitocrossal 200* series. The manufacturer's efficiency curves for the boiler are given here. These curves are used to develop efficiency surfaces that are included in the simulation platform (developed in Chapter 3) and the prioritised-objective MPC approach formulated in Chapter 5.




Figure A.1: Boiler efficiency curves (*Viesmann Vitocrossal 200*) given for different return temperatures

Appendix B

Quantum-behaved Particle Swarm Optimisation

In this section, additional detail of the quantum-behaved particle swarm optimisation (Q-PSO) algorithm (as implemented in Chapter 3) is provided. Whereas using traditional PSO methods, each agent is assigned a position and a velocity at each iteration, with a Q-PSO algorithm, a probabilistic approach is taken.

The state of each D -dimensional agent i in the swarm of Ω agents, is depicted as a wavefunction, given here at the q^{th} iteration as $\Psi(z_i^q)$. This differs from standard PSO formulations, in which an exact position is used. Each agent of the Q-PSO swarm can be associated with a local attractor point denoted p_i^q . The agents move through the D -dimensional search space with a δ -potential well located for each at the associated local attractor point. The wavefunction for the $(q + 1)^{th}$ iteration is governed by a Laplace distribution (Sun et al. 2004), given for the j^{th} dimension as:

$$\psi(z(i, j)^{q+1}) = \frac{1}{\sqrt{L_{i,j}^q}} e^{-\frac{|z_{i,j}^{q+1} - p_{i,j}^q|}{L_{i,j}^q}} \quad (\text{B.1})$$

where $L_{i,j}^q$ is a parameter associated with the diversity of the distribution. Hypothetical probability density functions (ψ^2) are shown in Fig. B.1, illustrating the effect of varying the diversity parameter. A larger value for $L_{i,j}^q$ implies a wider distribution.

The Monte Carlo inverse transform method is used to obtain the position of the j^{th}




Figure B.1: Probability density functions of hypothetical wavefunction for different diversity parameters

dimension of the i^{th} agent from the associated wavefunction, using the equation:

$$z_{i,j}^{q+1} = p_{i,j}^q \pm \frac{1}{2} L_{i,j}^q \ln \left(\frac{1}{u_{i,j}^{q+1}} \right) \quad (\text{B.2})$$

where $u_{i,j}^{q+1}$ is drawn from a uniform distribution in the range $(0, 1)$. This is analogous to collapsing the wavefunction of the agent by measuring the position. The diversity parameter is calculated as:

$$L_{i,j}^q = 2\alpha |M_{b_j}^q - z_{i,j}^q| \quad (\text{B.3})$$

where α is a specified contraction/expansion coefficient (chosen for the purposes of stability to be < 1.781 (Sun et al. 2012)) and $M_{b_j}^q$ denotes the mean-best position of the agents in the swarm, given where P_{b_i} denotes the previous best position of the i^{th} agent as:

$$M_{b_j} = \sum_{i=1}^{\Omega} (P_{b_{i,j}}) \quad (\text{B.4})$$

The result is that while the wavefunction associated with each agent remains centred on the previous best value of the agent, the diversity tends to be larger and so the distribution tends to be wider for agents that are further away from the mean-best position. Furthermore, the influence of these agents on the mean-best value has the effect of slowing the convergence of the closer agents, allowing for a greater level of exploration to be carried out in the region around the mean-best. Using a traditional PSO formulation, these further away agents exert no influence unless they can independently improve upon the global-best. This limits the exploration potential in a particular search-space.

References

- Sun, J., Fang, W., Wu, X., Palade, V. & Xu, W. (2012), ‘Quantum-Behaved Particle Swarm Optimization: Analysis of Individual Particle Behavior and Parameter Selection’, *Evolutionary Computation* **20**(3), 349–393.
- Sun, J., Feng, B. & Xu, W. (2004), ‘Particle swarm optimization with particles having quantum behavior’, *Proceedings of the 2004 Congress on Evolutionary Computation (IEEE Cat. No.04TH8753)* **1**, 325–331.