NUMERICAL OPTIMISATION OF BUILDING THERMAL AND ENERGY PERFORMANCE IN HOSPITALS

Andrew Richard Cowie

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

The University of Leeds

School of Civil Engineering

March 2017

Intellectual Property and Publication Statements

The candidate confirms that the work submitted is his own, except where work which has formed part of jointly authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

The work detailed in chapter 5 has appeared in publication as follows:

> Cowie, A.C; Noakes, C.J; Sleigh, P.A. and Toropov, V.V. 2013. Meta-model based Optimization of Building Thermal Performance incorporating Local Comfort Analysis. In: *Proceedings of the Fourteenth International Conference on Civil, Structural and Environmental Engineering Computing, Cagliari, Sardinia*; 3 – 6 September 2013.

The candidate was responsible for the work detailed in the publication, the creation of the publication and the presentation of the publication. The three subsequent authors took advisory and oversight roles.

Elements of the work detailed in chapters 4, 6 and 7 have appeared in publication as follows:

> Cowie, A.C; Noakes, C.J.; Sleigh, P.A. and Toropov, V.V. 2015. A meta-model based method for thermal comfort and energy use optimisation at a local resolution. In: *Proceedings of the Fourtheenth International Conference of the International Building Performance Simulation Association, Hyderabad, India*; 7 – 9 December 2015.

The candidate was responsible for the work detailed in the publication and the creation of the publication and presentation. Due to unforeseen circumstances the candidate was unable to deliver the presentation at the conference, so this was delivered by a colleague (Jon William Hand B.Sc., M. Arch., PhD) who was briefed by the candidate. The three subsequent authors took advisory and oversight roles.

Acknowledgements

It is a necessity and a pleasure to acknowledge, and to dedicate this thesis to, the three parties without whom it would not exist.

Firstly my supervisors Dr. Andrew Sleigh and Prof. Vassili Toropov, and particularly my principal supervisor and mentor Prof. Catherine Noakes, who has been the architect of so many of the opportunities laid before me. I feel sure I would have far less passion for academia if not for the tireless support and enthusiasm of these three, as well as many others at the University of Leeds. My sincere gratitude goes out to all of you. Secondly my parents and family in general, for reasons that need no words. Finally Sam, for putting your dreams on hold while I pursued mine.

Furthermore I would like to extend many thanks to the denizens of the ESP-r mailing list, for their invaluable guidance in navigating dragon-infested code. Also, grateful thanks to both the DeDeRHECC Project for the use of their monitoring data, and the estates department of Bradford Royal Infirmary for the construction and operation information. Finally I would like to gratefully acknowledge the support of EPSRC in funding this work under grant EP/G029768/1 “Integrated Design of Hospital Wards for a Safe and Sustainable Patient Environment”.

Abstract

This thesis details the development and testing of a metamodel-based building optimisation methodology dubbed thermal building optimisation tool (T-BOT), designed as an information gathering framework and decision support tool rather than a design automator. Initial samples of building simulations are used to train moving least squares regression (MLSR) meta-models of the design space. A genetic algorithm (GA) is then used to optimise with the dual objectives of minimising time-averaged thermal discomfort and energy use. The optimum trade-off is presented as a Pareto front.

Adaptive coupling functionality of the building simulation program ESP-r is used to augment the dynamic thermal model (DTM) with computational fluid dynamics (CFD), allowing local evaluation of thermal comfort within rooms. Furthermore, the disconnect between simulation and optimisation induced by the metamodeling is exploited to lend flexibility to the data gathered in the initial samples. Optimisations can hence be performed for any combination of location, time period, thermal comfort criteria and design variables, from a single set of sample simulations; this was termed a “one sample many optimisations” or OSMO approach. This can present substantial time savings over a comparable direct search optimisation technique. To the author’s knowledge the OSMO approach and adaptive coupling of DTM and CFD are unique among building thermal optimisation (BTO) models.

Development and testing was focussed on hospital environments, though the method is potentially applicable to other environments. The program was tested by application to two models, one a theoretical test case and one a case study based on a real hospital building. It was found that variation in spatial location, time period and thermal comfort criteria can result in different optimum conditions, though seasonal variation had a large effect on this. Also the sample size and selection of design variables and their ranges were found to be critical to meta-model fidelity.

Contents

List of Figures

List of Tables

Abbreviations

BRE	**B**uilding **r**esearch **e**stablishment
BSI	**B**ritish **s**tandards **i**nstitute
BTO	**B**uilding **t**hermal **o**ptimisation
CIBSE	**C**hartered **i**nstitute of **b**uilding **s**ervices **e**ngineers
CFD	**C**omputational **f**luid **d**ynamics
DOE	**D**esign **o**f **e**xperiments
DoH	**D**epartment **o**f **H**ealth
DTM	**D**ynamic **t**hermal **m**odelling
EIA	**E**nergy **i**nformation **a**dministration
ESP-r	**E**nergy **s**ystems **p**erformance – **r**esearch edition
LPPD	**L**owest **p**redicted **p**ercentage **d**issatisfied
MLSR	**M**oving **l**east **s**quares **r**egression
NHS	**N**ational **h**ealth **s**ervice
OSMO	**O**ne **s**ample **m**any **o**ptimisations
PMV	**P**redicted **m**ean **v**ote
PPD	**P**redicted **p**ercentage **d**issatisfied
SAP	**S**tandard **a**ssessment **p**rocedure
T-BOT	**T**hermal **b**uilding **o**ptimisation **t**ool
TSENS	**T**hermal **sens**ation

Nomenclature

θ	Temperature	°C
θ_o	Operative temperature	°C
θ_a	Air temperature	°C
θ_r	Mean radiant temperature	°C
θ_{db}	Dry bulb temperature	°C
v	Velocity	m/s
Θ_{rm}	Exponentially weighted running mean of mean daily external temperature	°C
Θ_{ed-1}	Previous days mean external temperature	°C
Θ_{rm-1}	Previous days running mean	°C
α	A constant	
$\theta_{o,max}$	Upper limit on operative temperature	°C
$\theta_{o,min}$	Lower limit on operative temperature	°C
ρ	Density	kg/m^3
C	Specific heat capacity	J/kg °C
V	Volume	m^3
t	Time	s
i	An identifier	
n	Number of items	
K	Heat flow conductance	W/°C
q	Energy	W
E	Error	
A, B, C, T and Z	Matrices	
u	Vector component of velocity	m/s
x, y and z	Axial distances	m
P	Pressure	Pa
μ	Viscosity	Pa s
μ_t	Eddy viscosity	Pa s
S	Source term	

λ	Thermal conductivity	W/m °C
k	Turbulence kinetic energy	J s^2/kg m^2
ε	Dissipation rate of turbulence energy	J s^3/kg m^2 s
σ_t	Turbulent Prandtl number	
l	Length	m
d	Thickness	m
R	Thermal resistance	m^2 °C/W
U	U-value	W/m^2 °C
SM	Solar metrics	W/m^2
F	Objective function	
D	Design variable	
Q	Volumetric flow rate	m^3/s
h	Hour	integer
E	Energy load	kWh
M	Metabolic rate	W/m^2
W	Effective mechanical power of work being performed	W/m^2
I_{cl}	Clothing insulation	m^2 °C/W
f_{cl}	Clothing surface area factor	
h_c	Convective heat transfer coefficient	W/m^2 °C

Chapter 1: Introduction

1.1 Chapter Overview

This chapter provides an introduction to the thesis, introducing key concepts and presenting a rationale for the work detailed herein. First, a general introduction is given to provide context for the present work. Following this four individual sections introduce and briefly discuss thermal comfort, the particular challenges of hospitals, building simulation and numerical optimisation. The next section brings these together and introduces the concept of building thermal optimisation. Finally, the aims and objectives and scope of the present work are detailed.

1.2 Requirements of a Building

Even since the early stages of the evolution of mankind, we have sought respite from the crueller excesses of the climate on Earth. Without shelter we could not have survived. As time has passed our shelters have evolved with us, from a necessity for survival to vast and in some cases luxurious structures that dominate the skyline in many places worldwide. They no longer merely provide respite from the weather, but often now seek to maintain comfortable conditions for the occupants at all times.

To accomplish this level of thermal control, in all but the most cunningly designed structures active energy input is generally necessary. Whether heating or cooling, via the air or through radiation, we have developed myriad methods of keeping ourselves at a pleasant temperature within buildings; some more energy efficient than others. With such a range of options available to building designers, and the daunting array of design parameters that must be decided upon, it can sometimes be difficult to ascertain the most energy efficient way of doing things both for new builds and retrofit. Until relatively recently this was not generally considered a major problem; we had the capability to keep ourselves comfortable and that was what really mattered. However with the growing global concerns of climate change and the increasingly diminishing availability of non-renewable resources upon which we have relied to supply our energy needs for so long, it is now more important than ever to improve energy efficiency and reduce the amount of energy we use wherever possible.

According to the US Energy Information Administration (EIA) in 2013 40% of the total national energy use of the USA was consumed in residential and commercial buildings [2014]. From the most recent surveys of these buildings, in 2009 in residential buildings 48% of energy use went toward heating and cooling [EIA, 2013]. In 2003 in commercial buildings the figure was 44% [EIA, 2008]. Similarly, 40% of national energy use in the UK was consumed by buildings according to UK government literature [Gov.uk, 2014a]. Clearly, a significant proportion of energy use in developed countries goes towards maintaining thermal comfort in buildings, so improving the energy performance of these buildings has a large potential to reduce global energy consumption.

So then we have two contrasting objectives; to maintain thermal comfort for occupants at all times, whilst at the same time reducing energy use. Often these objectives are mutually exclusive; improved thermal comfort necessitates more energy input, and less energy input results in poorer thermal comfort. However, by considering all pertinent design parameters and applying a formalised framework to establish the optimum trade-off between these two objectives, this goal can be realised as far as is practicable.

Such a procedure is dependent on the ability to accurately predict the performance of buildings; this is accomplished through building simulation (discussed in greater detail in section 1.5). Building simulation models the physical processes in buildings and hence allows prediction of the performance of buildings including thermal comfort and energy use. It is widely used in modern design and assessment processes; for example building simulation is effectively mandated for conformance to UK building performance criteria in the form of a standard assessment procedure (SAP), developed by the building research establishment (BRE) [2014a].

The remit of the present project is to use building simulation to evaluate thermal comfort and energy performance of buildings, particularly hospitals (discussed in greater detail in section 1.4), and to apply numerical optimisation procedures (discussed in greater detail in section 1.6) to examine the optimum trade-off between the two.

1.3 Thermal Comfort

Generally speaking, thermal comfort is a measure of a person's comfort in terms of temperature. If a person is too cold or too warm, they are not in a state of thermal comfort. It is dependent on the individual under consideration and the environment to which they are subjected.

The input parameters to thermal comfort models are generally established using building simulation (introduced in section 1.5). There are four main environmental indices known to significantly influence thermal comfort [Fanger, 1972]:

- Air temperature,
- Air velocity,
- Radiant temperature,
- Humidity.

These indices describe the thermal environment to which the occupant is subjected, and hence are taken as inputs to most thermal comfort models. Environmental parameters are detailed further in section 2.2.3. It is worth noting here that under typical indoor conditions these parameters will often vary spatially, so an environment that provides a comfortable environment at one location within a room may not do so at another point in the room.

Other than environment indices, there are incidental parameters that also affect thermal comfort. Probably the most widely used of these are:

- Metabolic rate,
- Level of clothing.

These are generally taken as averaged equivalent values based on research; the literature in the area is reviewed in section 2.2.

An important factor to consider is what exactly constitutes thermal comfort. This has been an active field of research for over 50 years, as it is not quite as simple as saying “this environment is thermally comfortable, and this one isn’t”. The literature on this subject is examined in detail in section 2.2; this section introduces the key concepts.

Thermal comfort is not an objective term; it varies from person to person even under the exact same conditions. For example, an environment that a healthy young person might find thermally comfortable will probably not be so for an unwell elderly person. The input variables to thermal comfort models often attempt to encapsulate some of this subjective information; for example metabolic rate is an input to many thermal comfort models and this varies depending on age and the activity level of the occupant, among other factors. Also some thermal comfort models attempt to encapsulate variation within populations implicitly, through empirical studies. The most well-known example of this is the predicted percentage dissatisfied (PPD) metric developed by Ole Fanger [1972].

Evaluation of thermal comfort is further complicated by the fact that it is somewhat difficult to express it quantitatively. It is all very well to say that an environment is, for the majority of a population comfortable, or slightly uncomfortable, or intolerably cold etc., but how is this then translated to a numeric metric? To define quantitative metrics, in most cases discretized qualitative scales are used, generally ranging from "extremely cold" to "extremely hot" with "comfortable" in the middle. Thermal comfort models generally predict these metrics by modelling the physical processes which affect human thermal comfort, both internal (eg. vasodilation) and external (eg. heat transfer to and from the skin)

1.4 Thermal Comfort and Energy in Hospitals

As the present work is focussed specifically on the building performance of hospitals, a brief section is included here to introduce the particular challenges of these unique establishments.

The particular hospital under consideration for the present project, Bradford Royal Infirmary, consumed a total of 43,698 MWh of energy in the year 2012-2013, of which less than 0.1% came from known renewable sources, according to the health and social care information centre (HSCIC) [2014]. BRE estimates that up to 20% of this energy is wasted [2014b], and according to the department of health (DoH) 44% of the energy use in a typical UK hospital is used for heating [2006]. Clearly then by improving the energy efficiency of the hospital building stock significant potential exists for reducing energy use, particularly from non-renewable sources, and hence significantly reducing carbon footprint. Given that the NHS is the largest public sector contributor to the UK's CO_2 emissions (with an overall footprint of approximately 18.6 million tonnes of CO_2) [BRE, 2014c], this goal sits nicely in line with current UK government targets to reduce carbon emissions by 80% from the 1990 baseline by 2050 [Gov.uk, 2014b].

Hospitals are more often than not populated largely by people who are unwell in some form or another. This presents its own challenges in relation to thermal comfort, as many illnesses are known to affect thermal sensation. Clearly it is not practical to assume specific illnesses in the modelling, so this can only really be taken into account by erring on the stringent side in terms of limits on thermal comfort. Moreover, these unwell occupants are in stark contrast to the staff of the hospital who have a far higher activity rate and are also more likely to be in a generally good state of health. We have then two contrasting main groups of occupants, with potentially wildly differing thermal requirements; and this is not to mention other groups such as visitors. The literature on this subject is reviewed in section 2.2.5.

A further complication that arises from the presence of many unwell occupants is the issue of cross-contamination; it is rather counter-productive to patient well-being if they are constantly exposed to other pathogens by the hospital environment. This is well studied in the literature, for example by King et al. [2013]. As a result of this ventilation requirements in hospitals are often far more stringent than in other buildings; the DoH has even published specific guidance on ventilation for hospitals [2007].

Moreover, hospitals tend to be housed in large buildings not dissimilar to commercial buildings such as office blocks (though perhaps lacking the more recent fashion for glazed facades). However the occupancy profiles for hospitals is completely different to many other large buildings, being inhabited and in operation 24 hours a day, 365 days a year. Practically, this somewhat simplifies the modelling of the building as casual gains and thermal controls can be assumed to be in effect constantly, whereas in many other buildings this assumption would generally not be valid.

As a final note, hospitals in the UK generally have a need for retrofit in order to maintain effectiveness in a changing climate, particularly with regard to overheating during summer as identified by Short et al. [2012]. Whilst increased provision of active thermal control systems is the most obvious solution to this, it is noted by Short et al. that this is incongruous with the need to reduce energy consumption. Attaining optimal energy performance with the active thermal control systems within hospitals is vital to minimise this conflict of interests.

1.5 Building Simulation

The literature of building simulation is reviewed in section 2.3, and the general methodology of the specific software used in the present work is described in section 3.2. This section introduces the concept of building simulation and briefly discusses its application to the present work.

Building simulation provides a very important step in the building optimisation process described above, as it provides the means to evaluate the objectives for any given permutation of the design parameters. This is generally accomplished by constructing a model of the building in question, and then simulating the performance of the building and its thermal control systems under appropriate climatic conditions. Performance of the building can be defined in terms of many metrics, including thermal comfort of occupants and energy use of the thermal control systems. The model can usually include any number of various modelling domains, which can encompass anything from heat transfer through the building envelope to

detailed simulation of the air flow within rooms. In the context of the present project, the salient functions of building simulation are twofold: firstly to evaluate thermal comfort under the dynamic conditions present in a building, and also to evaluate the energy use of thermal control systems (eg. radiators and HVAC).

For the present work, the building simulation software ESP-r was selected. This was used to simulate the thermal and energy performance of single-bed hospital rooms under varying climatic conditions, providing the data needed to optimise room performance.

1.6 Numerical Optimisation

In order to formally define numeric optimisation, the terminology must first be established. First, the metrics by which each solution is ranked are termed **objective functions**. These functions must be dependent on the design parameters, which are represented as variables in the objective functions and are hence termed **design variables**. Design variables and/or objective functions are often subject to limitations, for example simply for practicality. Consider a radiator temperature design variable; are temperatures of -5°C or 150°C sensible? The optimisation process will not know that these values are not practicable unless you tell it so; these are termed **constraints**. Note that constraints may be more complex than simply defining limiting values, constraints on further functions of the design variables are possible.

Given one or more objective function(s) that depend on the design variables, and subject to the constraints, the purpose of numerical optimisation is to find the minimum possible value of the objective function(s). Note that minimization is merely a convention within the field, maximisation is also possible with by simply inverting the objective functions (i.e. objective function to be minimised is 1/F, where F is the function to be maximised). Where there is more than one objective function, this becomes rather more complex as there are number of ways of handling this. This is discussed in more detail in section 2.4.3.

This is accomplished in numerical optimisation using some framework of decision making. The simplest example is random trial-and-error; try one solution, then select another at random and try that, if it is better than the first then discard that one, and so on. Once no better solution can be found in a certain number of tries you have probably found the optimum. However being entirely random, this method may not be efficient for covering the design space, and may take some time to find the optimum. Most numerical optimisation algorithms take a more structured approach to selecting the next solution, based on the available

information from previous solutions. There are a wide variety of algorithms available for this; for the present work a genetic algorithm was selected. This is described in detail in section 3.3.

Numerical optimisation provides the automated analysis and decision making functionality in the process of optimising building performance. Given the inputs and outputs from the building simulation, this allows each permutation of the design variables to be ranked in terms of the objective functions, and hence decisions made regarding their usefulness. It is worth noting that whilst optimisation is an umbrella term used for many different things, the term numerical optimisation is used throughout the present thesis to represent the formalised and automated process of optimising performance based on the objective functions, in terms of the design variables. The literature on this subject is reviewed in section 2.4.

Numerical optimisation is extremely useful in the context of building design, as it allows the design process to be turned on its head. Instead of building design starting from the design parameters, it allows design to begin from the objectives. So instead of building design following the process of identifying the parameters to be designed, and then finding some permutation of these parameters that meets the performance criteria, design may begin by establishing how the building would perform ideally, and then using numerical optimisation to find a permutation of the design parameters that gives this outcome. Essentially, it removes the human trial-and-error process from design, and when implemented properly ensures that the design always gives ideal performance. However, this is extremely sensitive to how the design problem is abstracted to an optimisation problem; the optimisation will only solve the problem it is given, it has no knowledge of the design intent beyond that encapsulated in the objective functions.

1.7 Building Thermal Optimisation

Bringing the concepts introduced in the previous sections together, we arrive at the research area within which the present project resides. This is the process of using building simulation to evaluate the objective functions of thermal comfort and energy use from the design variables, and then using numerical optimisation to determine the optimum trade-off between them. In the present thesis this is termed **building thermal optimisation (BTO)**. This is a growing research area; the literature on the subject is reviewed in section 2.5.

Whilst BTO is simple in its idealised form, in practice it is rather complex. Firstly, building modelling is not perfect, and there are many assumptions and simplifications generally made in building simulation. No building model available today can hope to take into account every

single variable and physical process upon which the performance of the building depends. It is imperative then that any model used for BTO is shown to be at least broadly representative of the actual performance of the building. Secondly, purely in terms of practicality and logistics it is a significant undertaking. With the leaps and bounds in computing power available today, it has only recently become practical to run a great many building simulations in a reasonable amount of time, as is required for BTO. In the context of the present work, the modelling was undertaken at a high level of detail with a correspondingly high computational requirement, so this issue was particularly acute.

In the present project, a BTO procedure is developed and implemented computationally. The resulting program is termed throughout the thesis **thermal building optimisation tool (T-BOT)**. In particular, T-BOT was intended to address a number of research gaps in the literature; this is further detailed in section 2.5.2.

1.8 Aims and Objectives

The aim of the thesis is to develop and implement a BTO methodology that allows quantitative evaluation of an optimum trade-off between thermal comfort and energy use, evaluated at a local level within individual spaces. The performance of this methodology will be explored by applying the program to progressively more complex cases.

The specific objectives of the project are as follows:

1. Establish the state-of-the-art art in the fields of thermal comfort, building simulation and optimisation and key research gaps in the field of BTO.
2. Identify appropriate approaches to modelling thermal comfort, building simulation and optimisation and combine these to create a BTO methodology.
3. Develop a computational program to apply the BTO methodology.
4. Assess the practicality and fitness-for-purpose of the BTO program through application to a simplified test case.
5. Further explore the performance of the BTO program in terms of practical application though a case study based on a real hospital environment.

1.9 Scope of Work

The following is a brief summary of the thesis content:

Chapter 1: Introduction

This chapter sets out the context for the project, and introduces key concepts of building performance, thermal comfort, building simulation and numerical optimisation, and briefly details the process of combining these elements to accomplish BTO. Also the specific challenges of hospitals are briefly discussed.

Chapter 2: Literature Review

This chapter reviews the literature of key aspects of the project; thermal comfort, building simulation, numerical optimisation and BTO. Implications are discussed and research gaps are highlighted.

Chapter 3: Simulation and Optimisation Tool Methodologies

This chapter discusses the functionality and details the methodology of the 3rd party software used to implement the simulation and optimisation elements of the project. The software used for this was ESP-r and Altair HyperStudy.

Chapter 4: Initial Model Development

This chapter introduces the building model that was used for testing and development in the initial stages of the project, and describes the developed BTO methodology. This is done in terms of each individual element; construction of metamodels to predict model responses for the purposes of optimisation, sampling of the design space using building simulation, and the optimisation parameters. Furthermore the implementation of the methodology is detailed.

Chapter 5: Proof of Concept

This chapter details the proof of concept study that was performed to assess the practicality and fitness-for-purpose of the methodology. This was performed with the test model introduced in chapter 4. Results are reported and discussed, and key outcomes are highlighted.

Chapter 6: Case Study Model Development

This chapter details further development performed in response to outcomes of the proof of concept study and to enable a case study of a real hospital environment.. The case study site and the representative building model are described, and the performance of the model is evaluated against measured data.

Chapter 7: Case Study Results

This chapter presents some results of the case study, demonstrating key outcomes. A large number of cases were optimised, so for the purposes of brevity not all results are shown in this chapter. Results for all cases are presented in Appendix B.

Chapter 8: Discussion

This chapter discusses the work as whole, and draws out key outcomes in terms of the results, program performance, use of the program in practice and further work.

Chapter 9: Conclusions

This chapter gives the conclusions drawn from the project, summarising the outcomes identified in Chapter 8.

Chapter 2: Literature Review

2.1 Chapter Overview

This chapter contains a review of the literature in any area of interest to the present project. Sections 2.2-2.4 review the individual subject areas that must be combined to allow building thermal optimisation (BTO). They detail, respectively:

- Thermal comfort and the models developed for its evaluation.
- Building simulation.
- Numerical optimisation.

Section 2.5 then reviews how these elements are brought together to enable BTO in the literature, and highlights research gaps, particularly those addressed by the present work.

2.2 Thermal Comfort Modelling

2.2.1 Origins of the field: Fanger, Gagge and Stolwijk

This is an aspect of building design that was generally judged by rules of thumb and empirical testing, until it experienced a revolution in the early 1970's. Around this time the growing popularity of computerised modelling motivated the collation of this empirical data into various models of human thermal comfort with differing levels of complexity. Generally considered the most important of these were the works of Ole Fanger [1972], summarized in his thesis first published in 1970, and subsequently republished in 1972.

Fanger applied logical reasoning to develop three conditions for thermal comfort, and then used his own and others' research to quantify each expression mathematically. The result was a comprehensive thermal comfort equation in terms of variables that fell under three categories; clothing, activity, and environment. Clothing variables are simplified values to represent the overall thermal properties of different types of clothing, and were originally introduced by Gagge et al. [1941]. Gagge is another very important figure in thermal comfort research; his work is covered in greater detail later. Activity variables depend on what tasks the people under consideration are doing at the time. As most buildings are designed and used for a specific purpose, these can often be approximated with some accuracy by reference tables. Rather more research was done in this area once Fanger identified its importance; an

example is the work of Garg et al. [1978]. Environmental variables are indices of the surroundings such as air temperature, humidity, etc.

This comfort equation is widely regarded as one of the most significant developments in the modelling of thermal comfort, and is still commonly used in modelling applications today. However, further to this, Fanger [1972] also collected together quite considerable amounts of experimental data on thermal comfort, and used statistical analysis to develop models of how people will react to thermal environments. Fanger developed three such models, each building on the last:

- Predicted Mean Vote (PMV): A seven point scale from "hot" to "cold" via "thermally neutral". Thermal neutrality corresponds to optimum thermal comfort as obtained by the comfort equation. This essentially is a method of assessing an environment for human habitation.
- Predicted Percentage of Dissatisfied (PPD): Based on statistical data from more than a thousand subjects, PPD is directly calculated from PMV. It is useful as it will generally be the thermally dissatisfied that are likely to control adaption measures (opening windows, turning on fans or heaters, etc.), so it is often more useful to impose limits on PPD than PMV.
- Lowest Possible Percentage of Dissatisfied (LPPD): This is an attempt to encompass thermal non-uniformity into the framework. Fanger reasoned that if you determine PPD at various points around a room, and take the average, you are very unlikely to get the minimum possible value of 5% (occurring at thermal neutrality) even if a PMV value of 0 is obtained for the room as a whole (i.e. modelled as a uniform environment). This discrepancy is down to thermal non-uniformity; the property of a room to make its internal environment differ between different locations within it. By varying the average temperatures used to calculate PMV, a minimum value of this average PPD can be obtained, which is the LPPD. The difference between the LPPD and the minimum PPD of 5% can then be used as an indicator for the non-uniformity of the room.

LPPD is rarely used these days, but PMV and PPD are still popular methods for assessing thermal environments, more than 40 years after their advent. The principal disadvantage of these models is that they were all formulated for steady-state conditions. Fanger was well aware of this limitation, and he states that [Fanger, 1972]:

> *"In practice, minor fluctuations in one or more of the variables will often occur, but as long as the mean value, taken over a suitable time interval, is reasonably constant,*

quasi-steady-state conditions exist, under which the comfort equation can also be employed."

A. P. Gagge is another key figure in thermal comfort research. Gagge et al. were researching thermal comfort along similar lines some 30 years before Fanger published his landmark thesis [1941], and their work informed much of Fanger's research. As a measure of this, almost an entire page of references in Fanger's book [1972] is devoted to papers with Gagge as the primary author. Gagge's teams work culminated in the development of his oft-called "two-node" model of the human thermoregulatory system [Gagge et al., 1971, 1986], which is more suitable for transient modelling of thermal comfort [ASHRAE, 2009].

The two node model is rather more detailed in terms of modelling human responses to their thermal environment. Many of the principles used originated from the research of Stolwijk and Hardy [1966], two names which often appear alongside Gagge in the literature. It considers the body as two concentric zones, an outer layer and a core. The outer layer represents the skin, and is considered as roughly 10% of overall body mass. All metabolic heat generation is assumed to occur at the core node. Modelling in this fashion allows use of various control functions to model processes such as vasoconstriction and dilation, sweating, and shivering. Gagge developed his model into two indices similar to Fanger's; the thermal sensation TSENS and thermal discomfort DISC. TSENS works on the same scale as PMV except that it has an extra two terms at the extremes, ±4 "very hot/cold" and ±5 "intolerably hot/cold". DISC works on a similar 11 point scale, but with various levels of discomfort from "intolerably hot" to "intolerably cold" via "comfortable". TSENS is calculated from the deviation of mean body temperature from hot and cold limits, which are defined by the limits of evaporative regulation. DISC is equal to TSENS when below the cold limit, and is calculated based on the body's ability to sweat out the excess heat otherwise.

As succinctly put in the ASHRAE handbook [ASHRAE, 2009], "The [two-node model] shows that comfort and thermal sensation are not necessarily the same variable ...". This is quite an important factor, as it is rather difficult to quantify this difference. Studies by Stolwijk & Hardy [1966], Hardy & Stolwijk [1966] and also together with Gagge [1967] highlighted that this difference is particularly evident during periods of steeply transient thermal conditions. These experiments generally involved transferring subjects from a thermally comfortable environment to a warm or cold one, and then back again. The results showed that when going from comfortable to uncomfortable environments, discomfort increases as the body slowly changes temperature and triggers various responses, notably shivering and sweating. On the other hand, when the transients were reversed (i.e. uncomfortable to comfortable), the

subjects felt comfortable much sooner than their thermal state would indicate. This indicated a distinct psychological element to thermal comfort that was not well understood at the time.

Another similar study by Gagge et al. [1969] investigated thermal transients caused by periods of sustained exercise, under various environmental conditions. The results showed that both comfort and thermal sensation follow mean body temperature as it rises at the beginning of the exercise, as would be expected. But after around 30 minutes of steady activity comfort becomes dominated by sensors in the skin, and loses its relation to metabolic rate. Warm discomfort is then mainly governed by regulatory mechanisms; sweating and blood flow.

Whilst Gagge's model has the advantage over Fanger's in that it is valid for transient conditions, there are many other limitations of both. For example, both models consider only average temperatures, so largely asymmetric thermal conditions can limit their accuracy. Both researchers, among many others at the time, investigated these effects with experimental studies. Some of them were rather extreme; notably Hall and Klemm [1967] exposed lightly clothed subjects to radiant temperatures differing by up to 110°C between the front and rear of the body. This resulted in a skin temperature difference of up to 9-10°C; this was postulated as a maximum allowable skin temperature variation for comfort. Gagge's experiments in the area [Gagge et al., 1965(1); Gagge et al., 1965(2)] focused on using asymmetric radiant comfort to balance uncomfortable air temperatures. The results were in agreement with Hall and Klemm's; even at air temperatures of 10°C the subjects were able to attain thermal comfort using radiant heat sources above them. It was found that large heat losses at parts of the body not exposed to the radiative heating were offset by the large heat gains at the irradiated parts, apparently without causing discomfort. Findings to the same effect were reported by other studies of the day, notably Schlegel and McNall [1968] and McNall and Biddison [1970]; although slight discomfort due to asymmetry was reported during a test where one wall was 30°C warmer than ambient temperature. Research by Fanger et al. [1985] suggested rather more stringent limits to maintain comfort; 10°C asymmetry for a cool wall, and 23°C for a warm wall. While it was clear that asymmetrical fields do not usually induce discomfort within practical limits, it was equally recognized that they can effect thermal sensation.

One potential way to get around this difficulty is to construct a more geometrically correct comfort model, considering each part of the body as a separate node. In general these models originated with Stolwijk [1971], with some significant work also being done by Wissler [1964]. Stolwijk's model is even more detailed than the two-node one, and of the few that were available at the time it was almost certainly the most complex. It considers the head, torso, arms, hands, legs and feet as separate segments, each consisting of 4 concentric layers; core, muscle, fat and skin. The thermal characteristics of each layer in each segment are modelled

individually, along with the heat transfer processes between them. Further to these 24 elements, blood is taken into account by a single core compartment which acts by convection to spread heat around the body. The model showed fairly good prediction of experimental physiological results, but was not specifically related to thermal comfort. It was originally developed in the context of manned space flight, where conservation of heat is clearly a rather more pressing concern than down on terra firma.

2.2.2 Further Development of Thermal Comfort Models

From these initial models sprang numerous variations over the next 40 years or so. The models can generally be grouped into three categories; "standard" models that predict comfort based on whole-body averages, multi-segment models such as previously described, and adaptive models.

Standard models

Within the realm of standard models, Fanger remained a leading figure up to his untimely passing in 2006. Air quality and ventilation were not considered in any particular detail in the original model, and he published many papers subsequently on these subjects [eg. Fanger et al., 1988; Fanger, 1989; Fanger, 2000 and Fanger and Toftum, 2002]. Zolfaghari and Maerefat [2010a] have done considerable work improving Gagge's two-node model by combining it with Pennes' bio-heat equations [1948]. Also, a few empirical models have sprung up fairly recently, for example Rollins et al. [2006] and Zhang & Zhao [2008]. Most of the recent development however has been focussed on multi-segment models.

Multi-segment models

A large number of multi-segment models have been developed over the last 40 years (more so in the last 10-20 years though); it seems that most research institutions concerned with thermal comfort developed their own variations at some stage. Two models that are often cited in the literature are the so called "Berkeley comfort model" and the "AUB model". The Berkeley model, developed at UC Berkeley, began in 2001 with the work of Charlie Huizenga et al. [2001], but is now largely attributed to his colleague Hui Zhang. The current incarnation of the Berkeley model is given in Zhang's three-part article published in Building and Environment [Zhang et al., 2010a, 2010b and 2010c]. Another member of the Berkeley team, Edward Arens, has also published significant supporting research [Arens et al., 2006a, 2006b and 2009]. The Berkeley comfort model is based upon Stolwijk's original multi-segment model [1971] with a few notable improvements; the model allows unlimited segments, considers clothing as a separate layer, and considers all thermoregulatory mechanisms explicitly [Huizenga et al,

2001]. This is a highly advanced and rather versatile model, and shows great potential for future development.

The AUB model was first developed by Salloum et al. [2007] as a collaboration between various universities in Beirut, and then subsequently improved by Al-Othmani et al. [2008]. This model is also based on Stolwijk's original [1971], but is rather more of a departure from Stolwijk's methods. In general, the AUB model gives much greater consideration to blood flow; the nodes for each segment used are core, skin, arterial blood and vein blood. The thermodynamic processes are modelled largely by calculating the blood perfusion into the tissue, which allows the heat transfer to be modelled more or less by first principles [Salloum et al., 2007]. The original model draws heavily on anatomical research and physiological data published by Avolio [1980]. The revised version [Al-Othmani et al., 2008] however included consideration of direction within the skin nodes, and a more recent blood-flow model for smaller vessels [Olufsen et al., 2000], which increased the accuracy somewhat.

Direct comparisons of the Berkeley comfort model and the AUB model were not found in the literature, and an in depth analysis is beyond the scope of this work. However it can be said in summary that the AUB model may be less accurate when comfort is governed more by sweating than body temperatures. Also, the Berkeley comfort model has been directly related to both local comfort [Zhang et al., 2010b] and whole-body comfort [Zhang et al., 2010c], making it on the whole probably a more versatile model.

Aside from these two, there are numerous other multi-segment models. Zolfaghari and Maerefat [2010b] have done other work in this area similar to that aforementioned ([2010a]); extending Gagge's two-node model to a three-node one to represent clothed and unclothed parts of the body. Also contained in this paper is a concise summary of human thermal models to date, which agrees with a similar (but earlier) review by Arens and Zhang [2006; section 16.8]. The following is a selection of models based on these literature reviews, supplemented by a further review by the present author. This is not an exhaustive list of multi-segment models, such a study is beyond the scope of this broad review.

A relatively early model was proposed by Fu [1995] in his PhD thesis , based on earlier work done by Jones and Ogawa [1992 and 1993] and Smith [1991]. This model was interesting because it used a finite-element approach to model the shape of the body. Other early work includes papers published by Imre et al. [1988], Xu and Werner [1997], and Yigit [1998], which all present thermal models for clothed humans. Murakami et al. [2000] and Tanabe et al. [2002] took the approach of combining various modelling methods as opposed to creating an entirely new comfort model; the work of Tanabe et al. informed the development of the AUB

model to some extent. Fiala et al. began developing their model over ten years ago [2001], and have recently developed it further [2011], including the development of various new comfort indices. Yi and Fengzhi et al. have published several papers detailing thermal models with a focus on modelling clothing [Yi et al., 2004; Fengzhi and Yi, 2005]. Kaynakli et al. built on the work of Yigit [1998] and combined it with Gagge's two-node model [Kaynakli et al., 2003; Kaynakli and Kilic, 2005]. Finally, Miyanaga et al. [2001] have done some work on relatively high accuracy geometric models of humans.

Adaptive models

Adaptive models are a type of empirical model based solely on observations of the adaptive actions of mankind; although it is worth noting that many of the core equations of other types of models are approximated empirically. They are based on the principal that humans have a certain degree of ability to bring themselves to thermal comfort with adaptive actions; opening a window perhaps, or taking off a jacket [ASHRAE, 2009]. Furthermore, longer term adaption has been shown to be a major factor, particularly with regard to the outside climate [Humphreys and Nicol, 1998]. Adaptive models generally predict temperatures or ranges thereof at which people will be comfortable, usually with a running mean of outside temperature as the principal input. The validity of adaptive models tends to increase with the ability of the occupants to control their own thermal comfort; personal fans for example, operable windows, or a relaxed dress code [ASHRAE, 2009]. The approach originated from ASHRAE funded research begun in the mid-80s [De Dear and Brager, 2002], and was formalised to practical models for inclusion into ASHRAE standards by Humphreys and Nicol [1998] and De Dear and Brager [1998]. Later, Nicol and Humphreys [2010] developed models for UK standards (see section 2.2.4).

As adaptive models are entirely empirical, deriving them requires substantial experimental data. There are two main databases of such data; the more comprehensive of these is the ASHRAE RP-884 database, founded in 1995 [De Dear and Brager, 2002]. This consists of many thousands of datasets from 160 separate office buildings across 4 different continents. The second was gathered by the Smart Controls and Thermal Comfort (SCATs) project, funded by the European Union and lasting from 1997 to 2000 [McCartney and Nicol, 2002]. This project was specific to Europe, gathering data from 26 office buildings in 5 European countries [Nicol and Humphreys, 2007]. Both databases include results from questionnaires of building occupants and measured environmental data; these form the basis of adaptive comfort theory and provided the data to develop the empirical models on which it relies. The fact that both databases only contain data for office buildings means that generally the approach is only fully

applicable to office buildings, though some models report to be applicable to buildings of a similar type and use (see section 2.2.4).

An interesting alternative adaptive model has been proposed recently by Haldi and Robinson [2010]. This model uses statistical equations to simulate conscious adaptive actions, and Haldi and Robinson also suggest ways that it could be applied to other elements of comfort.

Applications in Building Simulation

Despite the significant advances in thermal comfort modelling that have been accomplished in the last 30 years, in building simulation it is still the norm to use the well-established models with greater precedent. It may be said that the PMV model and variants thereof are still the most widely used comfort models even today. As such, implementation of advanced deterministic comfort models within building simulation software is limited, and it was decided that implementing them was somewhat outside the remit of the present project. However, the more advanced models do show potential to significantly advance the modelling of local thermal comfort; further work such as this that could follow from the present project is discussed in detail in section 8.4.

2.2.3 Environmental Parameters

Here the various environmental parameters used as metrics of thermal comfort, or inputs thereto, are introduced and explicitly defined.

The most important of these is operative temperature. Operative temperature is a measure of the temperature experienced by an occupant, taking into account the radiant temperature field, air temperature and air velocity. It is defined in CIBSE Guide A [CIBSE, 2006] as:

$$\theta_o = \frac{\theta_a\sqrt{10v} + \theta_r}{1 + \sqrt{10v}} \tag{2.1}$$

Where θ_o is operative temperature (°C), θ_a is air temperature (°C), θ_r is the mean radiant temperature (MRT) (°C) and *v* is the air velocity (m/s). At air speeds below 0.1m/s this is assumed to be:

$$\theta_o = \frac{\theta_a + \theta_r}{2} \tag{2.2}$$

Air temperature is straightforward, but MRT requires definition itself. MRT is a measure of the aggregate "radiant temperature" experienced by an occupant, taking into account all surfaces in view. All surfaces in view of one another are at all times instantaneously exchanging thermal radiation, and the net flow of energy between them is proportional to the temperature difference. By calculating a mean of the temperatures of surfaces around the

occupant, weighted by the relative view factors (ie. the proportion of the 3-dimensional view field that is occupied by the surface in question; this by necessity uses the concept of solid angles), the aggregate flow of radiant heat between the occupant and their environment may be simply represented.

In a more technical sense, MRT is defined in CIBSE Guide A [CIBSE, 2006] as:

> *'The uniform surface temperature of a radiantly black enclosure in which an occupant would exchange the same amount of radiant heat as in the actual non-uniform space. (see BS EN ISO 7726 for derivation). (Note: if the surface temperatures of the internal surfaces of the enclosure are unequal, mean radiant temperature varies throughout the enclosure and depends upon the posture and orientation of the occupant.)'*

This is perhaps easiest to conceptualise by considering the device by which operative temperature is commonly measured; a globe thermometer. Ordinary thermometers are not suitable for measuring operative temperature if the air temperature and MRT differ significantly [CIBSE, 2006]. However by placing the thermometer at the centre of an approximately 40mm diameter sealed sphere made of metal or plastic, painted grey or black to approximate the correct reflectivity, the instrument is then turned into a globe thermometer and its responses should closely approximate operative temperature [CIBSE, 2006]. This is because the sphere effectively acts as a practical form of the MRT averaging procedure.

It is worth noting that in previous revisions of CIBSE Guide A, operative temperature was known as dry resultant temperature. The two are identical in all but name; the name was simply changed in the latest update to CIBSE Guide A in the interests of international consistency of nomenclature [CIBSE, 2006].

Air temperature and velocity can be averaged over the whole room, but it is far more apt to evaluate them locally when considering local thermal comfort. This can be accomplished through CFD; see section 3.2 for an in-depth discussion of this.

2.2.4 Thermal Comfort Models in Legislation

The role of building regulations is to ensure that buildings are fit for purpose within reasonable limits. As maintenance of thermal comfort is a principal role of buildings, it is incumbent that said regulations include some guidance on evaluating the thermal performance of buildings. As regulations are generally set by different governments largely independently, they differ from country to country. The following brief review is principally focussed on UK regulations, commonly termed British Standards.

The most recent British Standard on thermal comfort in buildings is BS EN 15251 [BSi, 2008]. This document codifies the criteria and methods whereby the indoor environment of buildings is assessed; this includes thermal comfort, but also takes into consideration other factors such as air quality and noise. The section of particular interest to the present work is Annex A, entitled "Recommended criteria for the thermal environment". This section contains two distinct sets of criteria; the first based on Fanger's PMV/PPD model (codified in BS EN ISO 7730 [BSi, 2006]), and the second based on an adaptive comfort model.

The PMV/PPD criteria are intended for use in buildings with mechanical heating and cooling, when the internal temperature can be assumed to be well controlled and hence not significantly transient when occupied. These criteria are presented as tables of recommended maximum and minimum operative temperatures for the summer and winter seasons respectively, for a range of building types and comfort categories. The comfort classes are defined by maximum deviation of PMV (or PPD) from ideal comfort, and represent cases of different sensitivity to thermal conditions. Table 2.1 shows these classes as defined in BS EN 15251 [BSi, 2008].

Category	Thermal state of the body as a whole	
	PPD %	Predicted Mean Vote
I	< 6	-0,2 < PMV < + 0,2
II	**< 10**	**-0,5 < PMV < + 0,5**
III	< 15	-0,7 < PMV < + 0,7
IV	> 15	PMV<-0,7; or +0,7<PMV

Table 2.1: Recommended comfort sensitivity categories (taken from BS EN ISO 15251 [BSi, 2008])

The recommended temperatures are obtained by assuming values for all other inputs to the PMV model (ie. Clothing, activity, air velocity and humidity), and back-calculating to obtain a range of temperatures that satisfy each comfort category using the limits given in table 2.1. Clothing is assumed based on season; 1 clo for winter and 0.5 clo for summer. Activity is assumed based on the use of the building; for example sedentary activities (such as sitting at a

desk working on a computer) have an assumed value of 1.2 met. Air velocity is assumed to be small, ie. less than 0.1m/s, and humidity is assumed to be 50%, which are both reasonable assumptions for most indoor conditions. As well as these assumptions, design loads and design weather conditions are assumed as detailed in BS EN ISO 15927- 4 [BSi, 2005(1)] and 5 [BSi, 2005(2)]. Clearly there are a great many assumptions involved in deriving these sets of criteria, however they are fine examples of how the PMV model can be used to provide broad guidance on the thermal design of buildings. The criteria are not hard limits for internal temperatures, rather recommendations that should ensure a reasonably comfortable environment under most conditions.

The second set of criteria, based on adaptive comfort, is intended for use in buildings without mechanical cooling; adaptive alternatives to mechanical heating are not considered. This is presumably because most adaptive actions available to people in normal daily life are typically focussed on cooling down in summer; operable windows and desk fans are common examples. The criteria are presented similarly to the PMV set, as limits on operative temperature with 3 separate comfort categories. However, the limits are in this case taken as empirical linear functions of an exponentially weighted running mean of daily external temperature (shown in equation 2.3). Equations 2.4 and 2.5 give these empirical functions, and Figure 2.1 shows them graphically.

$$\Theta_{rm} = (1-\alpha)\Theta_{ed-1} + \alpha.\Theta_{rm-1} \tag{2.3}$$

$$\theta_{o,max} = 0.33\Theta_{rm} + 18.8 + c \qquad c = \begin{cases} Category\ I: & 2 \\ Category\ II: & 3 \\ Category\ III: & 4 \end{cases} \tag{2.4}$$

$$\theta_{o,min} = 0.33\Theta_{rm} + 18.8 - c \qquad c = \begin{cases} Category\ I: & 2 \\ Category\ II: & 3 \\ Category\ III: & 4 \end{cases} \tag{2.5}$$

Where:

- Θ_{rm} is the exponentially weighted running mean of mean daily external temperature
- Θ_{ed-1} is the previous days mean external temperature
- Θ_{rm-1} is the previous days running mean
- α is a constant between 0 and 1 (recommended value 0.8 [BSi, 2008])
- $\theta_{o,max}$ is the upper limit on operative temperature
- $\theta_{o,min}$ is the lower limit on operative temperature

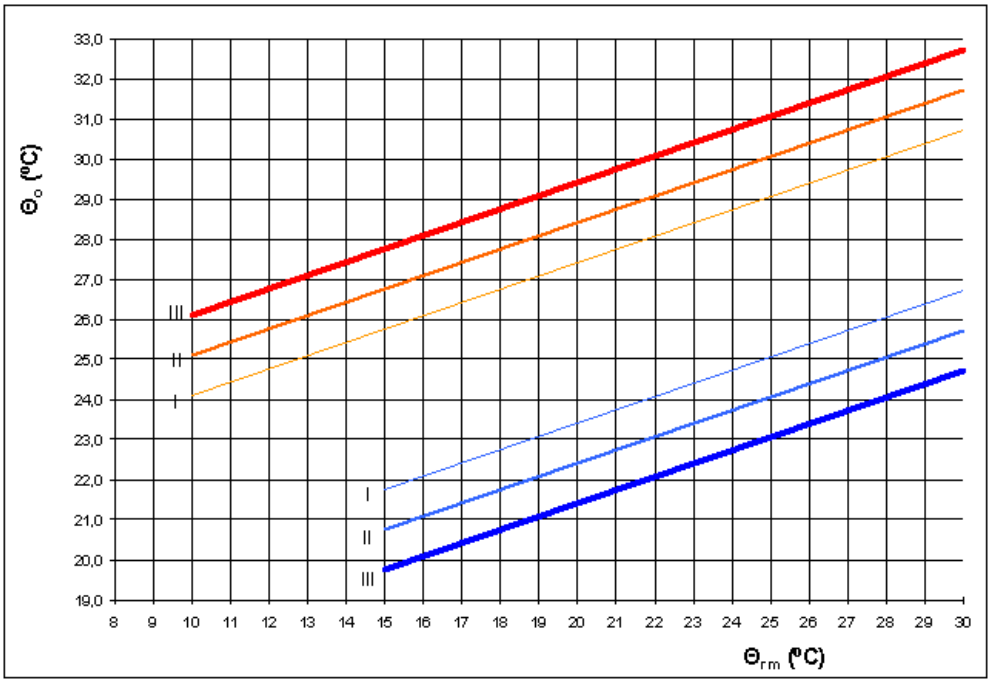

Figure 2.1: Adaptive limits on operative temperature as a function of running mean external temperature (taken from BS EN 15251 [BSi, 2008])

Adaptive comfort is a fairly recent development, and there is not enough experimental data available to define the empirical relationships for all contexts (see the *Adaptive comfort* sub-section of 2.2.2). In view of this it is made clear in BS EN 15251 [BSi, 2008] that:

> *"The (summer) temperature limits presented in this annex are primarily based on studies in office buildings. Nevertheless, based on general knowledge on thermal comfort and human responses, the assumption can be made that the limits may apply to other (comparable) buildings with mainly sedentary activities like residential buildings."*

An important companion to legislative standards is the literature of governmentally recognised professional institutes. In the UK the most relevant of these is the Chartered Institute of Building Services Engineers (CIBSE); the analogous organisation in the USA is the American Society of Heating, Refrigeration and Air Conditioning Engineers (ASHRAE). These organisations publish their own material, with content broadly similar to the standards. However they and their literature tend to have a slightly different remit than standards, focussing a little more on best practice than regulation.

In CIBSE Guide A [CIBSE, 2006] guidance is given in line with that given in BS EN 15251, but in more detail. The tables of recommended operative temperature limits are more extensive, having a section specific to hospitals. In the present project, comfort limits were adapted from here and used as comfort criteria in objective functions (see section 4.6.1 for details). The

guide also contains a section on adaptive comfort, though it is less explicit than BS EN 15251 in that the comfort criteria are stated to be only applicable to office buildings, and further states that "there are insufficient data to provide similar advice for houses" [CIBSE, 2006].

Finally, a third party whose legislature is relevant to the present work is that of the Department of Health. This government body presides over healthcare in the UK, and has published its own guidelines relating to thermal comfort in hospitals. In particular, HTM03-01 [DoH, 2007] sets out broad guidance for the design of ventilation systems in hospitals; this document gives the recommendation of 6 ACH ventilation rate for patient areas, and states that:

> "*Calculations and thermal modelling should be undertaken to ensure that, during the summertime, internal temperatures in patient areas do not exceed 28°C (dry bulb) for more than 50 hours per year.*"

However in recent research this criterion has proven to be less appropriate than the adaptive temperature criteria of BS EN 15251 [BSi, 2008] (see section 2.2.5).

2.2.5 Thermal Comfort in Hospitals

In section 1.3 the specific challenges of maintaining thermal comfort in hospitals were introduced and briefly discussed. In this section, the available literature on the subject is reviewed in the context of the present project.

Field Studies

In terms of field studies of thermal comfort in hospitals, there were a number found in the literature, but none that were wholly applicable to hospital bedrooms in the UK. A large amount of the literature presents results for operating theatres, whilst ward areas and bedrooms seem to be less well studied. Of those that were found, few were conducted in climates comparable to the UK. Only a few studies were found that were within Europe, one in Italy and one in Belgium.

De Giuli et al. [2013] performed a field monitoring study of three wards at Padua General Hospital in Italy, and examined correlation between PMV calculated from measured conditions and results of surveys of both staff and patients; these are shown in Figure 2.2. The primary Y axis units of "T," are not explicitly specified in the paper, but from context it is assumed this represents air temperature. Results revealed non-uniform correlation, in that the PMV index did not always accurately predict the thermal comfort of occupants. However, the survey results revealed a marked difference between comfort of patients and staff as shown in the

insets of Figs. 2.2(a & c); this was broadly in agreement with trends in PMV (calculated from measured air temperature) as shown in the main panels of Figs. 2.2(a & c), induced by differences in activity and clothing inputs. Patients were found to be much more accepting of the indoor environment, and de Giuli et al. [2013] suggested that this may be related to the different role the hospital plays to the two parties. Patients generally experienced only a short stay and were likely to be preoccupied with the purpose of their visit to the hospital; however for staff it was a permanent place of work, so indoor conditions are likely to have a far greater capacity to affect general well-being.

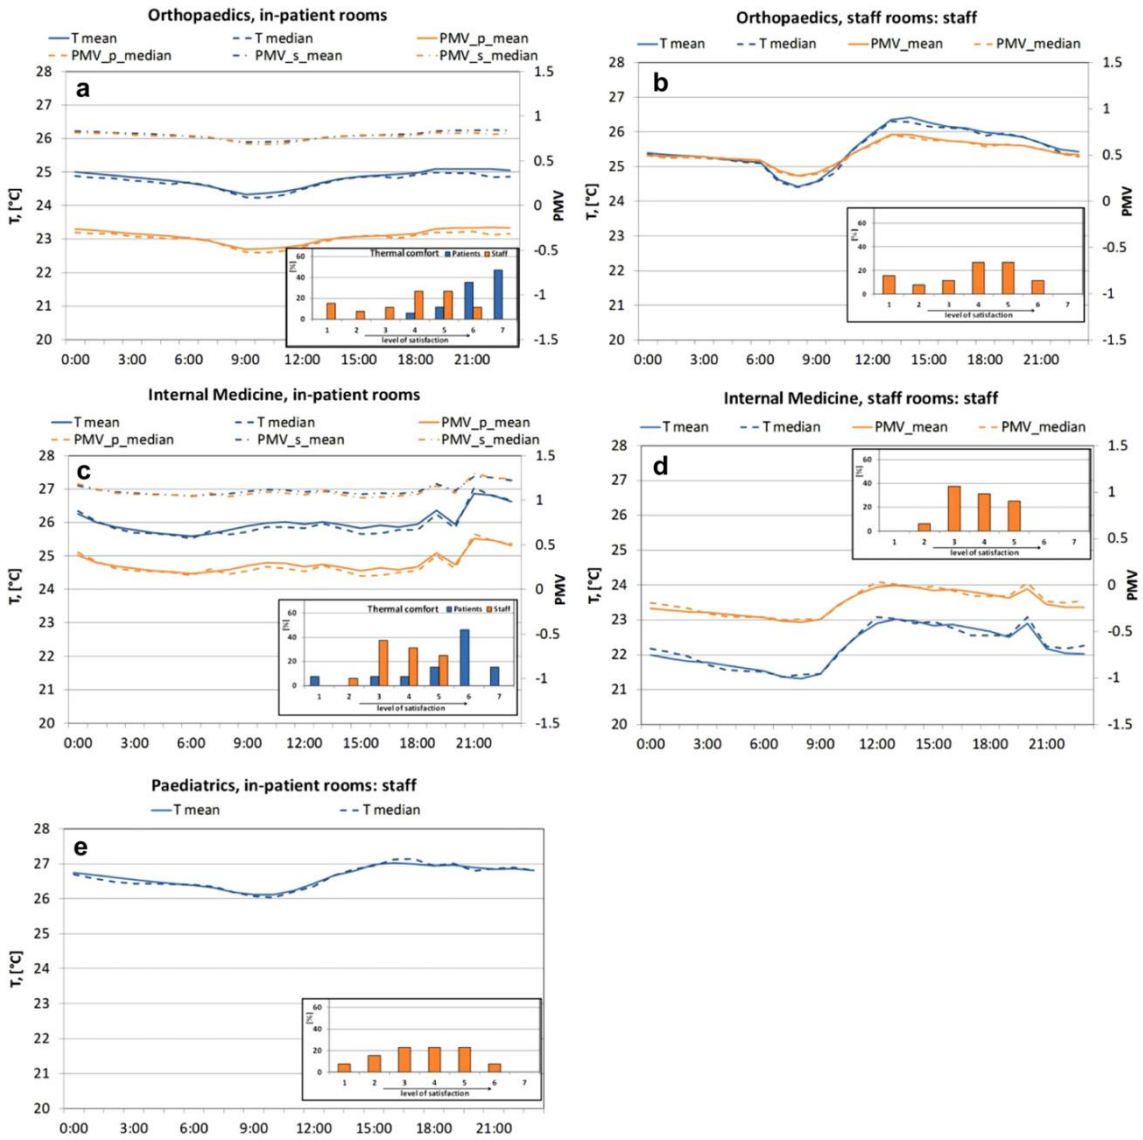

Key: PMV_p= patients' PMV; PMV_s= staff's PMV.

Figure 2.2: Comparison of subjective and objective comfort in an Italian hospital, taken from de Giuli et al. [2013].

Verheyen et al. [2011] performed a broadly similar survey in a clinical healthcare facility in Belgium, though this only included patients and was entirely focussed on thermal comfort as opposed to general indoor environment as in de Giuli et al's [2013] study. The study was

intended to highlight any differences in thermal response of patients in different wards of the hospital; results are shown in Figure 2.3. Meaning of data points in this figure are not explicitly reported in the paper, but from context it is assumed that points with solid centres are mean values, points without solid centres are recorded values, and bars indicate spread. The analyses found very strong correlation between calculated PMV and survey results, though subjective thermal sensation was systematically underestimated by objective measurements, indicating that patients had a lower neutral temperature than fully healthy individuals. However Verheyen et al. [2011] noted that this discrepancy could also be due to incorrect assumptions for activity and clothing parameters, among other possibilities. The study also found that patients in the neurology ward had uniquely skewed perception of thermal comfort; this was reconciled by reasoning that neurology caters to injuries of the brain, bone marrow and nervous system, all important elements in human thermoregulation. A notable difference in the calculation methodology employed by Verheyen et al. [2011] (as opposed to that of de Giuli et al. [2013]) was that PMV for patients lying on a bed was not calculated by the traditional method, but rather a specialised variant thereof proposed by Lin and Deng [2008].

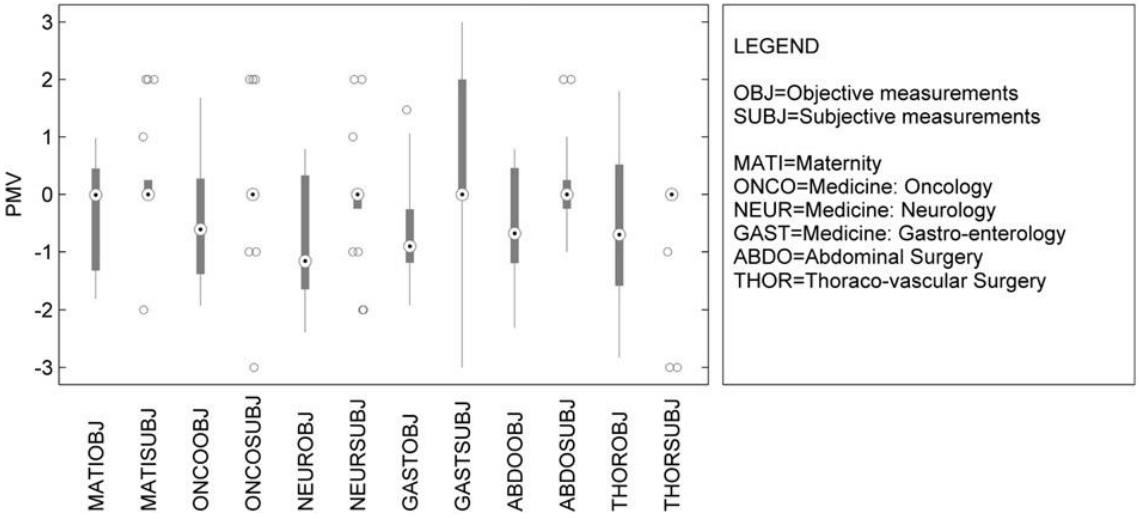

Figure 2.3: Comparison of subjective and objective PMV in a hospital in Belgium, taken from Verheyen et al. [2011].

Outside of Europe, some substantial field studies of thermal comfort in hospitals have been performed. Perhaps the most rigorous of these was a study of a medical centre in central Taiwan, Japan [Hwang et al., 2007]. This study compared subjective data gathered from questionnaires of patients to the thermal comfort criteria given in ASHRAE Standard 55 [ASHRAE, 2004], including direct comparison of subjective and objective PMV; these results are shown in Figure 2.4. In this instance, the PMV model overestimated patients' thermal sensitivity, exhibiting a rather greater gradient as can be seen in Figure 2.4. Analyses of this

data also revealed that the predicted neutral temperature was 2.2°C higher than that derived from the subjective data. The study was also able to draw a number of other conclusions;

- Only 40% of the environments assessed conformed to the criteria of ASHRAE Standard 55 [ASHRAE, 2004], however most of the areas that did not meet the standards only violated the humidity constraints. This was not borne out in the subjective results, as more than three quarters of participants rated their thermal sensation as neutral and were not concerned about humidity. Such a large discrepancy highlights the importance of cultural and climatic adaption in thermal comfort.
- Analyses of the year-round range of acceptable comfort temperature suggested that whilst the criteria of Standard 55 were approximately appropriate, occupants expected a more moderate temperature.
- Chi-squared tests of parametric significance revealed that physical strength had a very strong influence on thermal sensation, whereas gender, age and acclimatization did not.
- When compared with studies of office buildings in Hong Kong, higher neutral and preferred temperatures were observed. Another notable discrepancy between the results sets was that patients in the hospital preferred a higher temperature than thermal neutrality, which is contrary to findings from the office buildings.

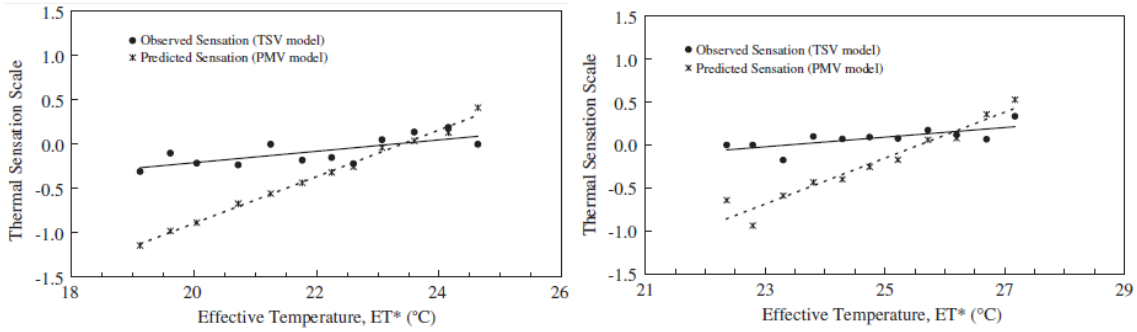

Figure 2.4: Comparison of subjective and objective PMV during winter (left) and summer (right) in a hospital in Japan, taken from Hwang et al. [2007].

There were a few other field studies found in the literature, mainly in significantly warmer climates than the UK. Both Azizpour et al. [2013] and Yau and Chew [2009] have studied thermal comfort in Malaysian hospitals; both papers adopt a similar approach to the studies detailed above, taking and comparing objective and subjective measurements. Azizpour et al. [2013] found reasonable correlation between PMV and subjective results, with R^2=0.88 using linear regression. Also the study found good agreement with adaptive theory, in that the subjective neutral temperature was significantly higher than that predicted by the PMV model.

Conversely, Yau and Chew [2009] found poor agreement between PMV and subjective results. Similarly to the findings of Hwang et al. [2007], only 44% of the environments assessed conformed to the criteria of ASHRAE Standard 55 [ASHRAE, 2004]; overall occupant satisfaction aligned well with this at 49%. It is interesting to note that comfortable temperatures found by Yau and Chew [2009] were consistently significantly higher than those found by Azizpour et al. [2013]. Both papers consider only staff, so this could be indicative of a difference in thermal control mode between the sites; acceptable comfort tends to be more lenient and fall more in line with adaptive criteria in free-running buildings than air-conditioned [BSi, 2008]. Unfortunately, very little clear information was found in the papers concerning the control modes of the case study sites.

Pourshaghaghy and Omidvari [2012] studied thermal comfort in a hospital in the city of Kermanshah, Iran. The results showed good agreement between subjective and objective measurements, though the calculations tended to slightly underestimate PPD. The study also found that PPD in males was slightly higher than in females, though this difference was less than 5%. Hashiguchi et al. [2005] conducted a study of a hospital in Japan, though the dataset was rather smaller than that of Hwang et al's [2007] study. Results were in agreement with previously described literature in that patients tended to be much more thermally comfortable than staff. In this particular instance though, comfort was governed much more by humidity, as there was a significant problem with excessively low humidity in the building which manifested as physical symptoms in occupants. Finally, Skoog et al. [2005] studied thermal comfort in a Swedish hospital in both summer and winter. This study also corroborated the large difference in thermal sensation between staff and patients.

In the context of the present study, these works are particularly salient in terms of the choice of comfort metric. Whilst some papers found good agreement between subjective and objective PMV and/or PPD, others found quite the opposite. Clearly there is a good possibility that the PMV and PPD indices do not necessarily represent comfort in hospitals particularly well. Furthermore, the difference in subjective comfort between hospital staff and others was highlighted a number of times. With this in mind it would seem to make sense to "roll back" the treatment of thermal comfort, closer to first principles. Instead of evaluating PMV and/or PPD in a particular simulation, the various environmental metrics from which PMV and PPD are calculated could be evaluated instead. In light of the conclusions from the literature, this has several advantages. Firstly by examining simpler metrics of thermal comfort, e.g. operative temperature, potential inaccuracies introduced by the modelling method of PMV can be avoided. However this is not to say that PMV should not be considered at all; a second benefit

would be that PMV and PPD can be evaluated for any combination of metabolic rate and clothing from the results of a single simulation.

Modelling Studies

Aside from field studies, there are also a small number of studies that examine hospitals in a more theoretical modelling context. A significant portion of this comes from a recent EPSRC funded research project entitled Design and Delivery of Robust Hospital Environments in a Changing Climate (DeDeRHECC). The present project had links to DeDeRHECC, and made use of data collected therein to inform the case studies; this is covered in more detail throughout chapter 5. Whilst temperature monitoring data was collected during the DeDeRHECC project, it was generally not twinned with subjective results; rather the data was collected to examine the current thermal environment and to calibrate models of the buildings in question. An important driver to the DeDeRHECC project was the development and implementation of the UKCP09 climate predictions program [UKCP, 2009], which provides probabilistic projections of future climate data.

Three papers relating to comfort in hospitals have so far resulted from DeDeRHECC; Lomas et al. [2012], Lomas and Giridharan [2012] and Short et al. [2012]. In the first of these, Lomas et al. [2012] investigated the performance of nightingale style wards at Bradford Royal Infirmary. This paper had a specific focus on examining the inherent property of these wards to resist overheating during periods of high external temperature, and how the changing climate would affect this property. Further to this, various “light-touch” measures were suggested to improve comfort with minimal, if any additional energy demand. The study concluded that nightingale wards are surprisingly resilient to summertime overheating, despite the vintage of the design philosophy. However, modelling predictions suggested that the studied wards in their current form would begin to overheat in extreme weather as soon as 20 years from now. The refurbishment recommendations developed during the study were predicted to mitigate this problem to various degrees, such that even a modest retrofit regimen would bring overheating within acceptable limits right up to 2080. These refurbishment measures included improved insulation, shading and natural ventilation; slow ceiling fans were found to further reduce the risk of overheating. Lomas and Yi (2009) published a paper a few years earlier that was a precursor to this work.

Lomas and Giridharan [2012] applied a similar methodology to a tower building at Addenbrooke’s Hospital, notably having a hybrid ventilation strategy as opposed to naturally ventilation. This study also made use of the UKCP09 climate projections [UKCP, 2009] to develop future climate data and examine resilience to the changing climate. The monitoring

results showed that the ward spaces were primarily regulated by opening windows, and that the current temperature ranges were generally within the adaptive criteria of BS EN 15251 [BSi, 2008]. However in other spaces of the hospital where provision for occupant thermal control was rather poorer, thermal conditions were somewhat uncomfortable. Similarly to the conclusions of Lomas et al. [2012] fans were found to be an effective yet relatively inexpensive aid to resilience; the modelling predicted that fans would maintain comfort in the wards even in extreme weather up to 2080. Interestingly, Lomas and Gridharan [2012] also recommended that single bed hospital rooms should be carefully examined to assess potential to improve thermal comfort and energy efficiency simultaneously; this falls within the remit of the present project.

The final of the three most relevant papers from DeDeRHECC is a companion paper to Lomas and Giridharan's [2012], being published sequentially in the same issue of the same journal. Therein, Short et al. [2012] provide much of the background information of the Addenbrooke's tower block, and also assess its resilience both currently and in the future; however this is done in somewhat more of a general sense than is Lomas and Giridharan's [2012] work, with a greater focus on CO_2 emission targets. Short et al. [2012] support the findings of Lomas and Giridharan [2012] in that the adaptive criteria of BS EN 15251 [BSi, 2008] was found to be more appropriate than fixed limits, such as those given in HTM03-01 [DoH, 2007]. The paper reports that from 2030 onwards, mechanical intervention may be required in extreme years to prevent overheating; it is further noted that in this situation the ventilation requirement of 6 ACH detailed in HTM03-01 [DoH, 2007] and the NHS energy demand targets of 55-65 GJ / 100 m^3 may be mutually exclusive.

Aside from these papers from DeDeRHECC, very little relevant literature was found that falls into this category. Khan et al. [2012] presented a model-based optimisation study of ventilation that is relevant to hospital premises; this is reviewed with other building optimisation papers in section 2.5. Also, Adamu et al. [2012] examined natural ventilation strategies for a new ward at Great Ormond Street Hospital through CFD modelling; the study concluded that significant improvements could be attained over the conventional design, and particularly advocated use of ceiling-based natural ventilation to provide natural personalised ventilation.

2.3 Building Simulation

Similarly to thermal comfort, prior to the 1960's and 70's building energy prediction was generally accomplished using well established simplified methods such as the degree day and the bin method. These methods generally relied on categorizing external temperatures, then using simplified calculations to determine energy consumption for the sum of each of these categories [Beausoleil-Morrison, 2000]. Whilst these steady-state methods have stood the test of time, and still give valuable approximations of building heat loss an energy performance, they do not take into account the detail of the processes which take place in buildings or the detail of external climate conditions which affect energy performance. These are mainly transient processes which are by definition difficult to encompass in steady-state modelling.

As the power and availability of computing resources increased over the 60's and 70's, so the state of the art in the area could move forward to more explicit modelling of thermal processes in buildings. This development began in the mid to late 60's, with models often employing a Loads-Systems-Plant (LSP) methodology to separate out building simulation into stages that could be more easily handled by the computing resources available at the time [Beausoleil-Morrison, 2000]. Stephenson and Mitalas [1967] first introduced the response factor method in 1967, which greatly facilitated the move toward transient modelling. This method used the principal of superposition to decompose the highly complex heat transfer mechanisms into a series of much simpler equations, allowing transient processes to be characterized and solved within a defined and manageable framework.

The beginnings of what we would now recognize as building simulation arose in the 1970's with the introduction of heat balance methods. As with most research, the impetus for these developments can be traced back to the allocation of funding; around this period the US Department of Energy allocated over $1 billion to research into energy conservation [Hong et al., 2000]. These were the first attempts to model heat transfer explicitly, without using approximated weighting factors. The methodology is also intrinsically time-stepped, furthering development toward fully transient simulation. An example of early implementations is the work of Kusuda [1976]. This method was later extended by the application of discretization and simultaneous solution, allowing building modelling to move toward a finite-volume approach, thus removing the need for many of the assumptions necessary for the response factor method [Clarke, 1977]. This also allowed extension of the heat balance methodology to other critical building components such as plant. It is this work by Clarke [1977] that ESP-r (the building simulation program used for the present project) in particular traces its roots back

too, in its first incarnation ESP. Around the same time other building simulation programs appeared which are still in use today eg. DOE-2 [Hong et al., 2000].

Development of building simulation continued with the introduction of airflow modelling. Originally this was entirely separate from building energy simulation, and progressed concurrently along two distinct paths; application of CFD [Jones and Whittle, 1992] and more macroscopic building airflow networks [Feustel and Dieris, 1992]. Early integration of airflow networks into general building simulation did not occur until the 1980's [Walton, 1983]. However whilst this allowed inclusion of the effects of airflow within the thermal simulation, they could not actually simulate the airflow from boundary conditions without user input of flow parameters; this remained exclusive to CFD for some time [Beausoleil-Morrison, 2000].

It is at this point in the early 1990's that building simulation began gaining support and widespread use in industry [Hong et al., 2000]. This was largely driven by growing global awareness of past and pending environmental crises, and hence the need to improve energy efficiency. Throughout the 1990's the International Energy Agency (formed in the 1970's in response to the oil crisis) performed research into this area [eg. Lomas et al., 1994- a, b & c], culminating in the (continuing) development of the BESTEST suite of building simulation program validation tests [Judkoff and Neymark, 1995].

The 1990's also saw the beginning of the coalescence of CFD and building simulation. This would seem like a logical coupling; CFD requires accurate boundary conditions to be effective, and building simulation is well equipped to supply these dynamically. It began with researchers including elements of building simulation into CFD codes, to allow specification of boundary conditions in a building context (e.g. external conditions and fabric as opposed to surface temperatures) [e.g. Holmes et al., 1990]. However, greater flexibility in this respect can be attained by approaching the problem from the other direction; integrating CFD functionality into more evolved building simulation models, or coupling distinct CFD and building simulation programs to achieve the same effect. Over the last 20-25 years a number of researchers have implemented such systems. For example, Schaelin et al. [1993] coupled the CFD package PHOENICS with the multi-zone air flow simulator COMIS. This coupled system was then used to investigate the validity of averaging contaminant sources spatially within zones. More recent examples include the work of Fan and Ito [2012], who coupled the CFD package FLUENT [ANSYS Inc., 2015] and the building simulation program TRNSYS [SEL, 2013] to investigate the placing of supply and exhaust openings in a ventilation system. Zuo et al. [2014] proposed an alternative, coupling an alternate method of solving the CFD equations termed fast fluid dynamics (FFD) with Modelica and achieving faster than real time coupled

simulation. A further alternative aimed at building control applications using reduced order CFD models has been proposed by Kim et al. [2015].

As noted by Beausoleil-Morrison [2000], there remains a common caveat among most coupled approaches. The interaction between the two domains is at the building surfaces, but this area is rather difficult to model for both domains. CFD conventionally employs empirical wall functions to represent near-wall flow, and similarly building simulation often uses empirically determined convection coefficients to represent heat transfer between surfaces and air. However this approximated approach may become far better informed using the information available in a coupled system, for example using local flow patterns to calculate appropriate convection coefficients. A comprehensive example of such a procedure is the adaptive conflation controller (ACC) developed by Beausoleil-Morrison [2000] to intelligently couple CFD and building simulation in ESP-r. No other comparable coupling procedures were found in the literature. The CFD in ESP-r was originally implemented by Negrão [1995] some years earlier; Beausoleil-Morrison's work essentially gave the coupling procedure some adaptability. By performing a preliminary CFD simulation before the actual simulation, general flow patterns can be assessed and the coupling procedure can be tailored to prevailing conditions each time it is invoked.

In general, building simulation has continued to evolve in line with changing requirements of both industry and research. Commercial packages such as IES-VE [IES Limited, 2014], Tas [EDSL, 2012] and ECOTECT [Autodesk, 2014] have emerged to better serve industry, offering much more user-friendly interfaces and more low-end and data entry automation along with the consultancy and technical support services that commercialization allows. Conversely, many research tools such as ESP-r have remained free and open-source, benefitting from a diverse research and development community that can keep them a little way ahead of the commercial tools in terms of flexibility and functionality, generally at the price of a cruder interface and steep learning curve.

The popularity of building simulation has grown exponentially over the last few decades, and growing global awareness of the climate change and energy crises over that time has resulted in many pieces of legislature being passed that essentially mandate their use in modern building design. This is evident in the form of current building standards, as reviewed in section 2.2.4; it would be impossible to assess buildings against these criteria at the design stage without building simulation. As such, there are countless examples of implementations of building simulation in the literature; an exhaustive review of this is beyond the scope of this review. A review of building modelling applications in hospitals can be found in section 2.2.5.

Despite a marked stabilization within the field of building simulation, it continues to be an active research area. The shear amount of building simulation packages available stimulates competitiveness within the field, and new methods are continually being developed to refine and expand building simulation, including:

- Alternatives to notoriously computation-hungry CFD [Jin et al., 2013],
- Integration with urban micro-climate models [Yang et al., 2012],
- Improved calculation of view-factors for radiation exchange [Clarke, 2001],
- Integration with optimisation algorithms, as is the subject of the present work. See section 2.5 for a detailed literature review.

2.4 Numerical Optimisation

The field of optimisation, herein termed numerical optimisation to differentiate from other uses of the word, refers to the algorithmic process of finding the minimum value of objective functions that are functions of certain design variables, possibly subject to constraints.

2.4.1 Iterative Methods

Iterative methods for optimisation come under two broad categories; line-search and trust-region methods. In essence, line-searches determine a search direction and then calculate a step size that will minimize the error; whereas trust-region algorithms attempt to model the function and then decide based on the fit whether to expand or contract the search area.

Quasi-Newton Methods

Newton's method, as the name suggests originated from Newton's work. Its simplest form basically involves modelling the function near the optimum by a quadratic, and then using the first and second differentials to find the minimum (or maximum). The principal disadvantage of this method is that it requires the Hessian matrix to be fully defined; a square matrix of the partial second order differentials. However for a function of any appreciable complexity finding the Hessian is no small task. To get around this problem, so called "quasi-newton" methods were developed around the 60's and 70's. This began with Davidon [1959], but his method was updated and popularized by Fletcher and Powell [1963]. This method is now known as the DFP method (Davidon-Fletcher-Powell). Broyden [1965] also proposed an early variation for solving nonlinear simultaneous equations. These methods only calculate the Hessian (or Jacobian for nonlinear equations) at the first iteration, and then sequentially

update them as opposed to recalculating them at each iteration, as would be required for Newton's method.

Interior-point Methods

These originated from the work of Karmarkar [1984], who developed one of the first practical linear programming algorithms that ran in polynomial-time (requiring computations of the order of $O(n^k)$ where n is the complexity of the inputs and k is a non-negative real number). The interior-point method comprises a function which tends to infinity at the constraint boundaries. By finding a point where the gradient of this "barrier function" is equal to 0, an optimal point is found. Most modern implementations are based on the work of Mehrotra [1992], which presented a "second-order primal-dual interior point method". This was an interior point method that used a second order Taylor series to approximate a primal-dual path. The primal-dual algorithms originated from the work of Megiddo [1989], and are basically a set of optimality conditions that have certain properties symmetrically between primal and dual formulations.

Sequential Quadratic Programming

This method was first proposed by Wilson [1963] in his PhD thesis. They are a dynamic programming implementation (and hence also known as quadratic programming), and solve a series of quadratic sub-problems based on the Langrangian function (thus complying with the KKT optimality conditions) with linearized constraints to determine search directions. This does mean that second order differentials are required, however in the SQP method they are generally approximated by finite differences. This method has proved to be incredibly versatile, and has been the subject of considerable research, continued even today. Han and Powell introduced two significant improvements to Wilson's basic method in 1976 and 1977 respectively [Han, 1976; Powell, 1977]. Firstly, they suggested using a positive-definite quasi-Newton approximation to formulate the sub-problems, making them convex. Also, they introduced a merit function based on a line-search technique, which continually estimates the overall solution. Gill and Wong [2010] state in a comprehensive review of SQP (used throughout this paragraph); "It may be argued that all subsequent developments in SQP methods are based on attempts to correct perceived theoretical and practical deficiencies in the Wilson-Han-Powell approach". Such developments include the use of unconstrained sub-problems, proposed by Fletcher [1982]. Another was the identification of the "Maratos effect" (named after the work of Maratos [1978]), which may manifest if the optimisation begins to converge close to the constraint, and can disrupt overall convergence due to the steep

gradient of the merit function. Both line-search and trust-region methods have significantly influenced the development of SQP.

Other Gradient-based Methods

Numerous variations on the above techniques exist; the main categories thereof are described here. One of the largest groups are conjugate gradient methods, which are more generally methods for solving systems of linear equations. Optimisation tends to map well onto such systems, so it is no particular surprise that conjugate gradient methods have been used to solve optimisation problems. They use the known properties of A-orthogonal (or conjugate) vectors to eliminate error in conjugate directions, and are closely related to the notion of Eigen-systems. The method was first proposed by Hestenes and Stiefel [1952]. More modern variations include preconditioning matrices, which basically make very large problems more manageable and improve convergence. The base method was only applicable to linear problems, however Fletcher and Reeves [1964] generalized it to apply to nonlinear problems. The method is a specialised form of the method of conjugate directions, defined by using the residuals (these are related to how far the current iterate is from the solution) to define search directions; which is a technique used in the method of steepest descent [Shewchuk, 1994].

Simplex Methods

These began with Dantzig [1951- a & b], but Dantzig's method did not technically involve the use of simplices in the geometric sense of N+1 sided shape in an N-dimensional space. Rather, Dantzig's method used matrix operations to traverse the edges of the feasible region defined by the constraints, to the optimum vertex. The method was only applicable to a particular form of linear problems where if an optimum solution existed in the feasible region, this solution would also exist as one of the vertices of the feasible region. Presently termed "simplex methods" are quite different in that they do use simplices in the geometrical sense, but have their roots in Dantzig's method.

Currently, the most widely used incarnation is known as "Nelder-Mead's sequential simplex method", after its originators Nelder and Mead [1965]; it speaks well of the method that it is still in circulation over 45 years later. This was based on the work of Spendley et al. [1962], who presented a similar method for tracking optimum operating conditions. An important stepping stone to this research was Zoutendijk's [1960] thesis, which presented the "methods of feasible directions", but similar ideas were developed by others around the same time.

The method is easiest to visualize graphically; consider a function to be minimized that is in terms of two variables. Plotting the functions response to these variables (or phase space)

would result in a 2-D contour plot, so the problem is 2-dimensional. Imagine a triangle of arbitrary size imposed upon the plot; this is the simplex. In 3-D space, a simplex would be a tetrahedron, and past 3 dimensions it becomes impossible to visualize in real terms. The method works by reflecting one point of the simplex in the plane (or hyperplane) formed by the other points; by evaluating the function at each point, an educated decision can be made as to which point to reflect (i.e. the search direction). There are other steps in the process to make it more versatile; once a point is reflected, the simplex may expand the reflected point further from the reflecting surface if it takes it closer to the optimum. This speeds up convergence, and allows the simplex to adapt to the local search environment. On the other hand, if no improvement can be found by reflecting points then the simplex will contract to increase search resolution.

2.4.2 Global Methods

The iterative methods described above all have a common failing; when applied to complex problems, for example with a repetitive wave-like phase space, they tend to fall into local minima (i.e. the local solution closest to the starting point). Some later developments were designed with this problem in mind, for example sub-gradient methods allow for increases in the objective function, but it can never be completely eliminated due to the nature of iteration. In response to this, much of the more recent development in optimisation has been focussed on non-iterative, or global search methods (sometimes called “heuristics”, as the algorithms are generally designed to improve their search as it progresses). The price for this improved functionality is time; global search methods are invariably far more computationally intensive than iterative methods, and thus take much longer to solve.

Evolutionary Algorithms

Most global optimisation comes under this category. These are based on the fundamental tendency of nature to over time move toward the “best” or “easiest” (usually lowest energy) solution to its problems. The most well-known implementations, genetic algorithms, mimic probably the most remarkable example of this tendency; natural selection. Other examples such as swarms model more esoteric processes.

Simulation of genetic processes began in the early days of optimisation; one of the earliest examples was the work of Barricelli [1954]. Research continued throughout the later 50’s, and Bremermann [1962] suggested how genetic models may be used for optimisation at a conference on self-organizing systems in Chicago. Later, Fogel et al. [1966] published a book on applications to artificial intelligence. Coley [1999] states that genetic algorithms specifically were in essence invented by John Holland in the 1960’s; Holland published a book in 1975

detailing his approach [Holland, 1975]. The field branched out to many other areas over time, but fairly recently a new wave of developments has extended its versatility. For example, both Fogel et al. [2001] and Rudolph [2001] have published papers related to self-adapting genetic algorithms. A few years later, Cioppa et al. [2004] introduced a method of ensuring diversity in the population. Historically, practical implementations of genetic algorithms were limited only by computer processing power, and they are a very popular method for the largest and most complex problems. Currently the de-facto standard genetic algorithm is the NSGA-II algorithm, developed by Deb et al. [2002].

Swarm optimisation techniques are a more recent development; for example particle swarm optimisation was developed by Kennedy and Eberhart [Kennedy and Eberhart, 1995; Eberhart and Kennedy, 1995]. The method begins along similar lines as genetic algorithms, with a randomly generated population of points, or particles. The method of progression toward the optimum is very different though. Each particle is randomly assigned a velocity and a direction, and they are set in motion. Convergence is achieved by applying acceleration to each particle in the direction of both it's own best point found and the global best point at regular intervals. With a reasonable initial population size, it stands to reason that at least one of the particles will find the global optimum.

Another form of swarm optimisation is ant colony optimisation, defined in its current form by Dorigo and Caro [1999]. Close predecessors were released by Dorigo et al. in the previous few years [Dorigo et al, 1996; Dorigo and Gambardella, 1997]. The method is quite closely based on the behaviour of real ants, with a few useful additions. Each "ant" makes decisions based on simple probabilistic analysis of the data available at the time. Once they have formed a solution, they will lay trails to attract other ants to the trail. One of the additions to real ant behaviour is that better solutions attract more strongly, thereby encouraging global convergence. Looking at it from a programming point of view, it is essentially using lots of very simple searches in parallel, which gives it strong similarity to particle swarms.

There are other evolutionary optimisation techniques, including simulated annealing, intelligent water drops, differential evolution and cultural algorithms. Some can be used in combination with others, for instance simulated annealing can be implemented within a genetic algorithm with a decreasing mutation rate. However the methods mentioned above provide a fairly concise cross-section of the area.

Stochastic Methods

These are a slight diversion from general optimisation methods, and apply principles of probabilistic uncertainty to random variables within an optimisation problem. Basically they allow optimisation over a range of possible scenarios, so a "robust" solution may be found; that is, one which will maintain its optimality in the majority of situations. They are extremely useful for optimisation of systems involving people, as human behaviour continues to confound numerical models. Also, the same techniques can be applied to studying the phase space of an optimisation problem, which for complex problems is invaluable information for selecting an appropriate way to solve it.

Stochastic methods are generally based around random sampling and approximation of the phase space. They are quite closely linked with evolutionary techniques, as the majority of these involve random elements to stimulate convergence. They originated from Robbins and Monro [1951], and were improved shortly afterwards by Kiefer and Wolfowitz [1952]. These early predecessors were more focused on creating a solver for an unknown function than what we would now call stochastic methods; although Kiefer and Wolfowitz did apply their improved technique to finding the maxima of a function. Another important development came from Blum [1954], who offered stochastic methods for multi-dimensional problems. Later, Martin and Masreliez [1975] suggested using stochastic methods to obtain robust estimates. More techniques and applications have emerged since then; for example "simultaneous perturbation stochastic approximation", presented by Spall [1992]. This method uses simultaneous perturbation of variables inside a vector to estimate the gradient, as opposed to the Kiefer-Wolfowitz method which uses a finite difference approximation to do this.

2.4.3 Multi-objective optimisation

It is useful to briefly expand on the challenges of multi-objective optimisation, as some of the concepts encapsulated therein are central to the present work.

Where more than one objective function is required, the user is presented with a choice of how to characterise this. Essentially there are two main options; to weight the objectives and thereby combine them to a single criterion, or to characterise the optimum as a Pareto front.

Weighting the objectives can be useful for convenience, as it provides a method to reduce a multi-objective problem to a single objective problem. This has the further benefit of always providing a single optimum solution. However it has the disadvantage of presupposing the relative importance of the various objectives, which may not be valid over the entire design space. Furthermore, a change in one objective value can be offset by an opposite change in another. In problems where two objectives are typically mutually oppositional (i.e. negative

correlation) this method can often mask important features of the optima, as when one component rises, the other falls. For example largely invariant weighted objective function values can lead to a conclusion that the design variables do not actually matter, where in actual fact the optima exhibit a clear trade-off between the two objective components.

In these cases, it is often preferable (and certainly more informative) to instead consider the Pareto optimum front. This defines the optimality of a solution as the property of having no superior candidate in terms of all objective functions. So in a problem of two objective functions, if one solution is better in terms of one objective function but another is better in terms of the other objective function, both will participate in the Pareto front. However if the first solution is better than the second in terms of both objective functions, only the first will participate in the Pareto front. The effect of this is to capture the optimum trade-off between the various objectives, composed of a set of optimum solutions as opposed to a single optimum. This is illustrated graphically in Fig. 2.5; note that the data in this figure is fictional and arbitrary and is presented for demonstration purposes only.

It is important to note that once a Pareto front is developed, in order to select a single optimum solution the user must apply their own judgement. Whilst this can somewhat reduce the value of the automation of an optimisation, in many engineering scenarios it is often preferable to be presented with a set of possible solutions and have the choice, as opposed to a single solution.

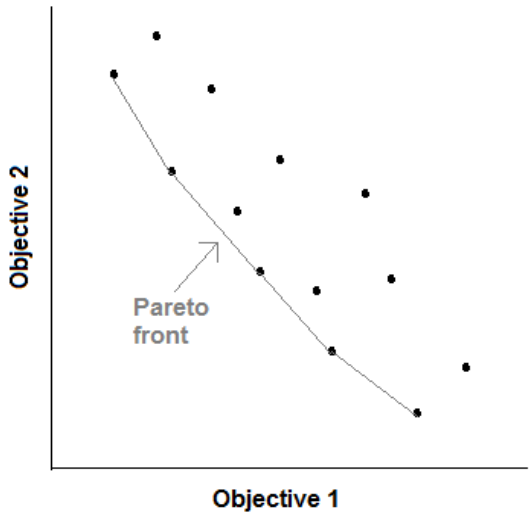

Figure 2.5: A graphical example of a Pareto front.

2.4.4 Metamodelling

As a final addition to this section, a brief review of metamodelling and its role within optimisation is presented. Metamodelling is the process of developing a simpler model to

approximate the responses of a more complex model; a model of a model in fact, hence the name metamodel. An exhaustive literature review of metamodelling techniques is rather beyond the scope of the present work, as metamodelling is essentially curve-fitting and there are many and varied methods for doing this. As such this review is mainly focussed on the emergence of the metamodelling method used in the present work, moving least squares regression; however for the sake of completeness a small section at the end is dedicated to other methods.

The principal of metamodelling can be said to have been born, along with so many other mathematical principles, with the great minds of the 18th and 19th centuries. Both Legendre and Gauss began using linear least-squares regression around 1800; it is not known for sure who exactly its true discoverer was.

From this most basic regression technique sprang many variants. The method of particular interest to this review, moving least squares regression, first appears in the literature around 1970 – 1980. The earliest reference that could be found was a paper by Lancaster and Salkauskas from [1981]. In this paper it is stated that "To date the theory of such moving least squares approximants and interpolants is meagre", suggesting that this was indeed one of the earliest works on the subject. It also suggests that the method originated from earlier work by D. Shephard, however this work could not be found in the available literature. References in this paper hint that development work was being performed by others in the area some 5-10 years previously. This paper focusses on examining a number of geometrical properties of moving least squares. From this point onwards, various applications began to spring up in the literature.

Another theoretical study was published some years later by Farwig [1986]. Following this, the first practical application found in the literature came from Otto et al. [1989], who applied the method to model chromatographic retention data. In the same year, Bos and Salkauskas [1989] published a paper linking general moving least squares regression to Backus-Gilbert theory [Backus and Gilbert, 1968].

Linkage of moving least squares with optimisation problems first began to appear in the literature in the 1990s. One of the first examples was the work of Toropov et al. [1993], which presented an optimisation procedure for structural optimisation using moving least squares to sequentially construct local metamodels from small samples. Around the turn of the millennium, the technique began to gain more wide-spread use in many diverse fields; for example Xiong et al. [2001] optimised fan components using CFD in a sequential process, calculating moving least squares metamodels and an optimum at each iteration and judging its

fitness by the discrepancy from a further simulation performed at that location. Also Ho et al. [2002] proposed a general optimisation method utilising moving least squares metamodels combined with a simulated annealing (SA) algorithm in a three step process; first an initial sample is calculated using the SA algorithm to determine sample distribution. Next, moving least squares metamodels are trained from this sample data and solved with a stochastic algorithm for the approximate optimum. Finally, a further optimisation is performed on the original objective function, using the simplex method starting from the approximate optimum. Kim and Arora [2003] considered moving least squares as a possible metamodelling technique in optimisation of force-displacement curves for tubes.

By this point moving least squares can be said to have become a well-established metamodelling technique, and the value of further review is limited. The use of moving least squares specifically in BTO is reviewed as part of section 2.5.

As noted by Gilkeson et al. [2013], approximation techniques such as moving least squares typically compare well with interpolation techniques such as kriging (the historical basis of this method is long and convoluted, the interested reader is referred to Cressie [1990]) in the context of metamodeling noisy responses. “Noisy” responses commonly refers to cases with highly non-linear equations, where the sample data may not be 100% accurate or may contain slight discontinuities, such as those typically obtained from CFD for example. Whilst interpolation techniques “force” the modelled response surface through each sample point, and attempt to develop the intermediate form of the surface by some form of interpolation, approximation techniques allow the modelled surface to diverge from the sample points. Whilst this does clearly reduce the accuracy of the modelled surface at the sampled points, it is generally preferable to have some small degree of error in the predictions than an over-fitted response surface, which can result in wildly unprecedented behaviour and false local optima. In effect, approximation methods allow smoothing of the response surface where interpolation methods do not; this is illustrated in the context of the present project in section 4.4.1.

Aside from the methods mentioned above, machine learning techniques are commonly applied to metamodelling. The main class of these are termed artificial neural networks (ANNs), as they generally simulate in a very basic manner the network topology of a brain; a network of very simple elements that when appropriately linked creates the desired model responses. These are widely accepted to originate from pioneering work published by McCulloch and Pitt in 1943. However Piccinini [2004] argues that there was already substantial research being conducted on the subject at the time, and the main novelty of McCulloch and Pitt’s work was infact the application of the notion of computational logic to

the subject, which was introduced some years earlier by Turing [1936-37]. The subject area advanced little until the introduction of the backpropagation algorithm by Werbos [1975]; this provided a means for ANNs to be effectively trained within the confines of computing capacity at the time. Despite their versatility, within a few decades newer techniques such as support vector regression (SVR) (originating from Cortes and Vapnik [1995]) had largely superseded ANNs in machine learning. However recent advances in data abstraction have sparked renewed interest in ANNs; this is typically termed "deep learning" as it pertains to efficient machine learning in many-layered (or "deep") neural networks. Generally, this revolves around formulating the neural network so it has distinct top-down and bottom-up connections between its layers; so instead of training the network in terms of recognizing data (ie. bottom-up), the network is trained to generate the data (ie. top-down). The model of the responses is then held in the weighting of the top-down connections, whilst the bottom-up connections allow the network to identify exactly what elements of each layer are important to generating the data. This bi-directional model may then be trained iteratively layer by layer [Hinton, 2007].

2.5 Building Thermal Optimisation

The topic of the present thesis, herein termed building thermal optimisation (BTO) represents the combination of the various fields reviewed in previous sections. Building design can be formulated as an optimisation problem, with performance metrics represented as objective functions and salient parameters of the building and its systems represented as design variables. This section reviews the literature available on this topic and highlights the research gaps which are addressed by the present work.

2.5.1 Previous Studies in BTO

There are three broad classes of objective functions typically considered in the literature; thermal comfort, energy use and cost. Many different models have been applied to optimise various combinations of these objectives. The general form of these models has been constant since the earliest work found in the literature [Gupta, 1970]; building simulation is applied to calculate objective functions given inputs of the design variables, and the problem can hence be solved by conventional optimisation algorithms. This has been applied in a number of different ways, generally characterised by the design variables considered.

For example, Gupta [1970] optimised thermal comfort with design variables pertaining to siting, insulation, mass, shading and glazing. This whole-building design optimisation approach

is common in the literature. Al-Homoud [1994] developed a model to optimise thermal comfort in terms of similar design variables as Gupta [1970], including additional parameters such as infiltration and operation schedule. This model was later applied to various case studies [Al-Homoud, 1997a, 1997b, 2001, 2005 & 2009], which generally demonstrated the models efficacy by predicted improvements in thermal comfort of up to around 50%. Bouchlaghem and Letherman [1990; also Bouchlaghem, 2000] explored different formulations of the thermal comfort criterion in terms of similar design variables, concluding that a time averaged deviation formulation provided robust functionality. More recent models have included additional objective functions, for example Chantrelle et al. [2011] also considered energy use and both financial and environmental cost, though with a focus on retrofitting.

A well as the form and fabric-focussed building optimisation, many other studies have been focussed on optimisation of the systems within the building, particularly HVAC. Many researchers have developed optimal control algorithms for such systems, for example Parameshwaran et al. [2010], Wang and Jin [2000] and Alcalá et al. [2005]. Notable examples of systems optimisation include the work of Djuric et al. [2007] which optimised objectives of energy use and cost, considering themal comfort as a constraint rather than an objective, with design variables pertaining to insulation and radiators. Wright et al. [2002] investigated the trade-off between energy use and thermal comfort in the form of a Pareto front, in the context of optimisation of a HVAC system. Park [2003] proposed an optimal control system for smart façade systems, and Henze et al. [2004, 2007] has studied optimal control in terms of thermal storage in both active and passive mediums.

Other studies have instead examined a specific aspect of the building design in greater detail, for example a number of studies have focussed entirely on optimising insulation such as Ozel and Pihtili [2007], Al-Sanea and Zedan [2011] and Anastaselos et al. [2011].

In general, it may be said that the increasing availability of computing power over time has allowed these models to progress to increasingly explicit representation and simulation of the building. As a case in point, Gupta [1970] used the response factor method, Bouchlaghem and Letherman [1990] used the admittance procedure, Al-Homoud [1994] used a relatively basic building simulation code called ENERCALC, and Chantrelle et al [2011] used the popular systems simulation program TRNSYS. The use of distinct building simulation programs in evaluating objective functions is currently the norm in the literature; examples include EnergyPlus [e.g. Eisenhower, 2012; Pantelic, 2012], TRNSYS [e.g. Magnier and Haghighat, 2010; Chantrelle et al., 2011], DOE-II [Bichiou and Krarti, 2011] and IDA-ICE [Hamdy, 2011]. This approach is epitomised by the GenOpt program, which provides a BTO interface that is independent of the building simulation program as detailed by Wetter [2000].

As well as dynamic thermal simulation, a few studies have used computational fluid dynamics (CFD) in BTO. Xu et al. [2008] used CFD to optimise the operation of a HVAC supply temperatures in a single room, and Stavrakakis et al. [2012] optimised window opening, both with thermal comfort as the objective. Khan et al. [2012] optimised air flows with consideration of pathogen transport and infection risk as well as thermal comfort. Also Klemm et al. [2000] used CFD to optimise the relative placements of two buildings with respect to local air flow. These approaches generally address the problem from the CFD side, introducing elements of building modelling into existing CFD codes, for example dynamic temperature boundary conditions. No BTO work was found in the literature that applies adaptive coupling procedures to independent dynamic thermal and CFD domains, such as the functionality available in ESP-r. The challenges of coupling CFD with more conventional building modelling are discussed in section 2.3, and this could be a potential reason why there are relatively few BTO studies that use coupled simulation approaches.

With the increasingly fine modelling resolution applied to BTO, the time taken with direct search optimisation techniques can become prohibitive. For example the most common optimisation algorithm applied to BTO, genetic algorithms (GAs), typically require hundreds if not thousands of objective function evaluations. Clearly then the simulation time of objective function evaluations becomes critical; assuming 1000 function evaluations, an increase in simulation time of 1 minute amounts to an increase in total run time of 1000 minutes, or 16.67 hours. When including CFD, which is particularly computationally hungry, this problem becomes particularly acute. Metamodelling is commonly applied in the literature to address this problem. For example Xu et al. [2008], Stavrakakis et al. [2012] and Khan et al. [2012] used metamodeling in order to accomplish their CFD optimisation studies. Metamodelling can also be applied to more conventional BTO, for example Eisenhower et al. [2012] proposed a generalised framework for whole-building optimisation with large numbers of design variables, performing Monte-Carlo sampling of the design space in terms of all design variables and using this data to construct metamodels. The most popular method used for metamodeling in BTO currently is neural networks (NN), for example Xu et al. [2008], Magnier and Haghighat [2010], Stavrakakis et al. [2012] and Zemella et al. [2011] all used NN metamodels. Other methods used include kriging [Gengembre et al., 2012], support vector regression (SVR) [Eisenhower et al., 2012] and moving least squares regression (MLSR) [Khan et al., 2012].

2.5.2 Summary of research gaps

As discussed in section 2.3, a common assumption made in building simulation is that of well-mixed room air; convection is a particularly difficult phenomenon to predict explicitly, and hence it is often assumed that the air in a room acts a bulk mass, absorbing and losing heat

uniformly. In practice however room air is typically not well-mixed and local temperature and air velocity gradients can be significant. The local radiant environment also varies spatially particularly due to heating and cooling devices and external weather conditions, which is also typically not explicitly taken into account. Whilst modelling studies have taken such spatial variation into account, for example the work of Fan and Ito [2012], it is typically not taken into account in BTO. Some notable exceptions do exist, for example the work of Xu [2008], but these approaches generally only consider a small number of design variables pertaining to ventilation, and are often modelled under steady conditions.

It is hereby suggested that since thermal comfort is usually of primary concern in BTO studies, these assumptions may limit the applicability of some existing models. It is therefore proposed to investigate the practicality and feasibility of a BTO methodology which does away with these assumptions, and is capable of evaluating thermal comfort locally within rooms and allows concurrent optimisation of the radiant and convective environments. This requires dynamically coupled building simulation and CFD, and hence the literature would suggest a metamodel-based method as being most appropriate. Whilst building simulation and CFD have each been applied individually to BTO in the literature, such a closely coupled procedure has not previously been applied to BTO to the author's knowledge.

ESP-r provides an ideal building simulation platform for this. As detailed in section 2.3, it is a well-established simulation program with substantial pedigree and precedent, and incorporates state-of-the-art coupling mechanisms between its various modelling domains, including CFD. The open-source model of ESP-r lends itself well to a study such as this that uses this functionality in a largely unprecedented manner, as the program can be adapted to specific requirements of the BTO architecture. Furthermore, being a natively Linux-based program it is ideally suited for command-line driven automation. Despite its apparent suitability, to the author's knowledge ESP-r has not been applied in any fully automated BTO context in the literature.

The metamodeling procedure generally creates an implicit divide between the simulation and optimisation parts of BTO; simulations are first performed to provide data to train the metamodels. Next, these metamodels are used to represent the design space and hence perform optimisation. It is further proposed that in this context, the practicalities of this disconnect can be exploited to provide flexibility in the optimisation functionality that can be extracted from a single initial sample, offsetting to some extent the large computational burden added by inclusion of CFD. To the author's knowledge, this has not been explored in the literature. Furthermore, such a method could then be used to efficiently explore how the

modelling resolution affects solutions in specific aspects, for example control periods; to the author's knowledge this has also not been considered in the literature.

BTO is often employed in a somewhat general context; application to specific environments with tailored constraints and design variables is not as common in the literature as general formulations. Where specific case studies are reported, for example by Eisenhower et al. [2012], Al-Homoud [1997a, 1997b, 2001, 2005 & 2009] and Chantrelle et al. [2011], they are usually reported as case studies demonstrating the application of a generalised BTO methodology. In the present thesis it is proposed to upend the perennial development paradigm. A BTO methodology shall be developed with a specific environment in mind, in this case hospitals. The functionality, objective functions and design variables of the BTO program shall be initially tailored to building models of hospital environments. It is postulated that this may allow more rigorous exploration of the performance of the program, as the functionality is notionally limited to a particular remit and therefore can be tested with greater focus on particular sets of conditions and design variables. Consideration of extending the applicability of the program to more diverse contexts can then be deferred to a later stage of development when data and experience relating to the performance of the program is available.

As discussed in section 1.4 hospitals present a unique design challenge, and as mentioned in section 2.2.5 are not well studied in the literature. Furthermore, the relatively high computational outlay required for BTO that includes CFD is well suited for optimisation of critical spaces. Hospitals typically cater to the most vulnerable of any population, so it is important that facilities are well equipped to provide comfortable conditions. Also in the UK in particular, the health service is currently under a great deal of pressure to maximise efficiency, so energy-efficient operation of facilities could be a critical factor.

Characterisation of thermal comfort in BTO literature is most often PMV and/or PPD, as this provides a straightforward and convenient method to bundle the most important factors affecting comfort into a single index. However results of field studies comparing objective and subjective measurements of PMV and related metrics in hospitals, such as the work of de Giuli et al. [2012] as detailed in section 2.2.5, have revealed that PMV does not always provide a good assessment of thermal comfort in hospitals. It is therefore proposed to "roll back" the characterisation of thermal comfort to a more primitive metric, in this case operative temperature (as defined in section 2.2.3). Some precendent for more primitive comfort metrics does exist in the literature, for example an objective function formulation identified as being efficient by Bouchlaghem [2000], was selected to characterise thermal comfort in terms of operative temperature. This could then allow some exploration of the impact of using

bundled metrics such as PMV instead, which to the author's knowledge has not been explored in the BTO literature.

Other potential research gaps include use of detailed thermal mannequin comfort models, comparison of a diverse range of optimisation methods, and a comparison of a diverse range of metamodeling methods. However, these aspects are not examined in great detail in the present thesis.

Chapter 3: Simulation and Optimisation Tool Methodologies

3.1 Chapter Overview

The methodology developed during the present project utilises several third party programs that were not developed by the present author. For the purposes of completeness in the reporting of the methodology of the present work, this chapter gives some details of the methodologies of these third party elements.

3.2 Building Simulation

The building simulation in the present work was accomplished using the open-source software package ESP-r (environmental systems performance, research version; ESRU, University of Strathclyde, UK). The various processes and methodologies of this program are described here. This section is split into four sub-sections; the first two describe general building simulation, and the final two describe the enhanced air-flow modelling.

3.2.1 Overview of dynamic thermal modelling

Dynamic thermal modelling (DTM) generally pertains to simulation of the building fabric and energy performance, and typically simulates in detail both conductive and radiative modes of heat exchange. Some consideration for convection is necessary, but in DTM this is most often subject to a bulk air volume assumption. Whilst this generally provides an acceptable approximation of convective heat exchange between surfaces, it does not allow evaluation of spatial variation in the convective field as was necessary for the present work. Hence the convection modelling here was conflated with computational fluid dynamics (CFD) to provide this functionality; this is described in detail in sections (3.2.3-4).

Conduction

Conduction is one of the primary means whereby energy is passively injected into a building from the environment (or rejected from the building to the environment). A time-varying heat balance (ie. of the general form of equation 3.1) exists over the building envelope, such that changes in external conditions are reflected in internal conditions but subject to a time lag, often termed phase-lag. This phase-lag is dependent on the nature of the building envelope; thicker layers of denser materials with high heat capacities take more energy to heat, and therefore take longer to transmit external fluxes to the interior, resulting in a greater phase-

lag. Remembering the approximately sinusoidal shape of diurnal temperature variations, a consequence of the phase-lag is attenuation of peak conditions; for example when external conditions reach a maximum turning point some of the energy injected into the fabric is lost back to the environment, resulting in reduction of peak internal conditions compared to external peaks. Generally the greater the phase-lag the more pronounced the attenuation effect, though this is not always a firm rule. Figure 3.1 demonstrates these principles graphically.

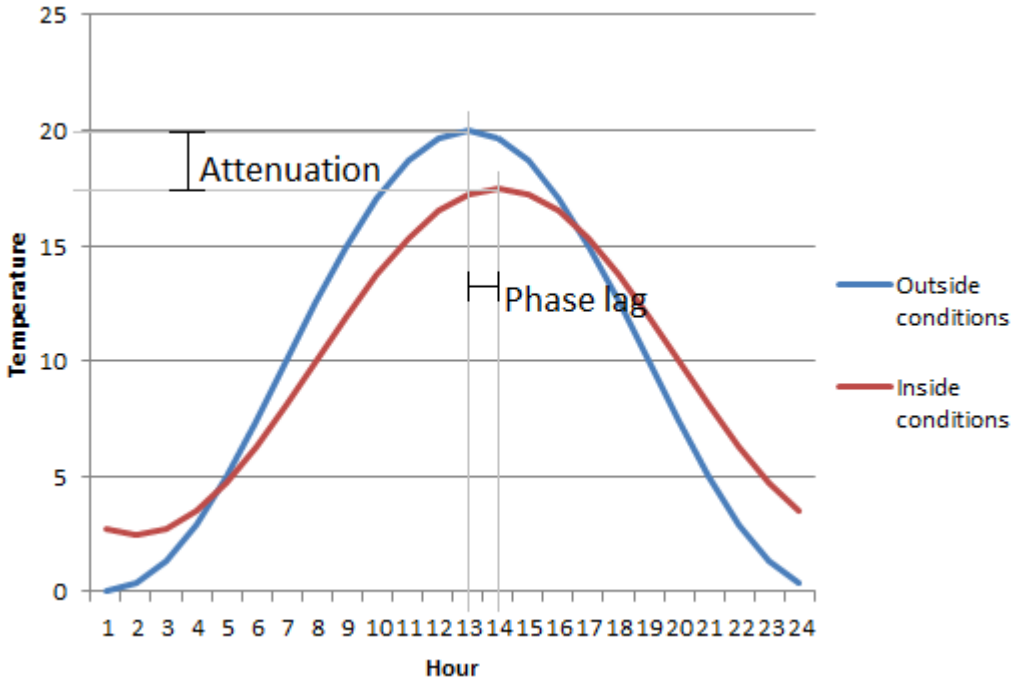

Figure 3.1: Demonstrating phase lag and attenuation in room performance with sinusoidal data (note that this data is theoretical and is shown for demonstration purposes only).

Phase-lag is the foundation of the phenomenon known as thermal mass; the property of a structure to offset external changes in order to better maintain comfortable internal conditions. As a practical example, consider an old cathedral. In the UK such monolithic structures generally have external envelopes of solid masonry, and can be anywhere up to around a meter thick; this results in very high thermal mass. If one were to touch the internal surface of the building envelope on a warm morning, it will invariably feel very cold. This is a result of the thermal mass of the building offsetting the cool conditions of the previous night into the day-time, which on a hot day can go a long way to maintaining comfortable conditions within the building.

As well as conduction through the external envelope, in some situations conduction within the building can also have significant effects on thermal conditions. For example, poorly insulated plant or boiler rooms can leech heat into adjacent spaces. However in most cases internal

conductive heat fluxes are negligible, which allows consideration of internal surfaces as adiabatic boundaries. Similarly, whilst direct conduction from the building fabric to occupants can affect thermal comfort (eg. excessively cold floors), under normal conditions it does not have any appreciable effect. Therefore this is normally not considered in evaluating thermal comfort. Nevertheless it is still necessary to simulate conduction in the building fabric accurately as it largely controls internal surface temperatures, which do affect thermal comfort through radiation as explained below. Fortunately, as conduction is intrinsically reliant on a heat balance it plays to the strengths of modern building simulation software.

In the context of building thermal optimisation (BTO), optimisation of the building fabric has the potential to improve passive thermal performance of the building, and was therefore considered worth investigating. When developing the models used in the present project, steps were taken to ensure that material properties of the external walls could be controlled by design variables. See section 4.3 for a detailed discussion of building model development.

Radiation

Radiation is the instantaneous exchange of energy between two surfaces in view of one another. It is occurring constantly all around us wherever we are, not only between other surfaces but also between those surfaces and our bodies. As a result, thermal radiation is one of the elements of our environment which directly affect thermal comfort. In layman's terms, thermal radiation operates under a similar principal to diffusion; areas of high concentration (ie. warmer surfaces) will radiate to areas of low concentration (ie. colder surfaces) until equilibrium is reached and both surfaces are the same temperature. In the context of human thermal comfort this is impractical to model explicitly, as it is would be unnecessarily complex to accurately simulate the surface temperature of all of an occupants clothing and skin. The radiative field is therefore characterised by radiant temperature; a mean of all surface temperatures in view of the occupant weighted by the proportion of the total view field occupied by each surface, known as view factors. This radiant temperature is one of the two primary inputs to operative temperature, the other being air temperature.

There are two principal ways to calculate view factors; angle factors and ray tracing. Angle factors are the more traditional method, being relatively easy to calculate by hand, whereas ray tracing involves much more intensive calculation but tends to result in greater accuracy. Angle factors utilise solid angles to determine the proportion of the view field occupied by each surface. As a theoretical example, consider an occupant suspended in the centre of a sphere. If one half of the sphere is black, and one half is white, the two halves can each be said to have a view factor of 0.5, as each occupies exactly half of the 3-D view field of the

occupant. The radiant temperature experienced by the occupant would therefore be the average of the surface temperatures of the two halves of the sphere, with equal weighting between the two

On the other hand, ray tracing derives exact exchange areas by Monte-Carlo sampling of notional rays from the surface or object in question. Notional lines are drawn from some sampling point, spreading in all directions and continuing until they impact a solid surface. By recording which surface each ray impacts on, and knowing the relative angle of each ray, radiant exchange areas can then be calculated; this process is repeated for many sample points per surface, for each surface in the room. Simple calculations repeated many times over clearly fall nicely into the remit of computers, so ray tracing is often the method of choice in building simulation software, however trigonometry is also well established in modern programming languages so angle factors are sometimes offered as a simpler alternative as is the case with ESP-r. For the present project ray tracing was used to determine view factors to provide greater accuracy.

As well as taking into account exchange between surfaces and occupants, radiation modelling also has another important role to play in building simulation. Solar radiation is a major contributor to passive energy gain in buildings by impinging upon the external envelope. As well as this, all buildings aside from a few specialised examples generally have windows, and glazing has the unique property of allowing light and hence radiation to pass to the interior of the building. As a result solar radiation can impinge on internal surfaces and heat them directly. A portion of impinging solar radiation is also reflected from surfaces, resulting in a diffuse field of solar radiation as well. Whilst this is usually of a considerably lower intensity than direct solar radiation, it acts in all directions and can still have significant effects on building thermal performance.

With radiation, the first of the mechanisms arises by which thermal comfort may vary spatially. Depending on the position of the occupant within a room the view factors for each surface will vary, and hence so will radiant temperature. As a practical example, consider a room with one external south-facing window and a radiant panel on the opposite wall. It is a warm summer's day with a clear sky and little wind. Clearly in conditions such as this the external environment is very warm, and internal spaces will generally require cooling; as such the radiant panel is set to a low supply temperature. Assuming the building does not have a large amount of thermal mass, we can assume that the internal surface temperature of the external wall and the window will both be significantly higher than internal walls due to the intensity of the solar radiation and the warm outside air. Someone stood close to the external wall and window will therefore perceive a greater radiant temperature as the warmer surfaces occupy and greater

proportion of their view field, as well as direct solar radiation that may be striking them through the window. Conversely, an occupant stood next to the cold radiant panel will experience a significantly lower radiant temperature as their view field is not so dominated by the warm surfaces, and has a greater influence from the radiant panel. From this we can surmise that the optimum radiant conditions will likely vary depending on where in the room thermal comfort is evaluated.

3.2.2 DTM calculation methodology

The theory behind ESP-r is well documented by Clarke [1977, 2001], and also by a wealth of PhD theses describing specific investigations and extensions of the system [Hensen, 1991; Aasem, 1993; Nakhi, 1995; Negrão, 1995; MacQueen, 1997; Kelly, 1998; Hand, 1999; Beausoleil-Morrison, 2000; Amissah, 2005; Samuel, 2006]. A brief overview of the calculation methodology behind it is presented here; for more detail the reader is referred to the relevant literature.

ESP-r approximates the energy flow in buildings by forming control volume heat balance equations at a system of nodes; a discretized representation that retains the thermodynamic contacts of the building topology. The resulting system of equations is then solved simultaneously at each time-step, providing the initial conditions for the next time-step.

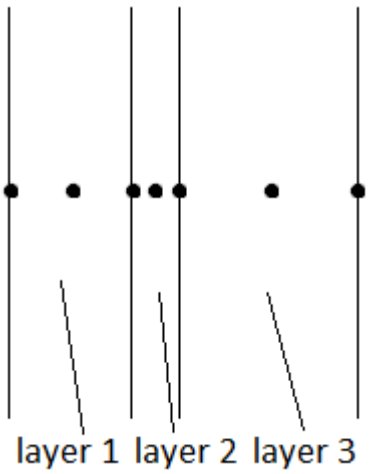

Figure 3.2: Showing positions of calculation nodes in a 3 layer construction, using the default convention of ESP-r.

By default, ESP-r represents each homogenous building layer (i.e. each different material within a wall or a floor) by 3 nodes; an internal node within the material, and two surface nodes at both edges of the material. The surface nodes are shared with neighbouring materials in the case of multi-layered constructions. Figure 3.2 illustrates this graphically. This

one-dimensional approach to energy transmission through the building fabric is the de-facto standard within the field of building simulation, and was considered acceptable for the present project. The number of nodes was informed by a parametric study conducted in the early stages of the development of ESP-r [Clarke, 2001], which found that 3 or more nodes per homogenous layer generally resulted in acceptable accuracy in the outputs. By default air volumes are represented by a single node, however this was not considered acceptable for this project, so a CFD domain was included in the model.

The heat balance equations are derived from the following general form, which assumes that thermal properties are time-dependent:

$$\frac{\rho_i(t)C_i(t)V_i}{\delta t}[T_i(t_c+\delta t)-T_i(t_c)] = \sum_{I=1}^{N} K_{i,I}(t)[T_I(t)-T_i(t)] + q_i(t) + \varepsilon \quad (3.1)$$

where $\rho_i(t)$ is the representative density at node i and time t (kg / m^3), $C_i(t)$ the representative specific heat capacity (J / kg °C), V_i the volume of the region represented by node i (m^3), δt the time step (s), $T_i(t)$ the representative temperature (°C), and $K_{i,I}(t)$ the heat flow conductance between nodes i and I (W / °C). Any heat generated inside the region is represented by $q_i(t)$, ε is the error resulting from the spatial and temporal discretization, N the number of energy flow-paths between node i and neighbouring nodes (represented by I), and t_c the current time. This equation effectively states that at any node i, the total change in temperature is proportional to the net energy flux over the period of the time step, the ratio being determined by the material properties.

ESP-r uses a weighted average of the explicit formulation (where $t=t_c$) and implicit formulation (where $t=t_c + \delta t$); by default equal weighting is used which results in a Crank-Nicolson formulation. In this way the best features of both formulations are retained, resulting in a numerically stable yet accurate system of equations. This system of equations is often augmented by the inclusion of further modelling domains, such as plant-driven thermal control systems (eg. radiators, HVAC); generally the augmentation process is simply a matter of expanding the matrix with further elements or replacing specific elements with more complex individual models.

By grouping the terms, it is possible to represent the system of equations by a matrix equation of the form:

$$\mathbf{A}\,\mathbf{T}_{n+1} = \mathbf{B}\,\mathbf{T}_n + \mathbf{C} \quad (3.2)$$

where matrix **A** contains coefficients pertaining to the future time step, **B** contains coefficients pertaining to the current time step, column matrix **C** contains known boundary interactions

such as heat exchange with the sky and ground, and column matrices $\mathbf{T}_{n+1}$ and $\mathbf{T}_n$ contain nodal temperatures and heat injections for the next and current time-steps respectively. Since **B**, $\mathbf{T}_n$ and **C** are known at the current time-step, it is appropriate to define a column matrix such that:

$$\mathbf{Z} = \mathbf{B}\,\mathbf{T}_n + \mathbf{C} \tag{3.3}$$

Substituting this into equation 3.2, the solution then becomes:

$$\mathbf{T}_{n+1} = \mathbf{A}^{-1}\,\mathbf{Z} \tag{3.4}$$

This equation is solved by computationally inverting the **A** matrix, which is a complex process and is performed by specialised matrix solver algorithms at the heart of ESP-r. A full and exhaustive explanation is given in chapter 4 of Clarke [2001].

The core purpose of the simulation is to extrapolate the boundary conditions to the rest of the model, and hence predict the performance of the modelled area. The boundary conditions for the problem are mostly provided by the external climate; in order to run an ESP-r simulation known climate data must be specified. This, along with specified interactions with other areas of the building, provides the external excitations that stimulate energy flow within the model; for example conduction through an external wall, sunlight entering through a window, and air flowing through an open door. Active thermal control systems in the model (eg. radiators and HVAC) provide further internal excitations, often aimed at counteracting the external influences and maintaining the thermal status quo in the modelled zone. However it is worth pointing out that depending on the control of the systems present in the room, this can be dependent on the conditions within the room at the time, so these cannot necessarily be considered as independent boundary conditions.

By necessity this method is discretized over time into distinct time steps, under the assumption that conditions remain constant throughout each time step. These can be any length of time, from fractions of a second to hours or even days. The size of the time steps plays a significant role in determining where the simulation sits in the trade-off between computation time and model accuracy. For example, a time step of 1 second would result in a fairly true to life simulation, as conditions will very rarely vary significantly within this time period. However, a simulation of 24 hours would then be split into 86,400 time steps, each one requiring a separate run through the process described above. Clearly, this will take a lot longer to simulate than if the time step was 1 hour, there being only 24 time steps to evaluate in this case. On the other hand, this larger time step may reduce the accuracy of the simulation as conditions may well vary within the period, and this variation would not be

encapsulated in the performance predictions given by the model. However by using effective values of time-dependant material properties and including an error term in equation 3.1, the method does seek to average out these sub-time step variations effectively, such that the performance predictions provide an acceptable approximation of results that would be obtained with shorter time steps. For the present project computation time was limited by the resources available; see section 4.3.2 for a discussion of this.

3.2.3 Overview of computational fluid dynamics

As mentioned previously, convection modelling in DTM is typically subject to a bulk air volume assumption. This models the air in the room as a bulk unit, implying that any heat transfer to and from the air is averaged out amongst the entire volume and ignoring any spatial variation in temperature and velocity. These variations will affect thermal comfort; an occupant standing in a patch of warm stagnant air, perhaps from a poorly designed HVAC supply vent, will generally be warmer than one standing in a patch of fast flowing cold air, perhaps from an open window. To get around this for the present project, the traditional "bulk" convection modelling of DTM was augmented with CFD simulations.

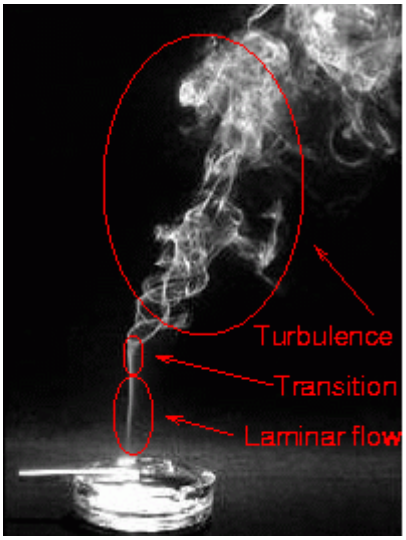

Figure 3.3: Photograph showing the different states of air flow.

CFD is essentially the iterative solution of the Navier-Stokes equations (given in section 3.2.4). It seeks to explicitly simulate convection, evaluating the flow patterns of air given a set of driving boundary conditions, and hence allows evaluation of spatial variation of air properties. Convection is typically the most esoteric of heat flow paths within buildings; this is largely due to the fundamentally chaotic and unpredictable way in which air usually moves, known as turbulent flow. Introducing smoke to air is an excellent way to visualise turbulence; figure 3.3 shows a smoke source (a cigarette) and the development from laminar to turbulent flow. It

can be seen that once in turbulent flow the air moves in a very unpredictable manner, hence the difficulty in accurately simulating turbulence. In DTM there is no need to take turbulence into account due to the bulk air volume assumption, however in CFD it is a critical part of the process, as detailed in section 3.2.4.

As all boundary conditions must be explicitly specified for any CFD simulation, they are typically subject to a high degree of assumption when CFD is used in isolation in the context of building modelling. In ESP-r, the need to give explicit values for CFD boundary conditions is largely eliminated when it is conflated with the DTM, as the boundary conditions are determined dynamically by the DTM. The DTM-CFD conflation within ESP-r was initially developed by Negrão [1995] and later improved by Beausoleil-Morrison [2000] who implemented the adaptive conflation controller (ACC); for an in-depth description of the mechanisms the reader is referred to these theses. They automate the process of passing boundary conditions from the DTM to the CFD on a time-step basis, and if applicable passing resultant heat fluxes back to the DTM. By default the ACC is active; this is a fully automated system that invokes the CFD, and controls the coupling procedure at each time step. The ACC has some decision making capability based on prevailing flow patterns evaluated by a preliminary CFD assessment.

One particular disadvantage of CFD is that the iteration is not guaranteed to converge to a feasible solution, and typically requires supervision and user intervention to be most efficient. In the present work however the process needed to be fully automated simply for practicality; this presented a significant problem in two respects:

1. The CFD simulations could not be monitored and coaxed to convergence as would normally be the case. However the conflation mitigated this problem considerably, as the CFD simulation worked in tandem with more stable simulation domains. Also, it was found that the procedure of the adaptive conflation algorithms went some way to compensating for the lack of user intervention; before the actual simulation is run a separate preliminary simulation is undertaken which assesses the general flow patterns and adapts the actual simulation to the conditions at that time-step, increasing the chances of successful convergence. Whilst this had the disadvantage of substantially increasing computation time, it was considered worthwhile to give a better chance of convergence and hence greater confidence in the results.
2. Because of the lack of user monitoring, the very human process of judging results to be “reasonable” or not was lost, and hence error was rather more likely to creep in. If erroneous results were relayed back to the DTM, it could propagate error throughout the simulation. To prevent this, handshaking (as it is termed within ESP-r) between

the DTM and CFD domains was limited to one way; boundary conditions were passed from the DTM to the CFD, but information from the CFD was not allowed to be relayed back to the DTM. Whilst this marginally reduced the value of the conflation, it was considered more important to eliminate the possibility of an error feedback loop developing within the handshaking procedure.

3.2.4 CFD calculation methodology

The foundation of CFD is the principal that any fluid (note that a fluid here is any non-solid matter, be it liquid or gas) must obey the laws of physics; particularly conservation of mass, conservation of momentum and conservation of energy. In general, any change in motion or state of a fluid has precedent in the surrounding fluid and/or the boundary conditions. To put it another way, a static fluid will only move or heat up or otherwise change in any way as a result of energy flux. It can be said that that the energy flux necessary to affect any change in a particular region of fluid will come from adjacent regions of fluid. This holds until we reach the edge of the fluid, whereupon energy flux may come from any surfaces the fluid is in contact with; these are the boundary conditions of the problem.

It is beyond the scope of this work to fully report the workings of CFD, the following are the fundamental equations by which CFD is accomplished. For more in depth information and derivations the reader is referred to the wealth of literature on the subject, for example Versteeg and Malalasekera [1995]. All of the following equations assume a standard Cartesian coordinate system (x and y horizontal, z vertical) and a Newtonian, incompressible fluid. Also, as CFD problems are typically very non-linear and impossible to solve analytically, the CFD domain is split up into a certain number of cells, each representing a region of fluid with notionally constant properties.

Conservation of mass

Otherwise known as the continuity equation, conservation of mass when applied to a fluid cell states that any change of mass within the cell is equal to the net flow rate of mass into- and out of- the cell. The assumption of an incompressible fluid simplifies this principal, as it necessitates constant density and therefore constant mass within the cell. Conservation of mass for incompressible flow can therefore be expressed as:

$$\frac{\delta}{\delta x}\rho u_x + \frac{\delta}{\delta y}\rho u_y + \frac{\delta}{\delta z}\rho u_z = 0 \qquad (3.5)$$

Where ρ is the density (kg/m^3), u_x is the vector component of total velocity in and out of the cell in the x direction (m/s), u_y is the vector component in the y direction (m/s), and u_z is the vector component in the z direction (m/s). [Beausoleil-Morrison, 2000]

Conservation of momentum

Often called the Navier-Stokes equations, these arise from Newton's second law applied to the fluid cell; the rate of change in momentum of the fluid cell is equal to the sum of the forces acting upon it. Expressing this in terms of vector components:

$$\frac{\delta}{\delta t}\rho u_x + \frac{\delta}{\delta x}\rho u_x u_x + \frac{\delta}{\delta y}\rho u_y u_x + \frac{\delta}{\delta z}\rho u_z u_x = -\frac{\delta P}{\delta x} + \frac{\delta}{\delta x}\left[\mu\left(\frac{\delta u_x}{\delta x} + \frac{\delta u_x}{\delta x}\right)\right] + \frac{\delta}{\delta y}\left[\mu\left(\frac{\delta u_x}{\delta y} + \frac{\delta u_y}{\delta x}\right)\right] + \frac{\delta}{\delta z}\left[\mu\left(\frac{\delta u_x}{\delta z} + \frac{\delta u_z}{\delta x}\right)\right] \quad (3.6)$$

$$\frac{\delta}{\delta t}\rho u_y + \frac{\delta}{\delta x}\rho u_x u_y + \frac{\delta}{\delta y}\rho u_y u_y + \frac{\delta}{\delta z}\rho u_z u_y = -\frac{\delta P}{\delta y} + \frac{\delta}{\delta x}\left[\mu\left(\frac{\delta u_y}{\delta x} + \frac{\delta u_x}{\delta y}\right)\right] + \frac{\delta}{\delta y}\left[\mu\left(\frac{\delta u_y}{\delta y} + \frac{\delta u_y}{\delta y}\right)\right] + \frac{\delta}{\delta z}\left[\mu\left(\frac{\delta u_y}{\delta z} + \frac{\delta u_z}{\delta y}\right)\right] \quad (3.7)$$

$$\frac{\delta}{\delta t}\rho u_z + \frac{\delta}{\delta x}\rho u_x u_z + \frac{\delta}{\delta y}\rho u_y u_z + \frac{\delta}{\delta z}\rho u_z u_z = -\frac{\delta P}{\delta z} + \frac{\delta}{\delta x}\left[\mu\left(\frac{\delta u_z}{\delta x} + \frac{\delta u_x}{\delta z}\right)\right] + \frac{\delta}{\delta y}\left[\mu\left(\frac{\delta u_z}{\delta y} + \frac{\delta u_y}{\delta z}\right)\right] + \frac{\delta}{\delta z}\left[\mu\left(\frac{\delta u_z}{\delta z} + \frac{\delta u_z}{\delta z}\right)\right] + S_M \quad (3.8)$$

Where P is the pressure (Pa), μ is the viscosity (Pa s) and S_M is the gravitational source term.

These three equations are all of similar form, however each equation describes conservation of momentum in a particular direction; equation 3.6 deals with momentum in the x direction, equation 3.7 in the y direction, and equation 3.8 in the z direction. The first term represents the rate of change of directional momentum within the control volume with respect to time. The remaining three terms on the left-hand side represent net out-flows of momentum from the control volume due to flows crossing the faces; this is often called convection but is technically termed advection. Moving to the right-hand side, the first term represents the pressure force component. The next three terms represent viscous forces in the relevant direction; these are essentially characterisations of diffusion across the control volume faces.

One additional term is added to equation 3.8, a source term representing buoyancy and gravitational forces that uniquely affect the z direction. [Beausoleil-Morrison, 2000]

Conservation of energy

Equations 3.6 – 3.8 describe the motion of the fluid, but it is also necessary to define the thermal field in order to simulate buoyancy, which can be a significant driver of air movement:

$$\frac{\delta}{\delta t}\rho C\theta + \frac{\delta}{\delta x}\rho Cu\theta + \frac{\delta}{\delta y}\rho Cv\theta + \frac{\delta}{\delta z}\rho Cw\theta = \frac{\delta}{\delta x}\left(\lambda\frac{\delta\theta}{\delta x}\right) + \frac{\delta}{\delta y}\left(\lambda\frac{\delta\theta}{\delta y}\right) + \frac{\delta}{\delta z}\left(\lambda\frac{\delta\theta}{\delta z}\right) + S_E \quad (3.9)$$

Where C is the specific heat capacity (J/kg K), λ the thermal conductivity (W/m K), θ the temperature (°C) and S_E is the heat generation source term. The terms in this equation are broadly similar to those in the momentum equations; on the left-hand side are rate of accumulation and convective terms, and on the right diffusion and source terms. [Beausoleil-Morrison, 2000]

With equations 3.5 – 3.9 we have a closed problem, with 5 equations and 5 unknowns. However due to the occurrence of turbulence, most flow patterns experience chaotic fluctuations that are typically very difficult to quantify; these are caused by instabilities between inertial and viscous forces [Beausoleil-Morrison, 2000]. To account for this, a separate turbulence model is generally applied to the problem. Turbulence modelling is an ongoing field of research, and is practically its own distinct subject area. It is beyond the scope of this thesis to examine the literature in detail, so this explanation focusses on the turbulence models available in ESP-r.

The k-ε model

This is by far the most popular turbulence model to date [Beausoleil-Morrison, 2000], and is the de facto standard model in most applications of CFD. The so-called "standard k-ε" model is the most applied accepted form, and was proposed by Launder and Spalding [1974].

An important concept in turbulence modelling is the eddy viscosity "μ_t"; this must be introduced before a clear explanation of the k-ε model can be given. It seeks to draw an analogy between turbulent diffusion and molecular diffusion, thereby providing a clear framework to model turbulence. First, the turbulent stresses are assumed to be proportional to the mean velocity gradients; this essentially represents the turbulence as an effective

viscosity on the time-averaged flow [Beausoleil-Morrison, 2000]. However, whereas traditional viscosity is a material property, the analogous turbulent quantity μ_t is instead a property of the flow, and can thus vary within the CFD domain. Similar assumptions are made for the thermal field. The practical upshot of this, once integrated into equations 3.5 – 3.9, is that the momentum and heat diffusion coefficients become variant depending on location within the domain, as opposed to being constant; momentum diffusion coefficients become $\mu - \mu_t$ instead of μ and heat diffusion coefficients become $\lambda + \frac{C_p \mu_t}{\sigma_t}$ instead of λ, where σ_t is the turbulent Prandtl number representing the ratio of turbulent and thermal viscosities [Beausoleil-Morrison, 2000].

The k-ε model is so named for its two principal variables; the turbulence kinetic energy "k" and the dissipation rate of turbulence energy "ε" [Beausoleil-Morrison, 2000]. These are used to define and evaluate the μ_t distribution in the domain, thereby allowing evaluation of equations 3.5 – 3.9. Like the velocity components u_x, u_y and u_z in equations 3.6 – 3.8, and the temperature θ in equation 3.9, both k and ε have transport equations similar in form to equations 3.6 – 3.9. Since μ_t is defined only in terms of k, ε, known material properties and empirical constants, this then results once more in a closed system of equations; 7 equations with 7 unknowns. It is worth noting however that it is only by introducing a number of empirical quantities into the equations for μ_t, k and ε that a closed system is obtained; a purely analytical formulation and application of the k-ε model would not result in a closed system [Beausoleil-Morrison, 2000].

Zero equation models

The primary distinction of these models from the k-ε model is that instead of calculating μ_t using k and ε, a fixed value is assumed or it is instead related to mean velocity distribution [Beausoleil-Morrison, 2000]. This significantly simplifies the turbulence modelling.

Two distinct zero equation models are provided in ESP-r; one with a fixed value of μ_t and one proposed by Chen and Xu [1998] where μ_t is related to local mean velocity as in equation 3.10:

$$\mu_t = 0.03874 \rho \bar{V} l \tag{3.10}$$

Where $\bar{V}$ is the local mean velocity and l is the distance to the nearest solid surface. It is reported by Beausoleil-Morrison [2000] that both Chen and Xu [1998] and Srebric et al. [1999] validated this model against experimental data and found good agreement with the predicted results; infact in some cases the zero equation model out-performed the k-ε model.

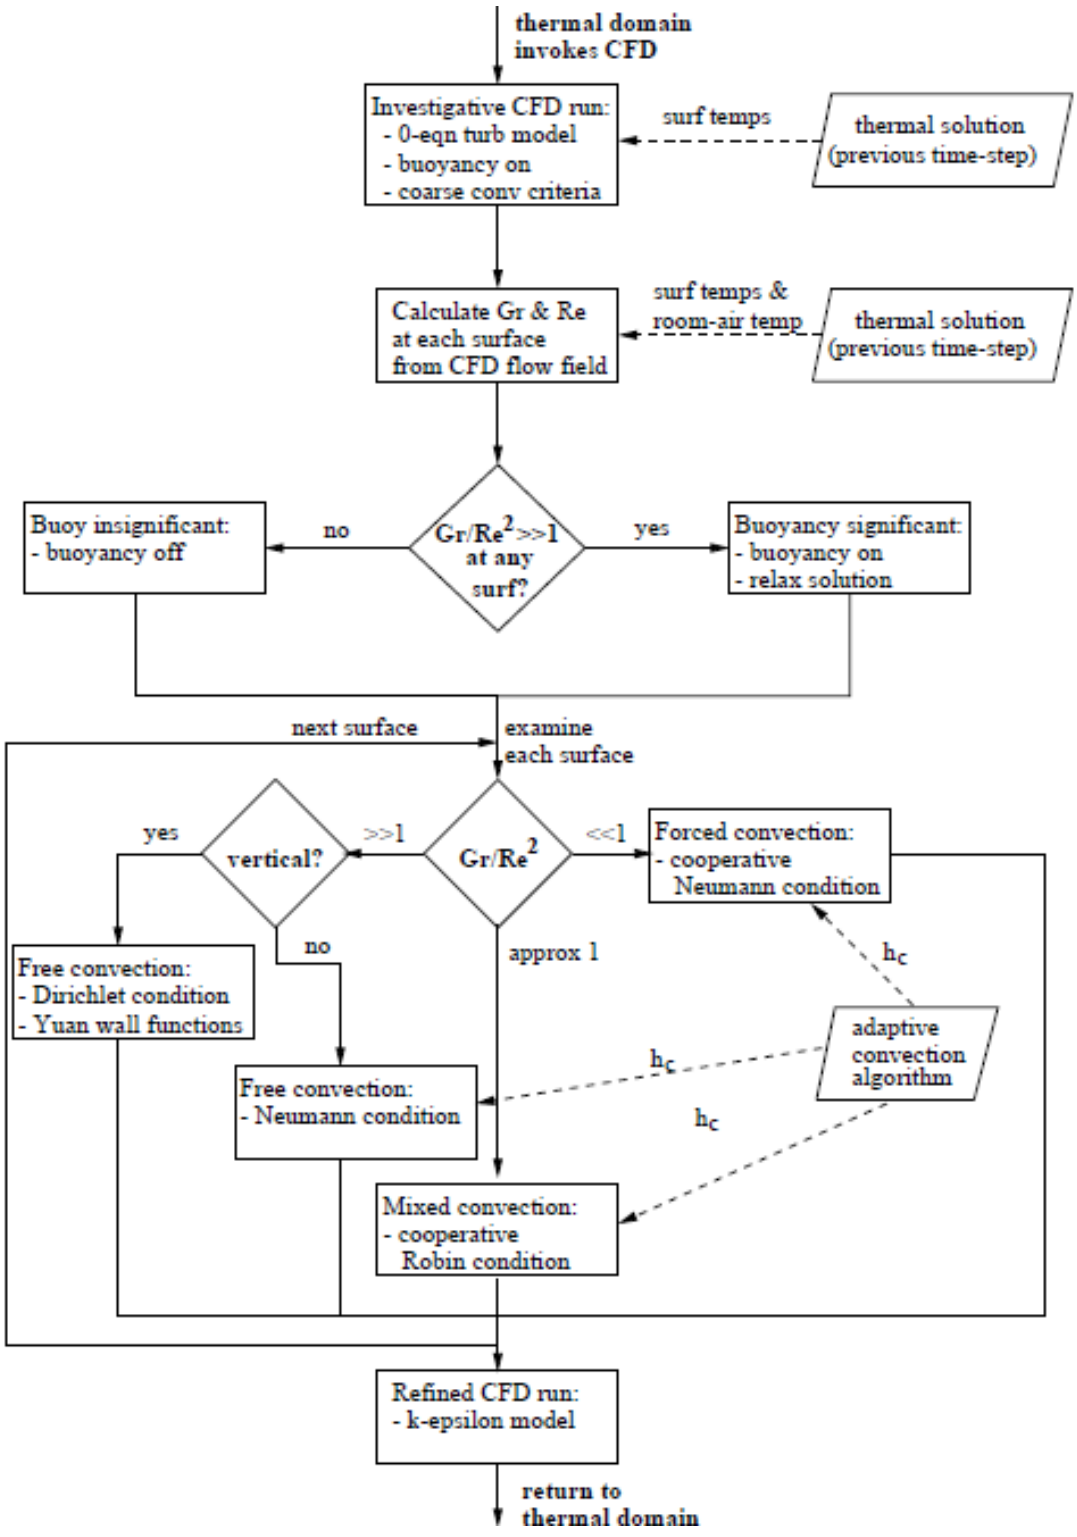

Figure 3.4: One-way adaptive conflation algorithm for solid boundary conditions.

Details of CFD solver in ESP-r

The CFD in ESP-r uses a finite volume approach. The solver algorithm is the tridiagonal matrix algorithm (TDMA) a.k.a. the Thomas algorithm; essentially this employs Gaussian elimination to solve the discretised matrix equation in sweeps, plane by plane. The SIMPLEC algorithm is employed for pressure-velocity coupling. A hybrid differencing scheme is used for spatial

discretisation, applying either first-order or second-order accurate methods depending on Peclet number. A staggered grid is applied to the velocity field to prevent potential discretisation errors arising from a "checker-board" pressure field. All of the above are well detailed by a wealth of literature; the author recommends Versteeg and Malalasekera [1995] for a comprehensive and accessible discourse.

Boundary conditions of the coupled CFD domain are assigned to surfaces or air flow openings in the building model, which are updated at each time step where appropriate. In the present project, inlet and outlet grills, window openings and doors fall under the category of air flow openings. These boundary conditions are implemented as a fixed velocity. Where the CFD is conflated with a mass flow network, the velocity is set dynamically by iteration between the CFD and mass flow network at each time step. Values of velocity can be negative (air leaving the domain) or positive (air entering the domain), and must satisfy conservation of mass as detailed above. All other boundary conditions fall under the category of solid surfaces. The ACC decides at each time step how to implement the boundary condition based on prevailing flow conditions. Heat flux can be negative (heat flow from the room air into the surface) or positive (heat flow from the surface into the room air) and must satisfy conservation of energy as detailed above. A flow-chart representation of the algorithm for one-way conflation is reproduced from Beausoleil-Morrison [2000] in Fig. 3.4. The boundary conditions mentioned in Fig. 3.4 are as follows:

- Dirichlet: Surface temperature prescribed by DTM domain, CFD resolves the boundary layer.
- Neumann: Heat flux prescribed by the DTM domain, enters the CFD domain in the source term.
- Cooperative Neumann: As the standard Neumann condition above, except mass-averaged air temperature from the CFD is used to calculate heat flux.
- Cooperative Robin: Local air temperature and heat transfer coefficient used to calculate heat flux, enters the CFD domain in self-coupling and source terms.

For more detail the reader is referred to chapter 5 of Beausoleil-Morrison [2000]. Boundary conditions for each model used in the present work are listed in tables X.X and X.X.

3.3 The Optimisation Technique – Genetic Algorithms

The following is a procedural description of genetic algorithms. This is given in lieu of a formulaic description as a mathematical representation of genetic algorithms is rather

complex and far less intuitive than the procedural representation. Little, if any, extra pertinent information is given by the formulaic description, however the interested reader is referred to Chapter 4 of Vose [1999] for such a discourse.

There are numerous variations on genetic algorithms, however the core procedure remains largely unaltered due to the very definition of genetic algorithms. The following are the basic steps involved [Vose, 1999]:

- Rank and pair up parent solutions.
- Perform mixing to create child solutions.
- Put parent(s) and/or child(ren) into next generation as desired.
- Repeat until next generation is full.

Vose [1999] defines a genetic algorithm, in its simplest form, as a special case of random heuristic search. Random heuristic search is a progression according to a transition rule τ, from an initial population P_0 of r samples taken from a search space Ω, to create a sequence of populations [Vose, 1999]:

$$P_0 \overset{\tau}{\rightarrow} P_1 \overset{\tau}{\rightarrow} P_2 \overset{\tau}{\rightarrow} \ldots \tag{3.10}$$

The transition rule τ is comprised of two elements; a heuristic function G that calculates a new sample distribution from that of the current population, and then r independent samples taken within the search space Ω according to the new sample distribution. These samples form the next population.

The definition of a genetic algorithm comes from specific definitions of Ω and G. Firstly, Ω is defined as a space of encoded representations of possible values of P. The simplest form of this encoding is binary representation of integers; this is used as an example throughout this section to demonstrate the procedure. However in practice, where the parameters of Ω are often continuous, more complex and efficient systems such as range encoding are often used instead [Coley, 1999]. As an example, consider the search space Ω={0,1,2,3,4,5,6,7}. This range of integers can be encoded into unique binary strings, each composed of 3 bits as shown in table 3.1 [Vose, 1999]. The bits may be thought of as analogous to genes.

Value	0	1	2	3	4	5	6	7
Binary string	000	001	010	011	100	101	110	111

Table 3.1: A simple case of binary encoding of integers.

To define G, inspiration is taken from the process of evolution by natural selection; hence the name genetic algorithm. The process is broken down into 3 distinct steps which each mirror a

specific element of natural evolution. First comes selection, whereby the population is ranked according to their fitness. Just as in the natural world, fitter individuals have a greater chance of influencing the next generation, whereas weaker ones have a greater chance of dying off without breeding. Next, in no particular order, comes mutation and crossover, collectively known as mixing [Vose, 1999].

At this point we can begin to relate GAs to the formalised optimisation definition given in Chapter 1; the functions that determine fitness are in fact the objective functions, and the dimensionality of the search space Ω represents the number of design variables. It is worth noting for clarity that in the integer example above Ω has dimensionality of only 1; the number of possible solutions (in this case 8) is the cardinality, which is a separate parameter entirely. A similar example with dimensionality 2 would have a search space Ω={ (0,0),(0,1),(0,2),…,(1,0),(1,1),(1,2),…,(7,7) } which would have cardinality 8^2.

Mutation functions exactly as the name suggests; random solutions in each generation are changed in some small way. In the binary integer example above, this would generally be accomplished by "bit-flipping"; selecting a single bit at random and flipping its value [Vose, 1999]. Mutation allows for evolution of subsequent populations outside the confines of the initial population. In the example, if the initial population by some quirk of fate did not include any individuals with the value "0", this could result in the corresponding region of the search space being excluded from the actual search. Mutation ensures that, regardless of the crossover method employed, the possibility exists for a "0" value to randomly appear with no precedent from previous populations. Crossover is analogous to sexual reproduction in the natural world; the process of combining two parents' genes to create children. Numerous methods exist for crossover, however for the simple example used here only two classical methods are described; 1-point crossover and uniform crossover.

Parents	Split at crossover point	Swap trailing bits	Children
011	0 / 11	0 / 01	001
101	1 / 01	1 / 11	111

Table 3.2: Illustrating 1-point crossover for the 3-bit integer example.

1-point crossover involves randomly choosing a point with uniform probability in the binary string and swapping all bits after this point [Vose, 1999]. For example, say two parents a=3≡011 and b=5≡101 have been paired up for crossover. Neglecting points at the edges of the strings (which result in the children being identical to the parents, known as cloning), two possible locations exist for the crossover point; between the first and second bits or between

the second and third bits. Let us assume that the first point (between the first and second bits) has been chosen; the probability is uniform so this is completely arbitrary. The parents, crossover process and resulting children are illustrated in table 3.2. We can see that from parents 3 and 5, offspring have been produced with values 1 and 7.

On the other hand, in uniform crossover each bit has a uniform probability of being exchanged [Vose, 1999]. So for the above example, neglecting the possibility of cloning by exchanging 0 or 3 bits, either 1 or 2 bits could be exchanged, with positions chosen randomly. It can be surmised therefore that uniform crossover provides a greater range of potential offspring from any given parents, whereas 1-point crossover is the more consistent method, involving less random selection. It is worth noting that cloning is still possible even when neglecting extremes as is done here; in the example above if (by either method) only the last bit is swapped, this would result in clones as the last bit is the same in both parents. Whether or not both children are included in the next population depends on the specific genetic algorithm; Coley [1999] gives an example in which only half the population are subjected to crossover (determined by fitness) and both children and also the parents are carried through to the next generation, whereas Vose [1999] suggests that an efficient crossover regime involves selecting only one of each pair of children to be used in the next generation. In the context of GAs, the principal of allowing a small percentage of the very fittest parents to persist to the next generation and breed again is known as elitism.

The process of evolution by natural selection provides an excellent heuristic; this is demonstrable simply by the success and diversity of life on Earth. However, there are certain aspects of the biological process that would be detrimental to the operation of a GA and are therefore not modelled. For example the notion of gender is not necessary; solutions are generally asexual and can “mate” with any other solution. Also in-breeding does not have the severe repercussions that it does in the natural world; whilst it is likely to homogenise the population in that particular area and hence stagnate the search somewhat, the characteristic genetic decay is generally not modelled.

Rationale

This technique was selected for several reasons; firstly, and perhaps most importantly, it is widely regarded as the de-facto standard optimisation method for large non-linear problems where gradient information is not available. There is substantial precedent for its application in most areas of optimisation, but particularly BTO; it is used almost exclusively in the literature (see section 2.4). Whilst there is some evidence that the latest trends are moving to

supersede GAs with more advanced techniques in some related areas [eg. Csébfalvi, 2013], it is still a safe option for most optimisation problems.

Secondly, the intrinsically population-based search procedure lends itself well to a multi-objective problem. This is due to the fact that when one has multiple objectives, the optimum is often not a single point but an optimum trade-off between the objectives, called a Pareto front. This is particularly true of BTO, as the objectives of minimising thermal discomfort and minimising energy use are almost always intrinsically oppositional; if they were not then building design would be much simpler and there would be little point in performing the present work. The definition of a Pareto front is a set of solutions, each of which having no superior competitor in respect of all objectives. So considering the population in an *n*-dimensional space, where *n* is the number of objective functions, a candidate will participate in the Pareto front as long as there is no other single solution with more optimal coordinates than the candidate in all axial directions. The population-based search means that the entire Pareto front can be identified simultaneously, as opposed to sequentially. It also makes the process resistant to converging on local optima, as once the population begins to coalesce to the global Pareto front, individuals and even groups that have fallen into local optima are increasingly likely to die off at each generation.

Chapter 4: Initial Model Development

4.1 Chapter Overview

This chapter describes the initial development of the methodology employed by T-BOT to accomplish BTO; subsequent studies motivated further developments of the methodology, however these are detailed in chapter 5. Section 4.2 is a brief overview of the general methodology and section 4.3 details the initial building simulation model development. Section 4.4 describes the metamodelling technique used, section 4.5 describes the samples that were used to train the metamodels, and section 4.6 describes the parameters of the optimisation. Finally, section 4.7 describes the programmatic architecture of T-BOT.

A portion of this work was presented in a conference paper written by the present author, with support from the project supervisors [Cowie et al., 2015]; specifically elements of sections 4.4, 4.6.1 and 4.7.

4.2 Methodology Overview

T-BOT uses a metamodelling based BTO methodology. This involves complete separation of the building simulation and optimisation, linked by a metamodelling procedure. Figure 4.1 shows a flowchart of the basic steps involved. The first part of the process, the simulation phase (shown on fig. 4.1 bordered in red) is simply simulating sample sets of the design space. The raw results of this sample set are then stored.

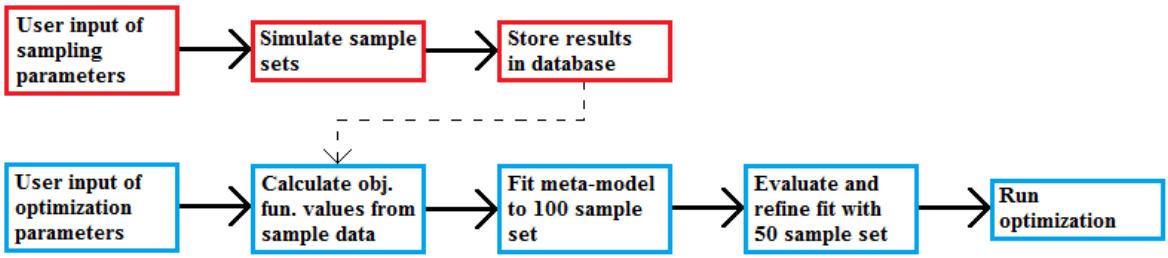

Figure 4.1: Flowchart of the basic elements in the methodology of T-BOT

The second part of the process, the optimisation phase (shown on fig. 4.1 bordered in blue), begins with reading in the results of the desired sample sets. This data is then processed and used to train metamodels which replace the building simulation for evaluating the objective functions. In this way, computation time is substantially reduced from what it would be if the objective functions were directly evaluated in the optimisation.

4.3 Thermal and Airflow Model Development

Building thermal simulation in the present work was accomplished using the open source software ESP-r. This section describes model development and discusses how ESP-r was used for the present project; for a description of the more general calculation methodology of ESP-r the reader is referred to section 3.2.

4.3.1 Test model geometry and boundary conditions

A relatively simple small room test model was constructed in ESP-r to facilitate development and testing of the BTO approach. The model is introduced here as it was instrumental in the development of the building simulation methodology, and it provides an exemplar for subsequent sections.

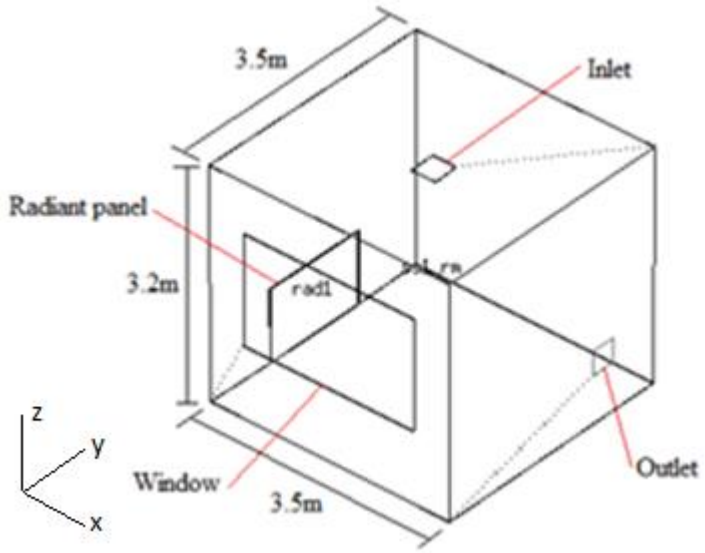

Figure 4.2: Diagram of the test model used during development

Figure 4.2 shows the test model geometry. The window was on the only external wall, which faced SW; the other boundaries were internal and considered as adiabatic due to the assumption of similar conditions on the other side. The geometry is intended to be representative of a fairly generic small room, such as a single-bed hospital room or an office. The radiant panel shown in Fig. 4.2 is in one of four possible positions, one at the middle of the base of each wall. The inlet and outlet shown in Fig. 4.2 represent a full mechanical HVAC system. The window was not modelled as operable; this is in contrast to the case study models detailed in chapter 6 which did consider the windows able to be opened. No door was modelled as all internal boundaries were considered adiabatic, rendering the construction of the internal walls largely moot. Also, the room was not considered to be in ventilation contact with any other spaces beyond that provided by the HVAC system, so the location of the door

was also moot. Table 4.1 gives details of the construction and material properties of all surfaces; the constructions and materials were all provided in databases within ESP-r, except the external wall which is detailed below. Table 4.2 gives details of casual gains assumed in the model; these values were assumed to be reasonable based on typical values and recommendations given in the help text of ESP-r.

Construction	Material	Thickness (mm)	Conductivity (W/mK)	Density (kg/m^3)	Specific Heat Capacity (J/kgK)	IR Emissivity (0-1)	Solar Emissivity (0-1)	U Value (W/m^2K)
Internal Walls*	perlite plasterboard	12	0.18	800	837	0.91	0.6	1.552
	breeze block	150	0.44	1500	650	-	-	
	perlite plasterboard	12	0.18	800	837	0.91	0.6	
Suspended Ceiling*	gypsum plasterboard	13	0.42	1200	837	0.91	0.5	4.976
Floor (Ground)*	earth	600	1.28	1460	879	-	-	0.699
	gravel fill	150	0.52	2050	184	-	-	
	heavy mix concrete	150	1.4	2100	653	-	-	
	air gap	50	n/a***	n/a	n/a	-	-	
	chipboard	19	0.15	800	2093	-	-	
	carpet	6	0.06	186	1360	0.9	0.6	
External wall	approximation**	325	0.137	1500	700	0.9	0.5	0.393
Window	plate glass	6	0.76	2710	837	0.83	0.5	2.811
	air gap	12	n/a***	n/a	n/a	-	-	
	plate glass	6	0.76	2710	837	0.83	0.5	
Radiant panel****	white painted aluminium	4	210	2700	880	0.82	0.32	n/a
* These surfaces were considered adiabatic so the material properties shown here are largely moot								
** This material was an approximation of a standard UK brick building construction (see *External wall* sub-section)								
*** All air gaps assumed to have R value = 0.17 m^2K/W								
**** This was modelled as a distinct zone. Insulation was added to all four edges to prevent excessive heat loss. Back face (ie. next to wall) was modelled as aluminium -> 10mm air gap -> wall. Front face was as shown above.								

Table 4.1: Constructions and material properties used in the test model

Source	Sensible (W)	Latent (W)	Radiant/Convective
Occupants	75	20	0.5/0.5
Lights	20	0	0.7/0.3
Equipment	50	5	0.5/0.5

Table 4.2: Casual gains assumed in the test model

External wall

Material	Thickness (mm)	Conductivity (W/mK)	Density (kg/m^3)	Specific Heat Capacity (J/kgK)	IR Emissivity (0-1)	Solar Emissivity (0-1)	U Value (W/m^2K)
brick	100	0.96	2000	650	0.9	0.7	0.393
glasswool insulation	75	0.04	250	840	-	-	
air gap	50	n/a*	n/a	n/a	-	-	
breeze block	100	0.44	1500	650	0.9	0.65	
* Air gap assumed to have R value = 0.17 m^2K/W							

Table 4.3: Material properties of default external construction in ESP-r

As the conductive path by which the external climate impacts on the thermal performance of the room, this was a fairly important aspect of the model. Material properties of the external

wall were considered as design variables; this is detailed in section 4.6. In order to facilitate application of these design variables, the multi-layer construction was approximated as a single material of averaged properties. This allowed the wall to be considered as a whole, as opposed to individual layers which would have greatly increased the number of design variables.

The construction upon which the approximation was based was the default external wall construction provided within ESP-r. This was, from outside to in, 100mm brick -> 75mm glasswool insulation -> 50mm air gap -> 100mm breeze block. Table 4.3 shows the material properties of this construction. Obtaining an averaged value of thermal conductivity for the approximation was straightforward, as a weighted mean could be taken which resulted in an equal U-value:

$$\bar{\lambda} = \frac{\sum_{i=1}^{n} \left(\frac{d_i}{d_T} \lambda_i \right)}{n} \tag{4.1}$$

Where $\bar{\lambda}$ is the weighted mean thermal conductivity (W / m K), n the number of layers, d_i the thickness of layer i (m), d_T the total thickness of the wall and λ_i the thermal conductivity of layer i. For the special case of the air gap layer, an effective thermal conductivity was calculated from the thermal resistance by the following formula:

$$\lambda = \frac{d}{R} \tag{4.2}$$

In the case of density and specific heat capacity, a representative average is not so straightforward. This is because heat storage in walls with air cavities is a special case; instead of heat propagating through the wall only by conduction as is normally the case, when it reaches the air gap, radiation and convection also play a part. The values given in table 4.1 for density and specific heat capacity for the approximation material were arbitrarily selected based on typical values of these parameters. To verify that this did not have a significant effect on room thermal performance, simulations were performed to compare operative temperature between models with the external wall construction shown in table 4.3, and that given in table 4.1. These were performed on slight variations of the test model with a simplified non-CFD ventilation model, for the mid-season case. The results of these simulations are shown in fig. 4.3.

It can be seen from fig. 4.3 that the approximated external wall construction does not affect room performance significantly, as the temperature profiles differ by only around 0.15°C at the maximum.

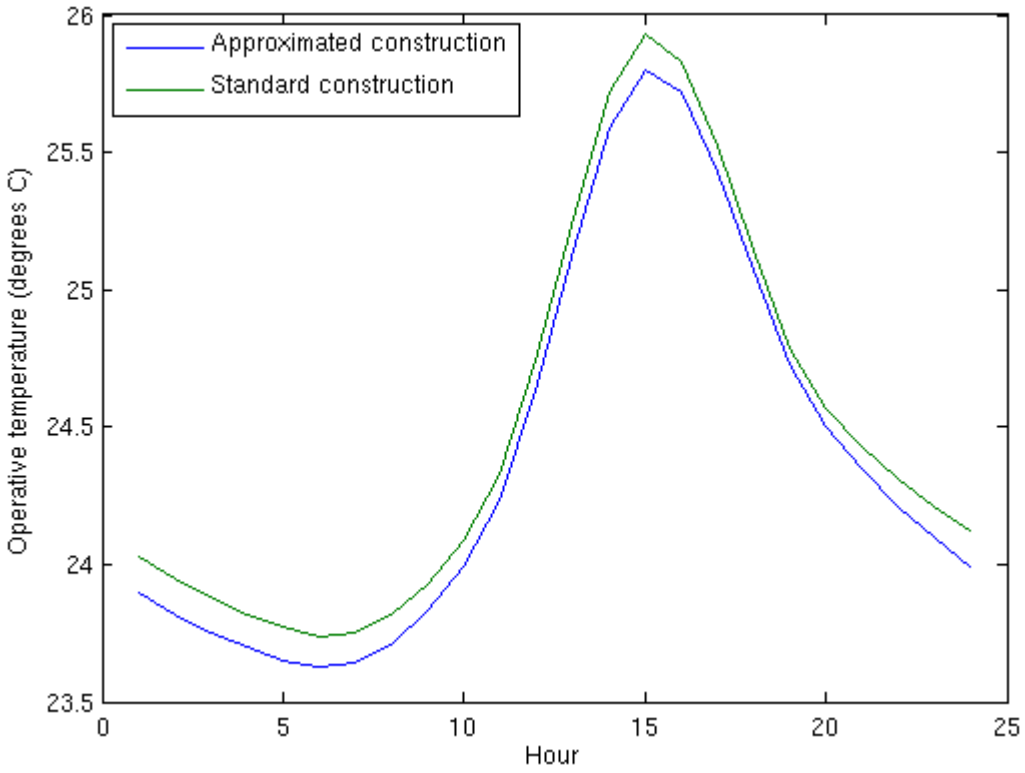

Figure 4.3: Results of simulations on the test model, for a mid-season case, using the approximated external wall material shown in table 4.1, and the actual construction shown in table 4.3.

Thermal Control

Thermal control of the room is achieved by a radiant panel and mechanical HVAC system. The radiant panel was intended to represent a conventional radiator; this is reflected in the choice of construction and materials as shown in Table 4.1. It was modelled using a thin-zone approach, whereby the panel is represented as a distinct zone with 6 surfaces as normal. This method has been shown to adequately represent a radiator in past studies using ESP-r, included with the source code as exemplars. Energy is delivered to the air point inside the thin-zone according to a basic heating-only, set-point based control algorithm, but the convection coefficients between the air point and the surfaces were manually overridden and greatly increased on the front face to simulate water inside the zone instead of air. Whilst simplifying the modelling somewhat, this approach has the disadvantage of introducing extra assumptions in the modelling in the form of the manual convection coefficients. No information was available with which to explicitly validate values, so values were assumed based on the aforementioned examples distributed with ESP-r. These produced reasonable results, and furthermore the solutions were not found to be particularly sensitive to these values, so the approach was also adopted for the case study models.

The room was assumed to be mechanically ventilated with air supplied and extracted at a constant volume flow rate. This was set at 6 ACH, as recommended for single patient rooms by Department of Health guidance [DoH, 2007]. An approximation of a four-way diffuser was modelled at the inlet, ensuring that the air was diverted through 90° to horizontal upon entering the room, instead of being propelled straight downwards.

Both systems were controlled with constant temperatures. For the radiant panel, energy was supplied to maintain a constant internal temperature. For the HVAC system, air was supplied at a constant temperature. Panel temperature and supply air temperature were two of the main design variables programmed with the test model; a detailed discussion may be found in section 4.6. It is noted that these control mechanisms are not necessarily representative of the majority of real situations, though constant set-points are used in some existing hospital wards [eg. Lomas and Giridharan, 2012]. In line with the purpose of the test model, the relatively simple control mechanisms allowed examination of model performance with somewhat greater clarity than a more complex system would allow.

Climate

For the purposes of developing and testing the model, theoretical climate data was used in lieu of real data. This was done in order to allow rigorous study of model performance without complications arising from random climate fluctuations. The climate data was defined in terms of six variables: dry bulb temperature, direct normal solar irradiance, diffuse horizontal solar irradiance, wind direction, wind speed and relative humidity. Dry bulb temperature and solar metrics were generated sinusoidally according to Equations 4.3 and 4.5 respectively, whilst wind was not considered (kept constant at 0) and relative humidity was kept constant at 50%. This was because humidity was not considered in the test model (only temperature), and wind was assumed to have no influence due to the non-operable window.

$$\theta_{db,h} = \sin(h_r) . \frac{(max - min)}{2} + mid \tag{4.3}$$

Here $\theta_{db,h}$ is the external dry bulb temperature at hour h (°C), h is the hour (integer in the range 1-24), h_r is the hour mapped into radians and is calculated according to Equation 4.4, *max* and *min* are the limits of dry bulb temperature (see Table 4.4 for values), and *mid* is the mean of *max* and *min*. Values of *max* and *min* were arbitrarily selected based on typical UK climates; given the purpose of the test model the climates were not considered to be important provided they were broadly representative.

$$h_r = h. \frac{\pi}{12} - \frac{\pi}{2} \tag{4.4}$$

Unlike dry bulb temperature, solar metrics do not vary uniformly throughout a 24 hour period; in general they increase during daylight hours, peaking around midday, and then drop off to 0 again with the fading daylight. This made calculation of sinusoidal variation a little more complex, as the rise-and-fall curve had to coincide with daylight hours; it was assumed that daylight hours were a 14 hour period from 6am to 7pm inclusive.

$$SM_h = \begin{cases} \sin(h_s) . \frac{(max-min)}{2} + mid, & h_{12centred} \geq 6 \\ 0, & h_{12centred} < 6 \end{cases} \quad (4.5)$$

Where: SM_h is the value of solar metrics at hour h (W/m^2), h_s is calculated according to Equation 4.6, and $h_{12centred}$ is calculated according to Equation 4.7.

$$h_s = (h_{12centred} - 5).\frac{\pi}{7} - \frac{\pi}{2} \quad (4.6)$$

$$h_{12centred} = \big||h - 12| - 12\big| \quad (4.7)$$

Case	Dry Bulb Temp. (°C)	Direct Normal Solar (W/m^2)	Diffuse Horizontal Solar (W/m^2)	Wind Direction (° clockwise from North)	Wind Speed (m/s)	Relative Humidity (%)
Winter	0-10*	0-100*	0-100*	0	0	50
Mid-season	7.5-17.5*	0-200*	0-200*	0	0	50
Summer	15-25*	0-300*	0-300*	0	0	50
* Sinusoidal variation within this range						

Table 4.4: Climate values for all climate cases considered for the test model

Three separate climate cases were considered; a winter case, a summer case, and a mid-season case between the two. Table 4.4 gives the values of the climate variables for each case. Solar incidence angle calculations which are coded within ESP-r are dependent on the time of year, so a broadly appropriate date had to be specified for each climate case. These were:

- 1st January for the winter case
- 1st April for the mid-season case
- 1st July for the summer case

4.3.2 Modelling domains

Here the various domains and other aspects of the building modelling are detailed, and specific challenges in the development of the initial methodology are highlighted. As a key interest in the model output is the distribution of thermal comfort with a room, both DTM and

CFD simulations are carried out in ESP-r. The interaction between these two domains is therefore an important consideration. All simulations in the present project were run with a time-step of 1 hour, unless otherwise stated.

The first and most obvious modelling domain is the basic building simulation, DTM, as described in sub-section 3.1.1. This is the backbone of the process, simulating the thermal interaction of the room with the external climate over time; the other domains serve to augment specific aspects of this core simulation. For the present project, the main role of this domain was to determine the radiant field of the room, encapsulating the spatial variation. In ESP-r, this was done by placing a grid of MRT sensors throughout the room. These sensors are notional bodies within the model that do not interact with the thermal environment of the room, but rather provide directives for the simulation to evaluate MRT at these specific locations. The program was limited to a maximum of 4 sensors per simulation, so in cases where more than 4 evaluation locations were required, more than one simulation was required. This would have increased computation time by orders of magnitude, were it not for the constraint to one-way conflation between the DTM and CFD. As discussed in sub-section 3.1.3, the one-way conflation prevented the results of CFD from influencing the results of the DTM. As such, the same simulations performed with and without CFD had identical radiant fields. Further considering that the location of the MRT sensors did not influence the simulation in any way, and only measured the radiant field, it follows then that only one simulation need be performed with CFD. Subsequent simulations, provided the only difference from the first was the location of the MRT sensors, could be assumed to have identical airflow patterns to the first simulation. Thus the CFD solver was disabled for runs of the same case beyond the first. The simulation time without CFD was greatly reduced to a matter of seconds, allowing the computation time of samples to remain largely unaffected with an increasing number of MRT sensor locations.

By default, solar incidence through glazing is distributed to the surfaces within the room according to area-weighting. To prevent this assumption, shading analysis was done for all models. This functionality is provided within ESP-r, and uses a ray-tracing technique to analyse room geometry in tandem with year-round solar incidence to provide data for the simulation to assign a more accurate distribution of solar gains. This gave a more accurate representation of local differences in MRT.

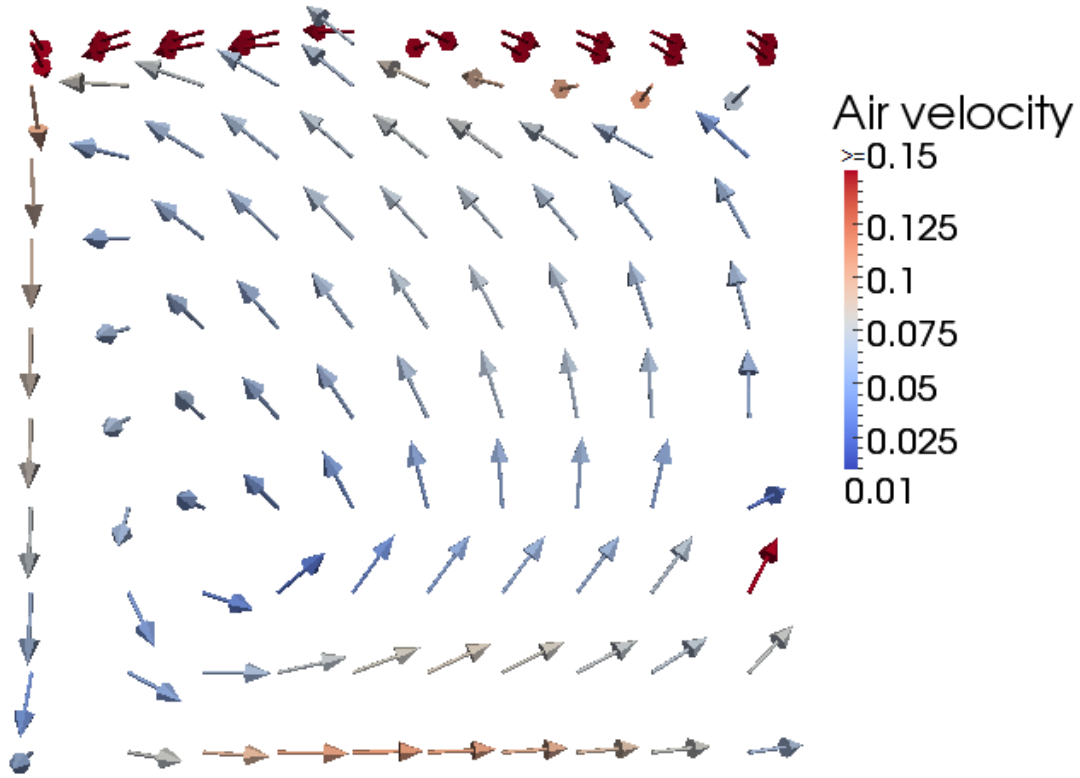

Figure 4.4: Airflow in the test model with a CFD grid resolution of 10x10x11 cells, on a plane normal to the Y axis at approximately the position of the outlet.

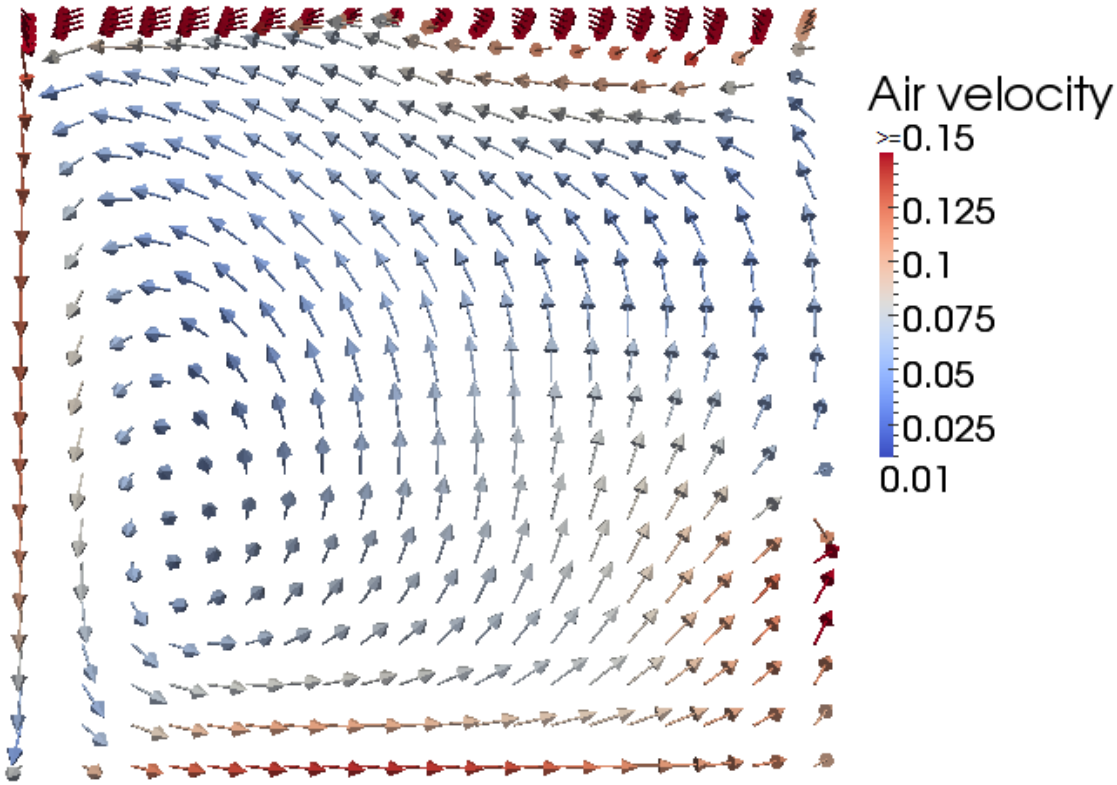

Figure 4.5: Airflow in the test model with a CFD grid resolution of 20x20x22 cells, on a plane normal to the Y axis at approximately the position of the outlet.

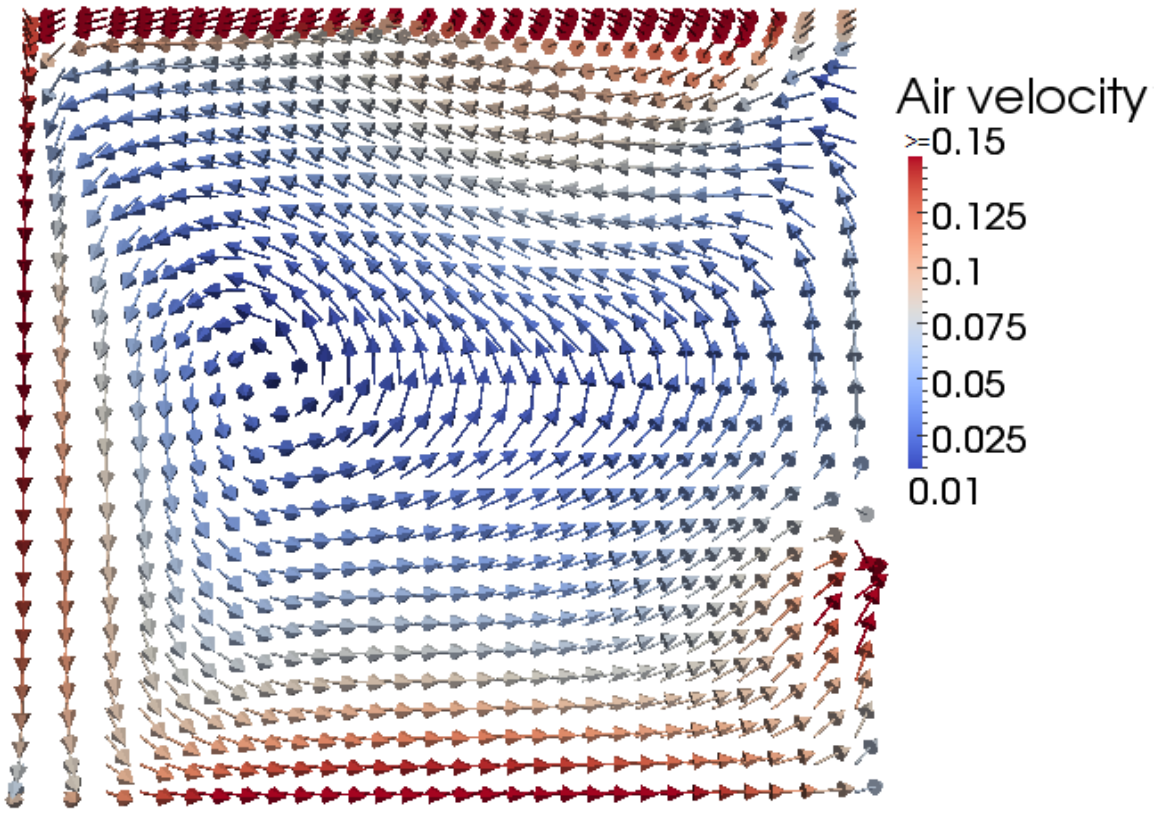

Figure 4.6: Airflow in the test model with a CFD grid resolution of 30x30x30 cells, on a plane normal to the Y axis at approximately the position of the outlet.

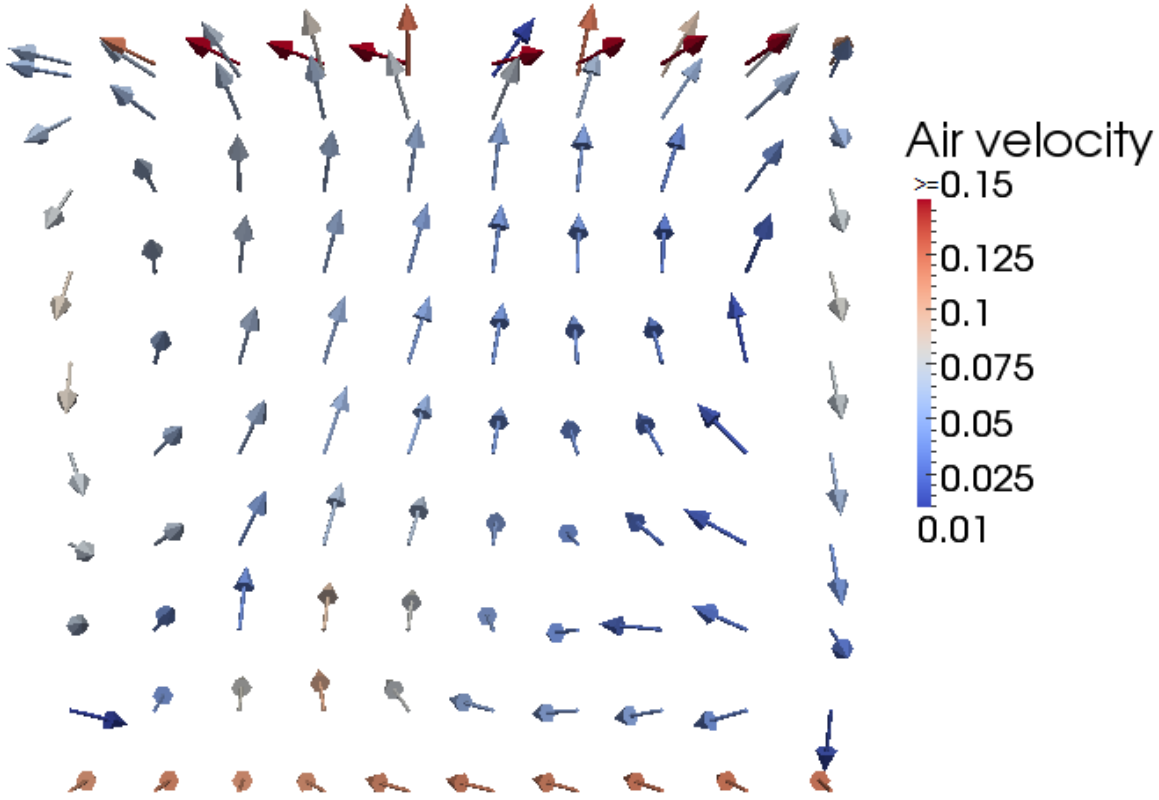

Figure 4.7: Airflow in the test model with a CFD grid resolution of 10x10x11 cells, on a plane normal to the X axis at approximately the position of the inlet.

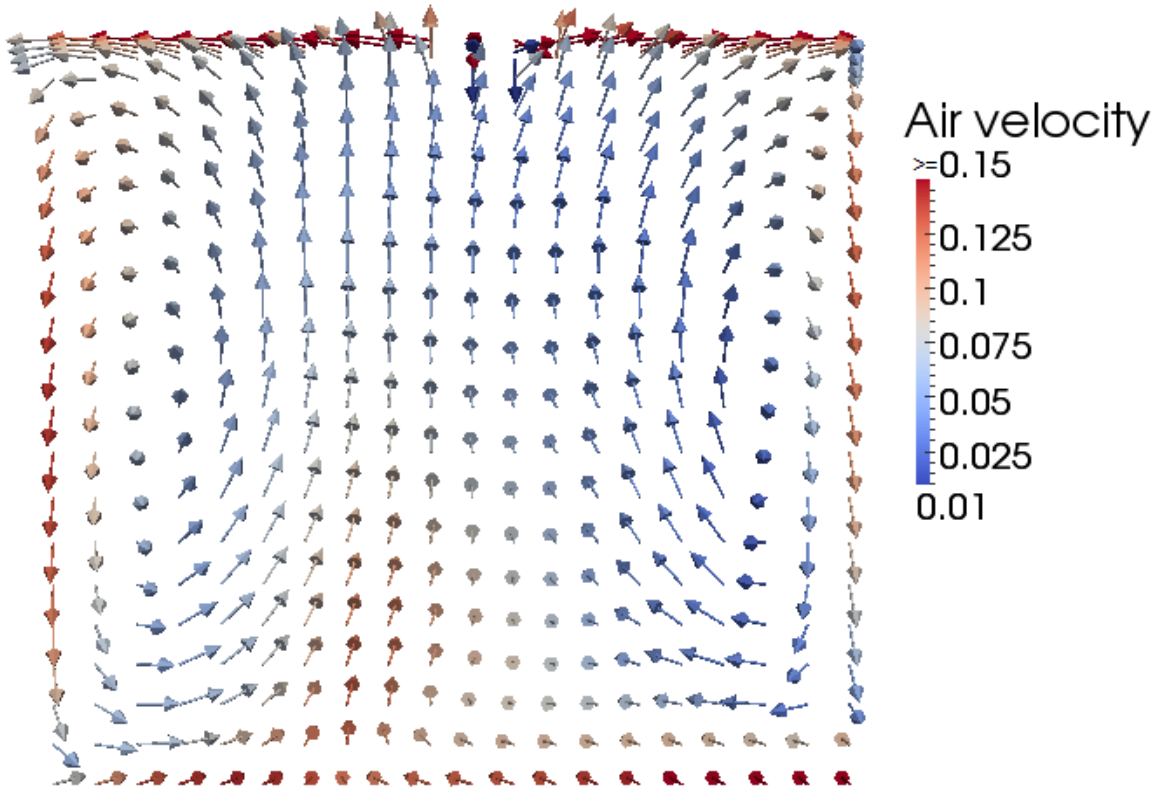

Figure 4.8: Airflow in the test model with a CFD grid resolution of 20x20x22 cells, on a plane normal to the X axis at approximately the position of the inlet.

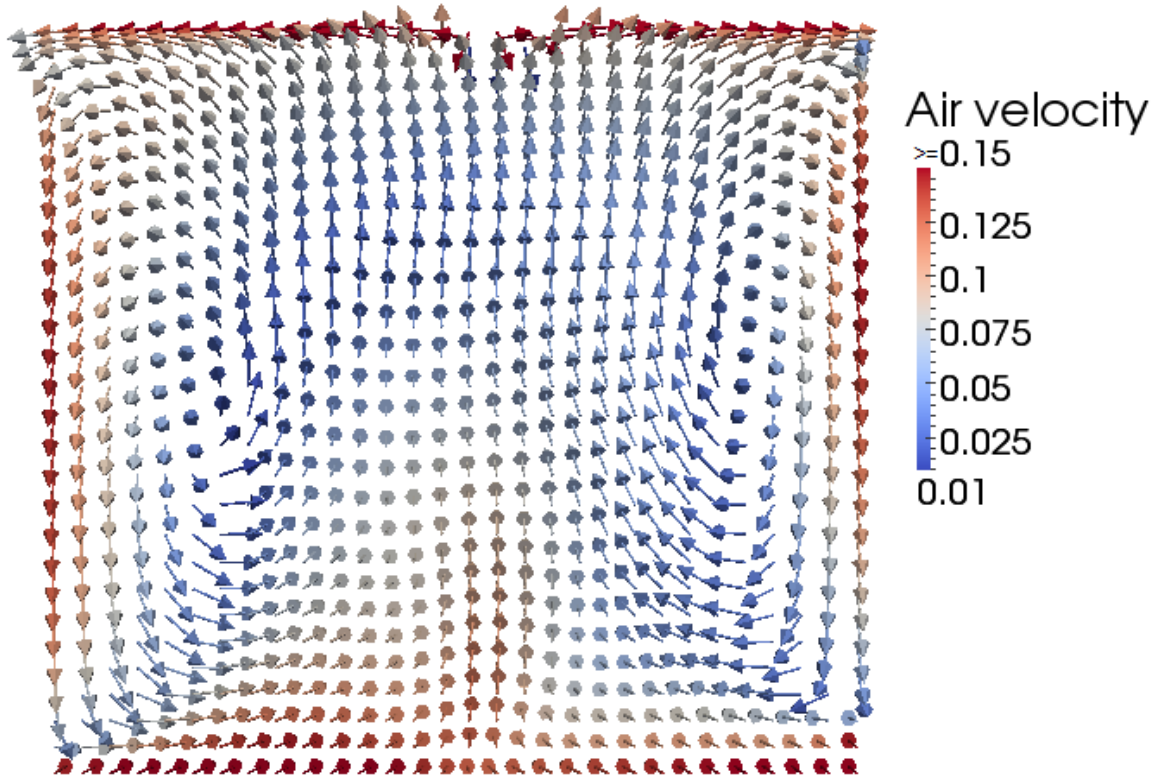

Figure 4.9: Airflow in the test model with a CFD grid resolution of 30x30x30 cells, on a plane normal to the X axis at approximately the position of the inlet.

Airflow

The airflow modelling within DTM simulations is generally highly simplified and only considers bulk flows, as discussed in section 3.1.1. As the present project called for evaluation of spatial variation of thermal conditions within a room, this aspect of the simulation was augmented by CFD analysis. Boundary conditions for the CFD domain are shown in table 4.5.

Surface	**Boundary Condition**	**Face****	**Area**
Floor	Solid surface	Bottom	12.25
Ceiling	Solid surface	Top	12.14
South wall	Solid surface	South	7.45*
East wall	Solid surface	East	11.08*
North wall	Solid surface	North	11.20*
West wall	Solid surface	West	11.20*
Window	Solid surface	South	3.75
Radiant panel	Solid surface	South, east, north or west	1.50
Inlet	Velocity (inlet)	Top	0.11
Outlet	Velocity (outlet)	East	0.12
* When radiant panel is on this face, area is given value minus radiant panel area			
** East is in the positive X direction and north is in the positive Y direction (as shown in Figure 4.2).			

Table 4.5: CFD boundary conditions of the test model.

The CFD had a very high computational requirement in comparison with all other modelling domains, making it by far the biggest contributor to computation time. This was an important consideration due to the sheer number of CFD simulations required. Each simulation of 24 hours required 48 separate CFD solutions, with two simulations per time-step; one pre-simulation and the actual simulation. For a 150 point sample, this meant 7,200 separate CFD simulations. Clearly it was necessary to minimize the computational requirement of the CFD wherever possible, whilst maintaining an acceptable level of accuracy. The most pertinent variable to this trade-off is the grid resolution (for a discussion of gridding applied to CFD see section 3.1.4), so a grid sensitivity study was performed with the test model to determine an acceptable grid sizing.

Figures 4.4 – 4.9 show airflow patterns on two different perpendicular planes, for three different grid resolutions. These results were gathered for a mid-season case without a radiator, and the results shown are those at 12 noon. It can be seen the general patterns of airflow are largely consistent in figures 4.4 – 4.6, and also in figures 4.7 – 4.9. Further to this, the air velocities for the sparsest grid resolution case (figs. 4.4 and 4.7) are broadly representative of the variation evident in the results of denser grids. The variation in velocity

is less consistent around the very edges of the room, however this was not considered a problem as conditions very close to the room surfaces are generally not as pertinent to thermal comfort as conditions closer to the middle of the room. To put it another way, it is likely to be far less useful to optimise thermal comfort of someone standing right next to the wall, than at positions closer to the middle of the room.

Likely due to the high air change rate in the room, very little temperature variation was evident in the room air aside from right next to the inlet. For the three grid resolution cases minimum and maximum temperatures were identical, and mean temperature differed by less than 0.2%.

The run-times of the cases were:

- 27 minutes for the 10x10x11 grid case,
- 6 hours for the 20x20x22 grid case,
- 38 hours for the 30x30x30 grid case.

These results exhibit a scaling of calculation time with grid resolution that is greater than linear, meaning that a lower grid resolution provides greater efficiency. As the results in figures 4.4 – 4.9 showed that the 10x10x11 grid case provides velocity and temperature fields and flow patterns that are representative of the more detailed simulations, it was decided that coarse grid CFD was sufficient for the purposes of the present work. To ensure that the level of detail in the CFD solutions did not fall below that shown in the results of the 10x10x11 case (figs. 4.4 and 4.7), the approximate grid size of this case, 0.35m, was set as a maximum for grid sizes in the case study models.

4.4 Metamodelling

A metamodel is a model of a model; a simple model used to predict the output of a more complex model with a reduced computational requirement. In optimisation this can be very useful, as it can potentially reduce the computation time of function evaluations by orders of magnitude. Metamodels are constructed from samples of possible solutions and are thus "trained" to predict the output from any given input. Metamodelling is a critical element of the BTO methodology employed for this work, as it provides the bridge between the simulation and optimisation phases; each function evaluation required by the optimisation is evaluated by the metamodel instead of a full building simulation. An accurate metamodel is essential to prevent misleading results from the optimisation, and a quick to evaluate

metamodel is also necessary to justify use of metamodelling as opposed to direct function evaluation.

4.4.1 The metamodelling technique – moving least squares regression

Selection

The simplest form of metamodelling is linear interpolation; drawing a straight line on a graph between two possible solutions and assuming that the host model behaviour follows this path. However for most applications of metamodelling this is far too simplistic to accurately predict the performance of the host model, hence more complex methods are used.

Since metamodelling is essentially curve-fitting, classical statistical methods such as least squares regression can be applied. However for multi-dimensional non-linear problems that are typically the focus of BTO metamodelling, more complex developments thereof are used. These include Kriging, a process originally developed for interpolating topological profiles, and moving least squares regression (MLSR).

At the more complex end of the spectrum of metamodelling techniques there are two methods that are popular in recent BTO literature; support vector regression (SVR) and neural networks (NN). Support vector regression is a statistical method that involves finding as flat a function as possible that satisfies all training data with an error of at most ε, selected by the user. It follows then that this method is best applied to closely linear data, so the method has been extended to cope with non-linear input data by pre-processing the data with a kernel function to map it onto some feature space such that the data becomes linearized [Eisenhower, 2012]. Neural networks are a general method of machine learning that can be applied to regression problems; however they have a significant disadvantage in that they typically involve random initialisation variables and hence exhibit rather poor reproducibility.

To select an appropriate method a brief investigation was conducted into the comparative performance of three of these techniques; MLSR, SVR and NN. These were selected due to their use in related BTO literature (see section 2.4). The three techniques were each used to construct metamodels from data gathered during the proof of concept study detailed in chapter 5; specifically a mid-season case, at a position in the centre of the room, for the full 24 hour period. The metamodels were constructed using a set of 100 samples, and then a separate set of 50 samples was used to assess the quality of the fit. Table 4.5 gives the results of the analyses in the form of maximum and average errors and R^2 values.

		Thermal Discomfort		Energy Use	
		Too cold	Too warm	Radiator	Ventilation
Neural Network	Max. error	0.139	0.033	0.029	0.022
	Ave. error	0.011	0.009	0.009	0.006
	R^2	0.992	1.000	1.000	1.000
Support Vector Regression	Max. error	0.862	2.791	1.557	6.204
	Ave. error	0.044	0.240	0.179	2.082
	R^2	0.747	0.717	0.867	0.452
Moving Least Squares Regression	Max. error	0.087	0.108	0.036	0.079
	Ave. error	0.037	0.013	0.012	0.002
	R^2	0.994	0.999	1.000	1.000

Table 4.6: Results of a comparison between metamodelling techniques for each objective function component.

As can be seen from table 4.6, the basic SVR used here provided quite poor results, whereas NN and MLSR gave far better approximations. It was beyond the scope of the present work to spend time developing a bespoke SVR implementation to achieve results comparable with those of Eisenhower [2012], so the method was discarded in favour of NN and MLSR. The approximations given by NN were generally marginally better fits than those given by MLSR, however a number of other issues also factored into the decision making process. Firstly, reproducibility and consistency; as mentioned above, neural networks generally utilise random initialisation variables in the method, meaning that two neural networks with identical non-random parameters, trained to approximate the same data, need not necessarily give the same approximations. On the other hand MLSR has no random parameters in the method, and therefore gives consistent and reproducible metamodels. Secondly, convenience and ease of use; the MLSR metamodels were constructed with the commercial optimisation program Altair HyperStudy, which provides a versatile GUI for the iterative process of analysing and refining the metamodels. NN was implemented within Matlab, meaning extra development work would have been necessary to create bespoke tools to allow time-efficient analysis and refinement of the metamodels. Further to this, the optimisation functionality of Altair HyperWorks was found to be rather more versatile than that provided in Matlab, so it made sense to group the metamodelling and optimisation together into the software that provided superior functionality. For these reasons, and due to the only marginally inferior (but still excellent) fits as compared to NN, MLSR was chosen as the metamodelling technique.

Implementation

In standard least squares regression, an approximation function is found that minimizes the sum of residuals:

$$\sum_{i=1}^{I}(g(p_i) - f(p_i))^2 \tag{4.8}$$

Where *g* is the approximation function, $\{p_i\}_{i=1}^{I}$ is a set of sample points, and $\{f(p_i)\}_{i=1}^{I}$ are the observed data at the sample points. Moving least squares extends the method by weighting the residuals by the distance of the sample points from the point at which an approximation is required.

A moving least squares approximation at a point *p* is therefore obtained by finding an approximation function that minimizes the sum of weighted residuals:

$$\sum_{i=1}^{I}\big(g(p_i) - f(p_i)\big)^2\beta(\|p - p_i\|) \tag{4.9}$$

Where β is a non-negative weighting function and $\|p - p_i\|$ is Euclidean distance between points *p* and p_i [Levin, 1998]. This means that instead of attempting to approximate the responses with a single function for the full range of *p*, the function becomes itself a further function of *p*, and the approximation can therefore describe the responses with far greater flexibility.

Forms of function *g* vary, but polynomials are a popular choice. Polynomials of order 1-3 were used for the present work, selected on an ad-hoc basis to achieve the best approximation. A Gaussian weighting function was used, of the form:

$$\beta(d) = Ae^{-\frac{(d-B)^2}{2C^2}} \tag{4.10}$$

Where *A* is a constant defining the height of the bell curve, *B* is the expected value and defines the position of the bell curve on the *d* axis, and *C* is the standard deviation and controls the width of the bell curve.

One common pitfall of moving least squares regression is the tendency to over-model the sample data resulting in rather noisy responses. This can present a significant problem for optimisation, as noisy responses often result in the formation of local optima. Whilst the optimisation method of choice went some way to mitigating the danger of local optima (discussed in greater detail in section 3.2), they could still significantly distort the Pareto fronts and so it was crucial to minimize this risk. The solution to this problem lies with the "closeness of fit" parameter within moving least squares regression, which in the case of Gaussian

weighting relates to *C* in equation 4.9. This parameter essentially defines the trade-off between how closely the data is modelled and smoothness of the responses. By defining two separate sample sets, one for initial fitting and one for validation of the fit, the optimum value of this parameter may be determined by solution of a simple optimisation problem. However for the present work it was found to be necessary to manually inspect the responses of the metamodels, as even at the optimum value of closeness-of-fit the responses would sometimes be over-fitted. This generally took the form of oscillatory or over-shooting behaviour at the point where thermal discomfort objectives became 0, as evident in figure 4.10; figure 4.11 shows the same responses with manually configured closeness-of-fit. As can be seen from comparison of these figures the manual configuration results in far smoother responses, which create a clearer optimum. This results in a smoother Pareto front.

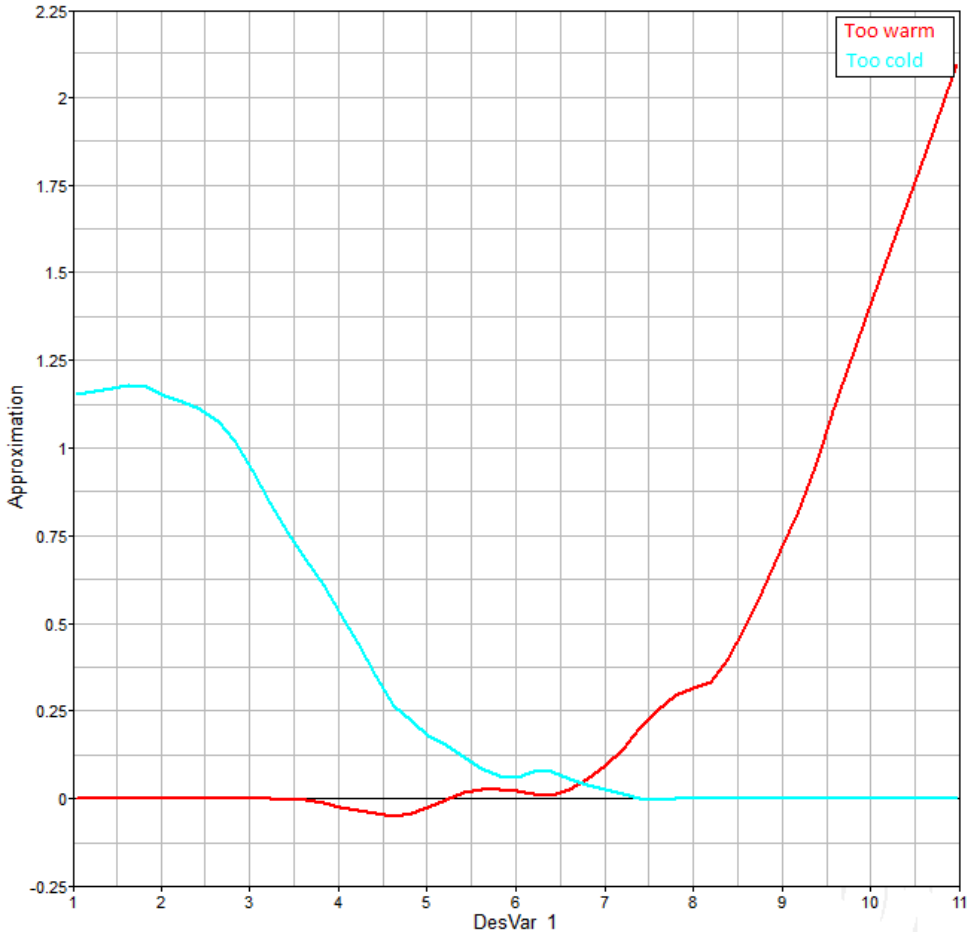

Figure 4.10: Over-fitted responses of MLSR metamodels of thermal discomfort objectives, with radiator temperature on the x axis, at ventilation temperature=1.

The responses were further improved by inclusion of an additional layer to the metamodelling process, a filter function that explicitly prevented values of metamodel responses below 0, setting any such values to 0. This was also necessary because negative values of the objective functions were mathematically implausible and had detrimental effects on the optimisation

process; negative objective components created false optima by falsely reducing the values of the objective sums.

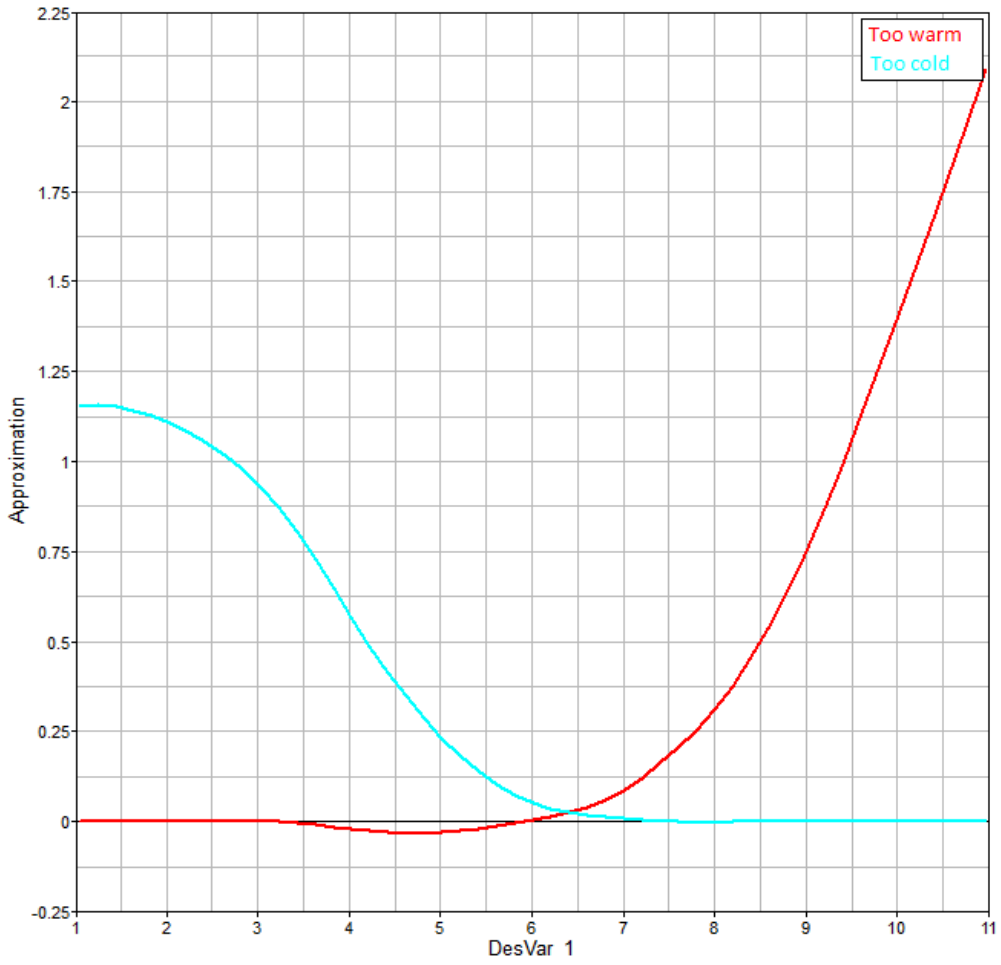

Figure 4.11: Smoothed responses of MLSR metamodels of thermal discomfort objectives, with radiator temperature on the x axis, at ventilation temperature=1.

4.4.2 The "one sample many optimisations" approach

The effective separation of the simulation and optimisation phases allows realisation of further economies, such as the facility to use a single sample for many different optimisations; a principal termed "one sample many optimisations" or OSMO in the present work. This arises from the fact that it is entirely academic where exactly the cut-off is made between the two phases. For example, it is intuitive to simulate a sample, store the resulting objective function values and construct a metamodel based on this data. However once this metamodel has been used to optimise the scenario, a whole new sample set is needed to optimise even a marginally different situation. Clearly this limits the usefulness of a metamodelling approach, as time savings are only made if the sample set is smaller than the number of function evaluations necessary for optimisation; granted this is often the case, but even greater time savings can be obtained simply by re-arranging the calculation order.

In an OSMO approach only the raw outputs from the simulations are stored, and calculation of objective function values from these data is shifted to pre-processing in the optimisation phase. This allows sample data to be independent of non-design parameters within the objective functions, and hence these parameters can be specified at run-time of the optimisation not the sampling. Objective function values can be calculated at any values of these non-design parameters, so an almost unlimited amount of different sets of metamodel training data can be calculated from the sample data, thus many different scenarios can be optimised using only a single sample set. In the context of the present work, these non-design parameters were:

- Evaluation location,
- Evaluation period, and
- Comfort criteria.

More information on how the OSMO approach was practically implemented can be found in section 4.7.

4.5 Sampling

In order to construct an effective metamodel, a sample of possible solutions is needed to "train" the metamodel. The distribution of this sample is known as a design of experiments (DOE). The inclusion of CFD in the building simulation greatly increased computational requirement, so it was found to be necessary to simulate sample sets in parallel on a High Performance Computing (HPC) cluster.

The DOE is important as the distribution of the sample determines how well it represents the sample space. If a sample is clustered in a particular area of the design space, that region will be better represented by the data. When metamodels are constructed from the data, it is important that the sample is evenly distributed so that metamodel fidelity is as consistent as possible across the design space.

The most basic DOE is a full factorial sample, which samples every point of an entirely discrete sample space. When this is applied to a continuous sample space, each dimension must first be discretized. If D dimensions are discretised to L levels each, the number of sample points is L^D. For example in a sample space of 3 dimensions (i.e. 3 design variables), with 10 values of each sampled (i.e. each dimension discretised to 10 levels), the number of samples would be $10^3 = 1000$. Fig. 4.12 shows a visualisation of such a design space with sampling points

highlighted, assuming dimensions are discretised uniformly and have ranges 1-11. Because the sampling points are orthogonal, it can be seen that each 2-D view (Figs. 4.12(a), (b) and (c)) only shows 100 (10x10) sampling points; one tenth of the total number of samples.

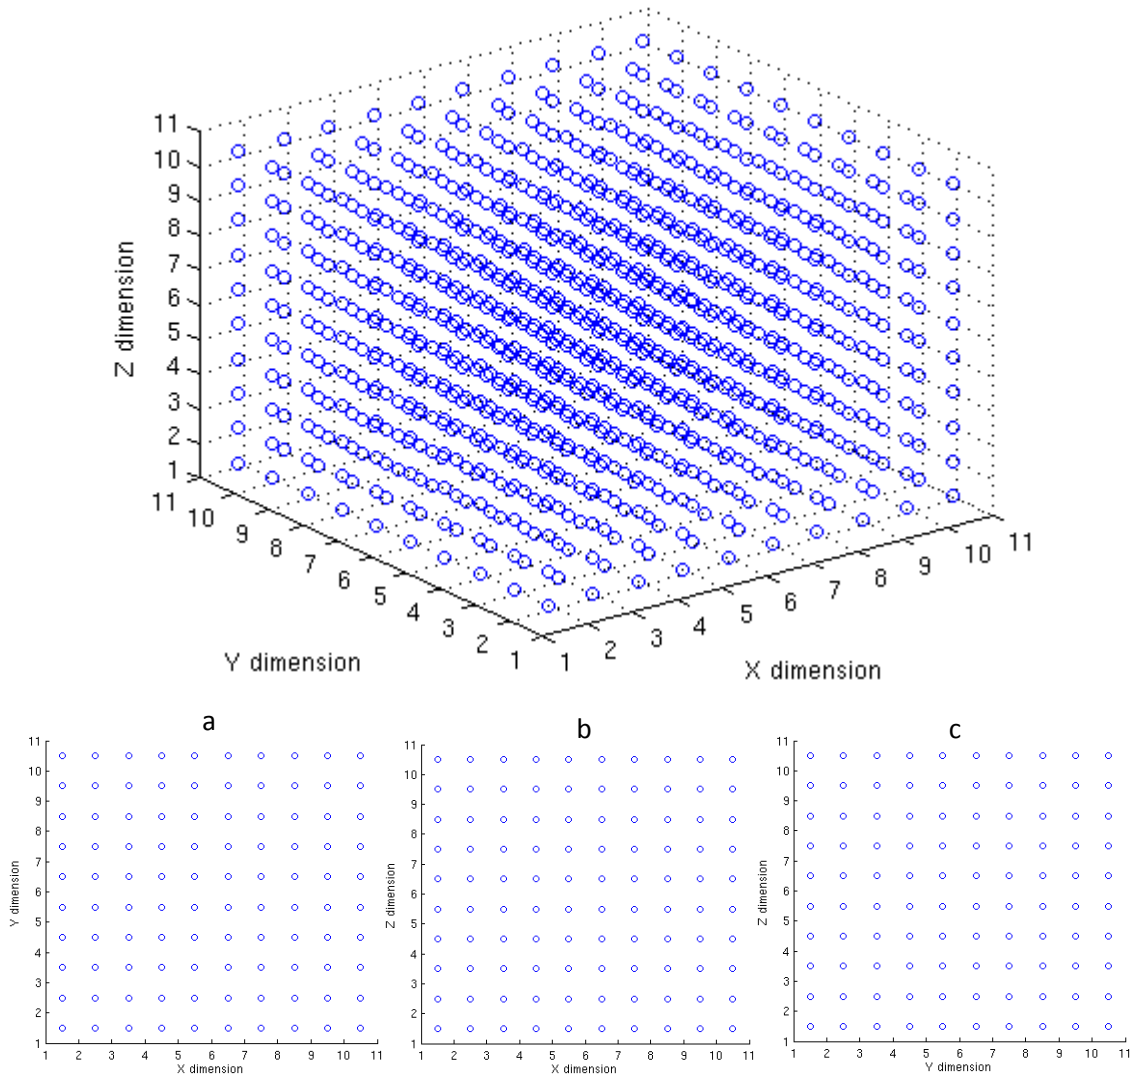

Figure 4.12: Visualisation of a full factorial sample of a 3 dimensional continuous design space; 2-D views from the Z, Y and X directions are shown in panels a, b and c respectively.

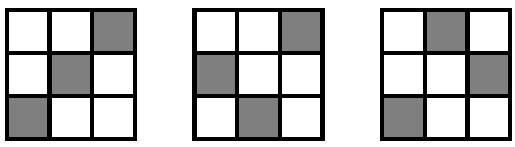

Figure 4.13: Examples of possibilities for a 3-point latin hypercube DOE of 2 variables.

The DOE used for the present work was a latin hypercube design; this is created by dividing each dimensional axis into *n* regions, where *n* is the number of points in the sample. The DOE must then conform to the rule that none of these regions may have more than one sample point within it; this is illustrated in Fig. 4.13. This design was iteratively generated 100 times,

and then the best DOE was chosen with the objective of maximising the minimum distance between sample points. Fig. 4.14 shows plots comparable with Fig. 4.12, for a latin hypercube design with 100 sample points. Because the design is by definition not orthogonal, it can be seen that each 2-D view (Figs. 4.14(a), (b) and (c)) shows all 100 sample points. This demonstrates a key advantage of latin hypercube designs in the context of optimisation; when the sample is used to model design spaces composed of subsets of the sample variables (i.e. 2 design variables out of 3 sample variables), a latin hypercube sample will tend to exhibit low redundancy in sample points. From comparison of Figs. 4.12 and 4.14 it is evident that this comes at the price of sample distribution; the factorial sample shown in Fig. 4.12 is clearly more evenly distributed than the latin hypercube sample shown in Fig. 4.14.

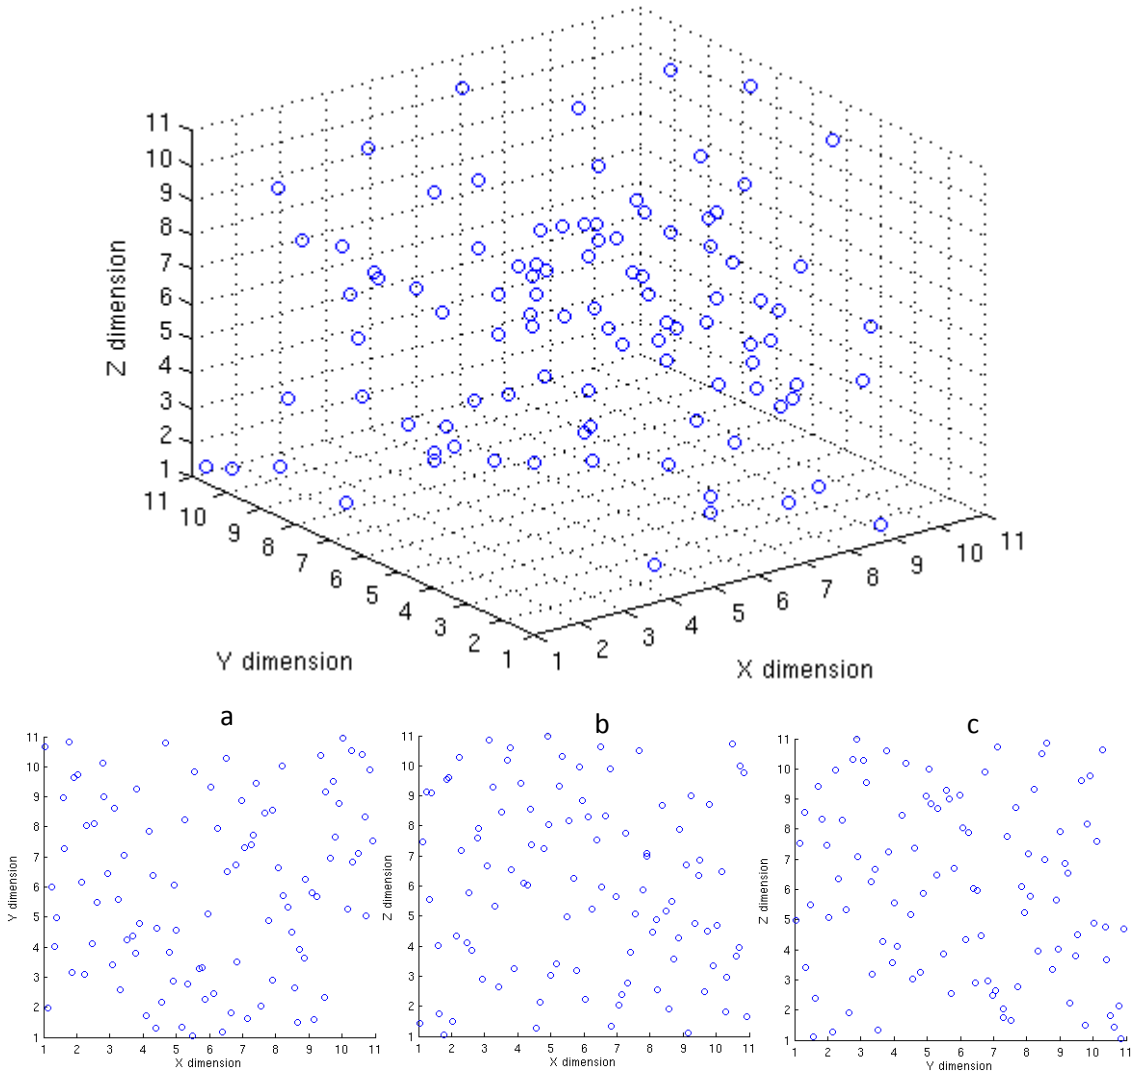

Figure 4.14: Visualisation of a latin hypercube sample of a 3 dimensional continuous design space; 2-D views from the Z, Y and X directions are shown in panels a, b and c respectively.

Also, though the latin hypercube sample exhibits superior scalability when considering subsets of design variables, it can be seen that the latin hypercube sample is considerably more sparse

in terms of the original (3 dimensional) sample space. Obviously this is to be expected with one tenth of the total number of sampling points, but it demonstrates another advantage of a latin hypercube sample in this context. An orthogonal sample design varies each parameter independently; any two axially adjacent sample points differ in terms of only 1 design variable. Where design parameters have well defined and independent effects on the objective functions, this can be beneficial as it allows clear assessment of the effects of each variable. However in the case of BTO design variable effects are likely to be strongly inter-related, due to the highly coupled nature of energy flow in buildings. A latin hypercube sample has no such restriction; indeed for any given sample point, by definition there can be no other sample point with the same value of any particular design variable. In this case of 3 dimensions for example, a factorial design with at most 100 sample points would sample $3\sqrt{100} = 4.6 \approx 4$ levels per design variable; clearly the latin hypercube design sampling 100 levels of each parameter will capture more of the inter-dependent variation.

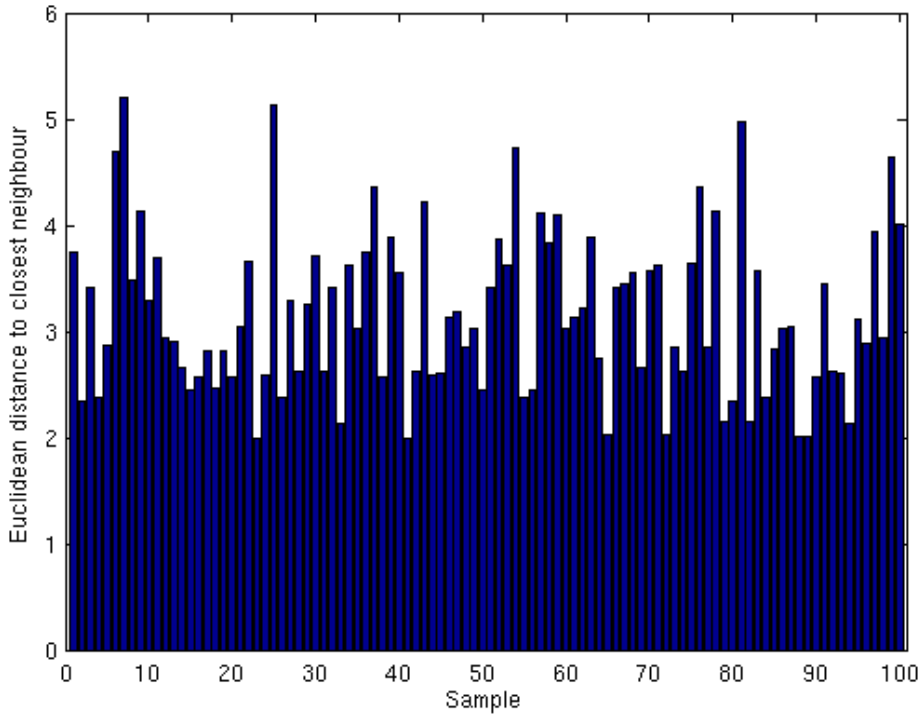

Figure 4.15: Sample dstribution bar-chart for a 100-point DOE of 5 design variables.

2 separate DOEs was used for the present project; first a larger sample that was used to initially fit the metamodels, and a second independent smaller sample that was used to evaluate and refine the fit. The shared memory parallelism of the HPC cluster necessitated all processes be on a single node. As there were 8 cores per node, this constrained the sampling to 7 parallel threads (the extra one was needed for the control process). Also, there was a 48 hour maximum run-time imposed on jobs. With the test model this resulted in a maximum sample size of around 100, keeping in mind that the categorical design variable of radiator

position effectively meant that every sample point had to be simulated 5 times, so a sample set of 100 points was actually 500 simulations (see section 4.6.2 for more details).

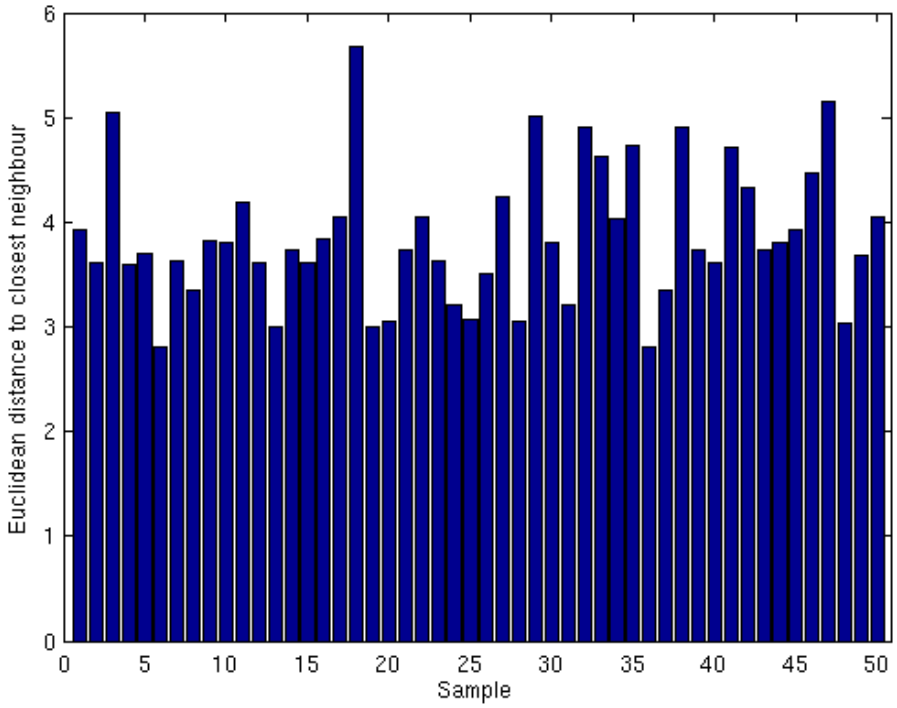

Figure 4.16: Sample dstribution bar-chart for a 50-point DOE of 5 design variables.

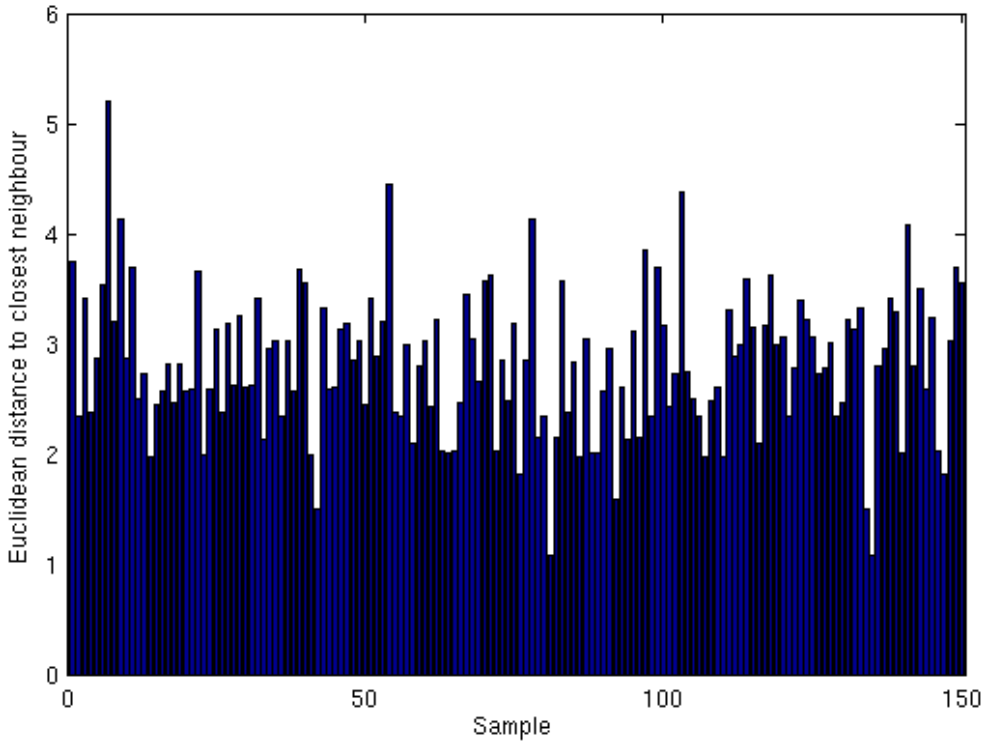

Figure 4.17: Sample distribution bar-chart for the combined 150-point DOE of 5 design variables.

The larger and smaller samples were generated entirely separately; no consideration was given to sample distribution with the two samples sets combined. Sample distribution bar-charts are given for sets of 100 and 50 samples in figs. 4.15 and 4.16, and the combined set in figure 4.17. These DOEs were in terms of 5 design variables (i.e. 5 dimensions), each having a range of 1 – 11.

The 100 point sample (shown in fig. 4.15) had a maximum variation in distance to closest neighbour of 62% of the maximum distance to closest neighbour. The mean distance to closest neighbour was 3.15, and the standard deviation 0.621 (20% of mean). For the 50 point sample (shown in fig. 4.16) the maximum variation was 51%, the mean 3.86 and the standard deviation 0.505 (13% of mean). The combined set (shown in fig. 4.17) was distributed markedly less uniformly, with a maximum variation of 79%, mean of 2.81 and standard deviation 0.510 (18% of mean). From these results it was surmised that whilst combining independent sample sets resulted in more extreme values of distance to closest neighbour, evidenced by the substantially greater percentage variation, the resulting distribution had mean and standard deviation comparable with those of the two individiual DOEs and so was considered acceptable.

Sample Size

Given the run-times for the test model and the restrictions on total run-time and number of parallel threads, a default sample size of 100 for the larger set and 50 for the smaller set was found to be appropriate. However the higher the dimensionality of the sample (ie. the number of design variables), the more sparse a sample becomes; this must be taken into account in minimum sample sizing.

A general rule of thumb for approximation using polynomials is that the number of samples should be, at the absolute minimum, double the number of coefficients in the polynomial. If *f(x)* is a polynomial of degree *d* with *n* variables, the number of coefficients in *f(x)* is given by the binomial coefficient:

$$\binom{n+d}{d} = \frac{(n+d)!}{d!\,n!} \tag{4.11}$$

So for a 2nd order approximation of a problem of 5 design variables, the number of coefficients is $\frac{(n+d)!}{d!n!} = \frac{(5+2)!}{2!5!} = 21$, so 100 samples should be sufficient to construct a viable initial metamodel. However, for a 2nd order approximation of a problem of 12 design variables, the number of coefficients is $\frac{(n+d)!}{d!n!} = \frac{(12+2)!}{2!12!} = 91$, so a set of at least 182 samples would be required as a minimum to construct a viable initial metamodel. Given that the maximum

number of design variables in the test model was 5, the default sample sizes given above were considered appropriate. However for the case studies the DOE technique was improved; this is detailed in chapter 6.

4.6 Optimisation Parameters

This section defines the parameters of the optimisations performed on the test model. This includes objective functions, design variables and input variables.

4.6.1 Objective functions

There were two objectives considered in the present project; thermal comfort and energy use. It is convenient for the purposes of optimisation to formulate the objectives such that they need to be minimised, so thermal comfort was considered inversely i.e. the objective function was a measure of thermal discomfort.

Thermal discomfort

Operative temperature was used as the measure of thermal comfort for the test model. The PMV model was also implemented as an objective for case studies; this is described in chapter 6. Operative temperature is a measure of the temperature experienced by an occupant, taking into account the radiant temperature field, air temperature and air velocity. It is defined in CIBSE Guide A [CIBSE, 2006] as:

$$\theta_o = \frac{\theta_a\sqrt{10v} + \theta_r}{1 + \sqrt{10v}} \tag{4.12}$$

Where θ_o is operative temperature (°C), θ_a is air temperature (°C), θ_r is the mean radiant temperature (°C) and *v* is the air velocity (m/s). At air speeds below 0.1m/s this is assumed to be:

$$\theta_o = \frac{\theta_a + \theta_r}{2} \tag{4.13}$$

The operative temperature formulation of objective function was taken as time-averaged deviation of operative temperature from user defined comfort limits:

$$F_1(\boldsymbol{D}) = \frac{\sum_1^n \left(f(\theta_o)\right)_i}{n} \tag{4.14}$$

Where $F_1(\boldsymbol{x})$ is the objective function value subject to set of design variables $\boldsymbol{D}$, n is the number of hours in the optimisation period, i is the evaluation location index, and $f(\theta_o)$ is calculated according to equation 4.15:

$$f(\theta_o) = \begin{cases} \theta_{o,min} - \theta_o, & (\theta_o < \theta_{o,min}) \\ 0, & (\theta_{o,min} \leq \theta_o \leq \theta_{o,max}) \\ \theta_o - \theta_{o,max}, & (\theta_o > \theta_{o,max}) \end{cases} \quad (4.15)$$

Where $\theta_{o,min}$ and $\theta_{o,max}$ are the lower and upper bounds of comfortable operative temperature respectively. Default values for these variables were taken from CIBSE Guide A [CIBSE, 2006]. For summer climate conditions these were 23-25 °C, for winter 22-24 °C, and for mid-season cases values were interpolated between the two at 22.5-24.5 °C. These values may be assumed hereafter where required in the present work, unless otherwise specified in the text. This objective function is very similar to one identified by Bouchlaghem and Letherman as being most efficient for BTO out of a range of possibilities [1990].

The two components of this objective, deviation above and below the comfort range, were programmed as separate functions, which were then summed to form the actual objective function. These components were separated for two reasons; firstly the metamodels could model the responses of the two components more accurately than they could the responses of the objective as a whole. Also the two components could be examined individually in the results of optimisation, allowing easy distinction between conditions that are too cold and too warm; this information would be lost if the objective was calculated as a single function.

Energy use

This second objective was rather more straightforward to calculate, as it was a simple sum of hourly energy loads for all active thermal control systems. Similarly to the thermal discomfort objective, it was broken down into two separate components that were then summed; radiant and convective loads:

$$F_2(\boldsymbol{D}) = E_r + E_c \quad (4.16)$$

The specific formulation of these components depended on the model under consideration, as the different case study models had different representations of thermal control systems. The test model was the simplest case, so this is taken as the base case with subsequent changes for case studies explained in chapter 5.

The radiant loads (E_r) were extracted directly from ESP-r, as it was found to be simplest to implement a radiant panel system with a thin-zone approach; a separate zone was defined to represent the radiant panel, and a control loop was specified to maintain the interior volume of the panel at the designated temperature by injecting heat flux into the thin-zone. The heat transfer co-coefficients within the thin-zone were edited to greatly elevated values to emulate the effective transfer of heat from the interior of the panel to its surface, and the radiator

temperature was defined by the design variable D_{sp_r}. The sum of the energy injected to maintain the temperature (in units of kW hrs) was taken as the value of the radiant component of energy use.

The convective component was calculated by modelling the HVAC heating coil as a simple heat exchanger:

$$E_c = \sum_{1}^{n} (\frac{Q}{\rho} C \Delta\theta) \tag{4.17}$$

Where E_c is the convective component of energy use (kWh), *Q* is the volume flow rate of the air (m^3/s), ρ is the density of the air (kg/m^3), C is the specific heat capacity of the air (kJ/kg °C), and Δθ is the change in air temperature required to reach the supply temperature from the external air temperature (°C). Fans were not included in the HVAC modelling for the test model. The supply temperature was defined by the design variable D_{sp_v}.

4.6.2 Design variables

There were 5 design variables programmed into T-BOT for the test model. These were:

- Thermal conductivity of the external wall (W/mK),
- Density (kg/m^3) and specific heat capacity (J/kgK) of the external wall,
- Radiator position,
- Radiator temperature (°C),
- HVAC supply temperature (°C).

These design variables were selected to characterise key parameters of the design and operation of the room with as few variables as possible, whilst also being relatively easy to implement as a design variable practically; it was relatively straightforward to vary these parameters in ESP-r models by script.

Thermal conductivity of external wall

This along with thickness determines the U-value of the wall, and hence largely governs heat loss to the environment through the external wall. It was taken as a design variable to control U-value, as thickness has wider structural implications that were not considered appropriate to optimise under the present remit. Limits of this design variable were:

$$0.1 \leq D_\lambda \leq 1.5 \tag{4.18}$$

Where D_λ is the design variable pertaining to thermal conductivity λ. These values were selected based on reasonable values for external constructions both with and without insulation.

Density and specific heat capacity of external wall

These material properties affect heat storage within the wall, and feature in the general DTM of ESP-r as described by eqn. 3.1. The two variables are controlled by the same design variable in order to reduce the number of design variables and hence the dimensionality of the design space, which implicitly improves the quality of the sample. This simplification may be made for two reasons; firstly it is always the product of these two variables that is used in thermal transfer equations such as eqn. 3.1, as this effectively converts specific heat capacity from units of mass to units of volume. It is therefore the product of the variables that the design variable should control. Secondly, the two variables with the units used here take comparable values, as can be seen in tables 4.1 and 4.3. The limits of this design variable were:

$$500 \leq D_{\rho,C_p} \leq 2000 \tag{4.19}$$

These values were chosen to represent conceivable values for both variables.

Radiator position

This was a discrete design variable with five possible values. These corresponded to four different radiator positions and a no-radiator case. Because the design variable is not numerically discrete (as with integers) but discrete in the sense of distinct cases (ie. apples, oranges, etc.), it is a special case of discrete variable termed a categorical variable. This effectively means that the variable cannot be mapped onto a continuous design space. As a result, instead of the metamodels taking account of variation in this design variable, separate metamodels are required for each value; each then becomes an independent optimisation problem. In this sense the variable does not conform to the definition of a design variable, however it is viewed as one in the present work simply for convenience and clarity in reporting.

Radiator temperature

This variable represents the internal temperature of the radiator, when present. The operation and control of the radiator is described in section 4.3.1. Clearly when a radiator is not present this variable becomes moot, and this was taken account of in the programming of T-BOT. The limits of this design variable were:

$$20 \leq D_{sp_r} \leq 40 \tag{4.20}$$

These values were chosen as safe operating limits for a conventional radiator in a hospital environment; it is stated in published Health and Safety Executive guidance [HSE, 2012] that "Where assessment identifies that vulnerable people may come into prolonged contact, [hot surfaces] should be designed or covered so that the maximum accessible surface temperature does not exceed 43 °C". Use of the radiator for cooling was not considered in the test model.

HVAC supply temperature

This represents the temperature at which air was supplied to the room by the HVAC system. The operation and control of the HVAC system is described in section 4.3.1. The limits of this design variable were:

$$15 \leq D_{sp_v} \leq 25 \tag{4.21}$$

These values were selected as typical operating limits for a conventional HVAC system.

4.6.3 Input variables

The optimisation method used was a genetic algorithm, as described in section 3.2. This takes a number of parameters as input variables; values for the most important of these are reported here:

- Maximum number of iterations (generations): 40 (it was found that an acceptable Pareto front had always been found by this point).
- Population size: Automatic (ie. allow the algorithm to control this variable based on how the optimisation is progressing).
- Chance of a solution mutating: 1%.
- Percentage of population considered "elite": 10%.

4.7 T-BOT Program Architecture

T-BOT represents the bespoke programs created during the course of this project to facilitate and automate the linkage between building simulation and optimisation software in order to accomplish BTO. A brief description of the basic steps involved is presented in section 4.2; this section is a detailed description of the procedural methodology implemented in the code. A schematic of the core programattic elements of T-BOT is shown in Fig. 4.18.

At the top level, the program is controlled by MATLAB scripts. All user interfacing and data entry is handled by the MATLAB controllers, though currently it is a relatively simple text based prompt-and-response interface. Most of the general process automation is also handled by

MATLAB scripts, mainly due to the simple and intuitive high-level parallelisation capability provided by MATLAB. Following this, there is a layer of Python meta-scripts which take input from the MATLAB scripts, then write and run customised shell-scripts which automate and fully replace the normal user input that would be necessary to run ESP-r. Some of these Python scripts also directly manipulate ESP-r model data files, particularly design variable entry scripts. The results of the ESP-r simulations are then automatically extracted and saved in ASCII text files, ready to be read back into MATLAB for pre-processing into results that can then be passed to the optimisation software to construct the metamodels.

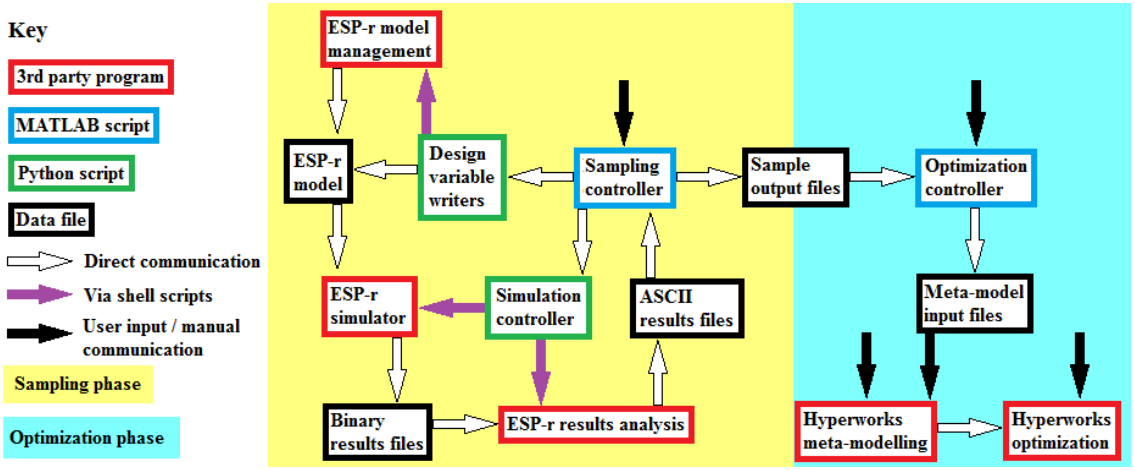

Figure 4.18: Schematic of programattic elements of T-BOT

In order to maintain some measure of flexibility for future developments, T-BOT was programmed in a modular fashion. For example, there is a separate Python script for transferring each different design variable to ESP-r. This is in contrast to more generic BTO programs such as GenOpt which tend to generalise the process of passing design variable values to just editing specific values within text files. Whilst this was necessary within T-BOT, this alone was not enough to fully accomplish design variable entry. With some design variables, it was necessary (or far simpler) to automate data entry through the interface as opposed to directly to the model data files. This has the disadvantage of making each design variable entry script specific to a model, but on the other hand the modular design means that scripts for new models can be modified from existing scripts and added into the system very easily.

The following is a step-by-step guide to the BTO procedure used by T-BOT.

Sampling phase

1. User input of parameters for sampling.

2. Create results directory and remove any leftover scratch files and libraries from previous runs.
3. If requested, generate sinusoidal climate data and import into ESP-r. If not, verify existence and applicability of existing climate data.
4. Take two independent latin hypercube samples from the n-dimensional design space where n is the number of continuous design variables; one of 100 samples, one of 50 samples. Note that if any categorical design variables are included in the problem, a different sample set must be simulated for each permutation thereof. For example, the radiator position variable had 5 possible values, so a sample set including this variable needed (100+50)*5 = 750 simulations.
5. Open pool of computation threads, usually 7 (8 processors per node, minus 1 for control process), in order to simulate samples in parallel.
6. Assign each simulation to a processing thread. This is done automatically but essentially randomly, working on a “first come first served” basis; when a thread becomes free it picks up the next simulation in a common job queue.
7. Start of parallel computations. Transfer design variable values for the sample to the ESP-r model; each processing thread has its own identical version of the ESP-r model to work on. This is accomplished in two different ways depending on the variable; direct manipulation of model data files or automation of data entry through the interface.
8. Run simulation. Note that due to constraints within ESP-r, mean radiant temperature can only be evaluated at 4 locations at a time, so in cases of more than 4 MRT sensors this entails multiple simulations. CFD is disabled for runs after the first; see section 4.3.2 for a detailed discussion of this.
9. Return ESP-r model to its default state i.e. undo all changes performed in step 7.
10. Retrieve mean radiant temperature results from simulation(s). Find the closest CFD node to each specified comfort evaluation location, and combine local air velocity, air temperature and mean radiant temperature to calculate operative temperature as in equations 4.10-4.11.
11. If necessary, retrieve/calculate values of energy use components.
12. To comply with the OSMO approach no further calculations are performed at this stage. Store data values at all time-steps and locations in ASCII text file in results directory.
13. Write out log file with additional information needed when sample results are retrieved and processed for metamodelling. This also serves as a reference guide for files within the results directory.

14. End of sampling phase. Perform general clean-up; close parallel pool, remove scratch files and libraries, etc.

Once this sample data is obtained and saved, because of the OSMO approach it can be used for many different optimisations. See section 4.4.2 for a discussion of this.

Optimisation phase

15. User input of desired results set and objective function formulations.
16. Read in data from log file, this provides all necessary information for reading in sample data.
17. Read in operative temperatures and energy loads.
18. Using secondary parameters input at step 15, calculate values of selected objective functions at selected location at all time-steps, for all samples.
19. Calculate time-averaged (or summed) values of objective functions over selected time period(s).
20. Linearly scale values of design variables to values of 1-11. This provides axial isotropy in the design space and marginally reduces the numerical complexity of the metamodels.
21. Format data into matrices and export to .csv file. This is to transfer the data easily from the MATLAB sampler on the HPC cluster running a Linux OS, to Altair HyperStudy on a desktop computer running a Windows OS.
22. Transfer data from sampling platform to optimising platform. From this step onwards the process is mostly not automated, as cross-OS automation would have required more development of the author's programming skills than was practical in the time available. However, as all necessary data was placed in a single file in step 21 this process was found not to be unduly impractical.
23. Use pre-programmed macro to separate out data into individual datasets and save in text files (necessary for use in HyperStudy).
24. Set up HyperStudy project; input design variables and required responses, link to sample data in text file, and create DOEs from data.
25. Create metamodels from DOEs; initial fitting is performed using the 100 sample set. The smaller 50 sample set is used to validate the initial fit and to select an appropriate value of closeness-of-fit parameter, and it is then added to the fitting data to further refine the fit.
26. Inspect metamodelled responses. Check for over-fitting and oscillatory or over-shooting behaviour; if this is evident close to the optimum region then it may be necessary to reduce closeness-of-fit parameter to smoothen responses. However this

must be done with care as it usually results in greater error; inspect residuals in tandem with responses to find an appropriate trade-off.

27. Set up optimisation problem in HyperStudy and run. This generally takes no more than 5-10 minutes depending on the nature of the problem; with problems with clearer optima (usually smaller time periods) the MOGA algorithm tends to identify the Pareto front more quickly.

Chapter 5: Proof of Concept

5.1 Chapter Overview

In order to examine the fitness-for-purpose of T-BOT and the usefulness of the outputs, a brief proof of concept study was performed with the test model and methodology detailed in chapter 4. In this chapter some results of this study are presented and discussed. Section 5.2 summarises conditions and intents of the study (for more detailed information the reader is referred to Chapter 4), section 5.3 presents results, and section 5.4 discusses conclusions of the study.

A portion of this work was presented in a conference paper written by the present author, with support from the project supervisors [Cowie et al., 2013]. Note that some of the results presented in this paper were subsequently superseded by results gathered with improved methodology; the results presented here are the most recent and hence may not match synonymous results in the paper.

5.2 Description and Rationale

Preliminary studies were run with the test model to investigate model performance. The problems were kept simple so that the performance of T-BOT could be examined with clarity. The optimisation parameters were as described in section 4.6. Model variables fell into two categories, design variables and contextual variables:

- Design variables:
    - Radiator temperature (continuous variable, range 20-40 °C)
    - Radiator position (discrete variable, 5 possible values [NW face, NE face, SW face, SE face, NO radiator])
    - Ventilation supply temperature (continuous variable, range 15-25 °C)
- Contextual variables:
    - Evaluation location (ie. position in the room where comfort is evaluated)
    - Optimisation period (ie. number of hours that objective functions are summed/averaged over)

Whereas design variables were varied implicitly in the optimisation process, contextual variables were independent for each case and are reported with the results below.

These preliminary investigations had three distinct goals:

1. To investigate if spatial variation does affect optimum conditions and if so to what extent,
2. To investigate whether optimising over shorter time periods provides useful information,
3. To qualitatively validate the fitness-for-purpose of T-BOT with the results.

Each of these goals was examined in turn.

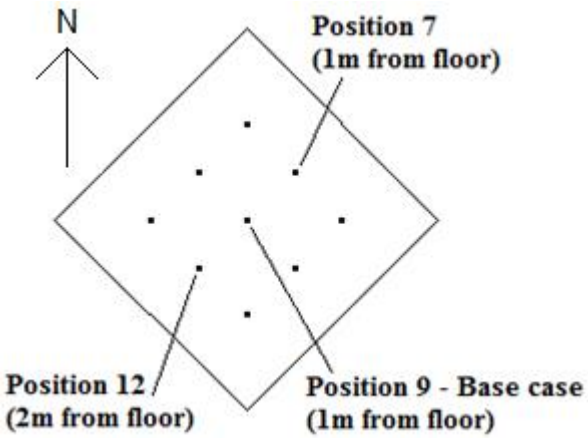

Figure 5.1: Showing positions of comfort evaluation points in the test model for the proof of concept study.

Fig. 5.1 shows the positions of comfort evaluation points used for this study; this is shown as a plan view, and relates to the model geometry as shown in Fig. 4.2.

5.3 Results

5.3.1 Base case

In order to have some benchmark, base case results were first obtained for comparison with subsequent data. Optimisation results are presented in terms of Pareto fronts. The theory behind these is given is given in section 2.4.3. The thermal discomfort criterion is on the horizontal X axis, with higher value signifying greater thermal discomfort (i.e. less comfortable conditions). Energy use is on the vertical Y axis. Lower values of both are more desirable. Ranges of both axes were scaled in all cases (both here and in Chapter 7) such that the shape of the Pareto front is clearly evident. *It is critical that the reader take account of the axial ranges, particularly when comparing such figures to one another*.

Though the Pareto fronts generally take the form of relatively smooth curves, they are made up of many individual points, each of which represents a potential design. The solution which achieves 0 discomfort is the highest energy solution, and conversely the point with the lowest energy use has the highest discomfort. Such negative correlation is a common pattern in all results. The Pareto front can therefore be considered as the optimum trade-off between the two objectives; each point represents a solution that is optimal for a given value of either objective.

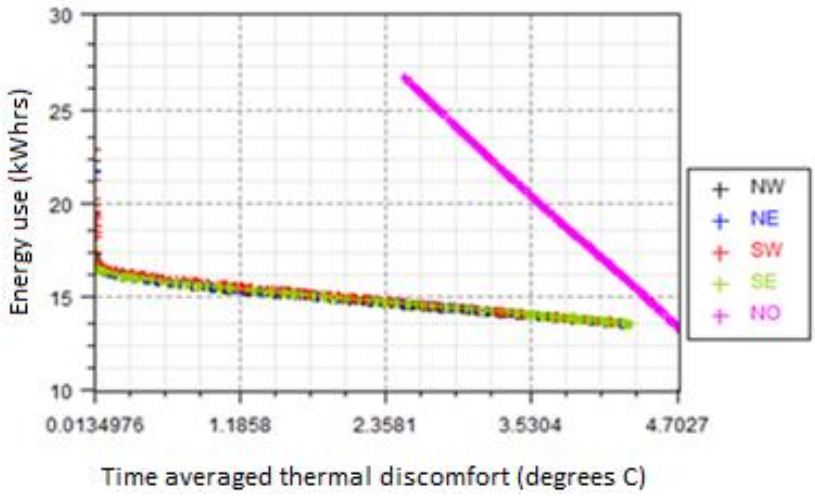

Figure 5.2: Pareto fronts at all radiator positions, for the winter case, at a position in the centre of the room, for the full 24 hour period. Legend refers to radiator position cases.

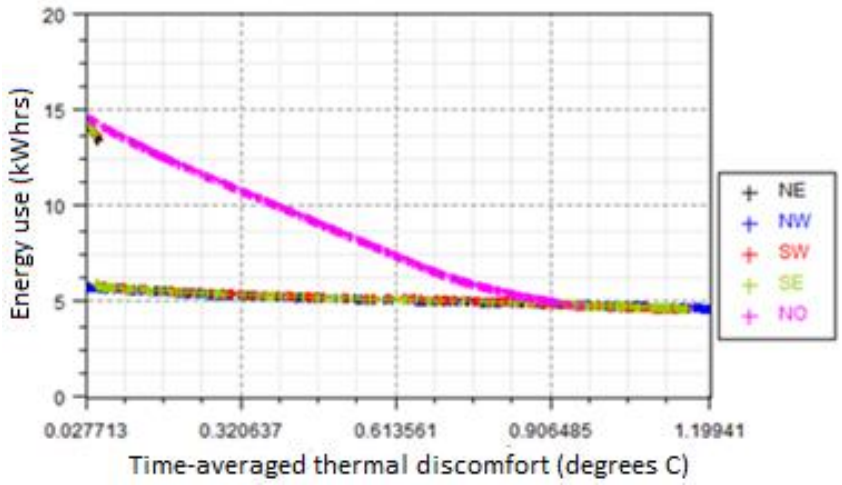

Figure 5.3: Pareto fronts at all radiator positions, for the mid-season case, at a position in the centre of the room, for the full 24 hour period. Legend refers to radiator position cases.

These results were obtained at position 9 as shown in Fig. 5.1; a position in the centre of the room, approximately 1m from the floor. Optimisations where performed over a 24 hour period. Comfort criteria were taken from recommendations given in CIBSE Guide A [CIBSE,

2006] as given in section 4.6.1, and sinusoidal climate data was used as described in section 4.3.1. Figures 5.2 and 5.3 show base case results for the winter and mid-season cases; the summer case was found not to be informative as the radiator was not configured to provide cooling or free-run, making all but the no-radiator case inapplicable.

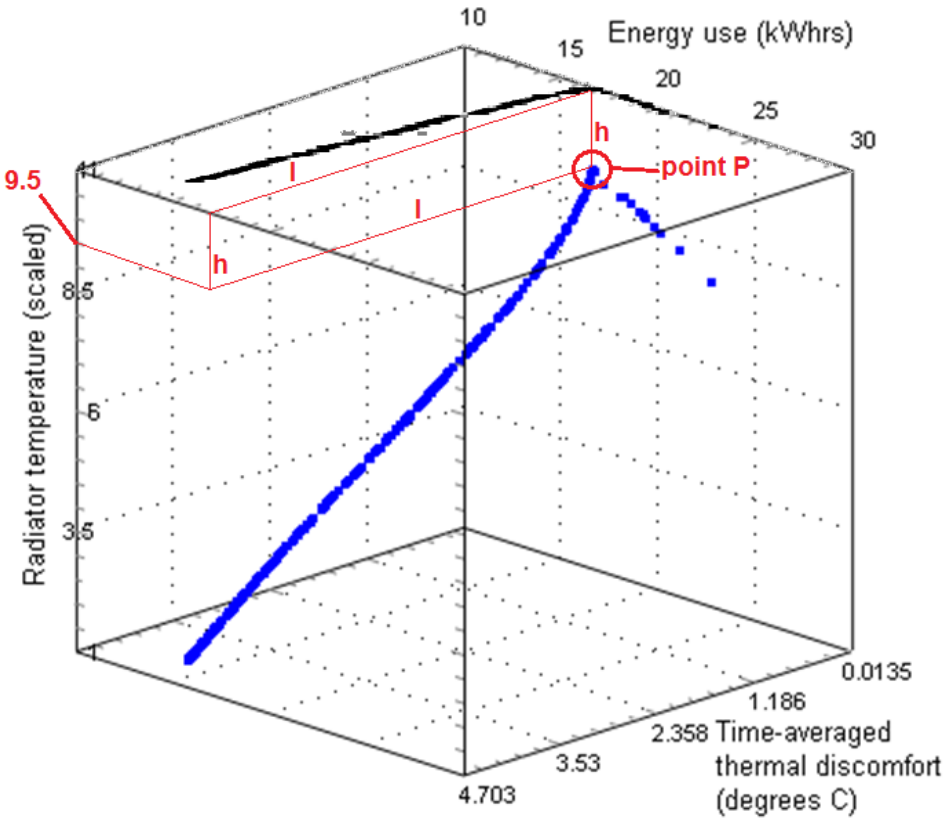

Figure 5.4: Pareto front plotted against a design variable with the front projected onto the top face, showing guidance on interpreting such 3-D plots.

It can be seen in Figs. 5.2 and 5.3 that all four radiator position cases have approximately identical Pareto fronts, as would be expected in the centre of the room. The no-radiator case for winter conditions cannot achieve acceptable comfort deviation due to the necessity for the extra heating that the radiator provides. On the other hand for the mid-season case the no-radiator case can achieve acceptable comfort, but the cases with the radiator can achieve similar levels of comfort at a lower energy requirement. Figs. 5.5 and 5.6 show fronts for the SW radiator position, for both climate cases, plotted against radiator temperature (a) and ventilation temperature (b) to examine the positions of the fronts in relation to the design variables, for the winter and mid-season cases respectively. Note that for the purposes of optimisation all design variables were scaled to values of 1-11 as detailed in section 4.6.2; actual limits of the design variables can be found in section 5.1. Fig. 5.4 shows guidance on interpreting these 3-D graphs; the Pareto front for the SW radiator position as shown in Fig.

5.2 is projected onto the top face of the cube, and it can hence be seen that the vertical axis value of example point P is 9.5.

It can be surmised from Figs. 5.5 and 5.6 that the energy requirement of the ventilation system is far greater than that of the radiator; this is clearly evident from the fact that the ventilation temperature is forced to minimum for all but the solutions closest to 0 °C discomfort, whereas radiator temperature is a smooth curve increasing with decreasing thermal discomfort. This also implies that fully optimal thermal comfort cannot be achieved using just the radiator; for solutions to obtain discomfort closer to 0 °C the ventilation system is required to supply warmer air also, though this increases energy use substantially.

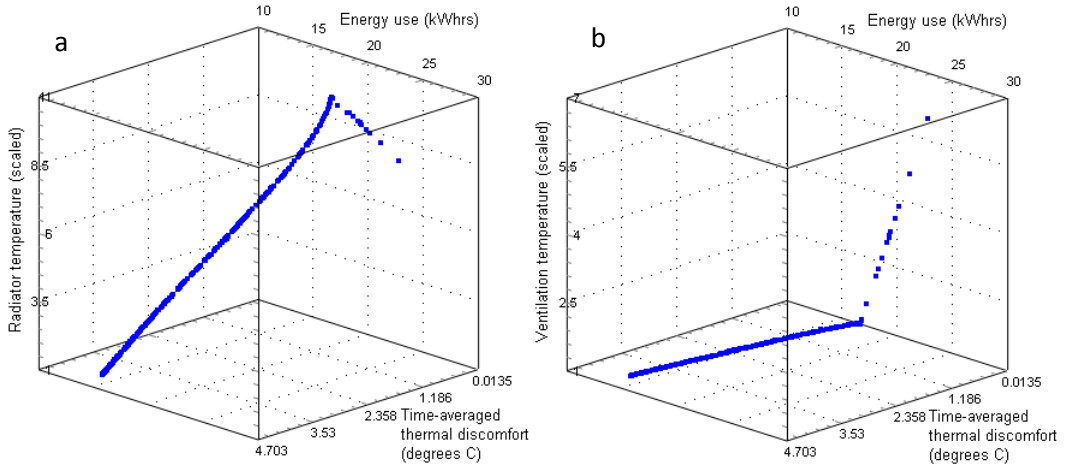

Figure 5.5: Pareto front for the winter case, SW radiator position, plotted against radiator temperature (a) and ventilation temperature (b) design variables.

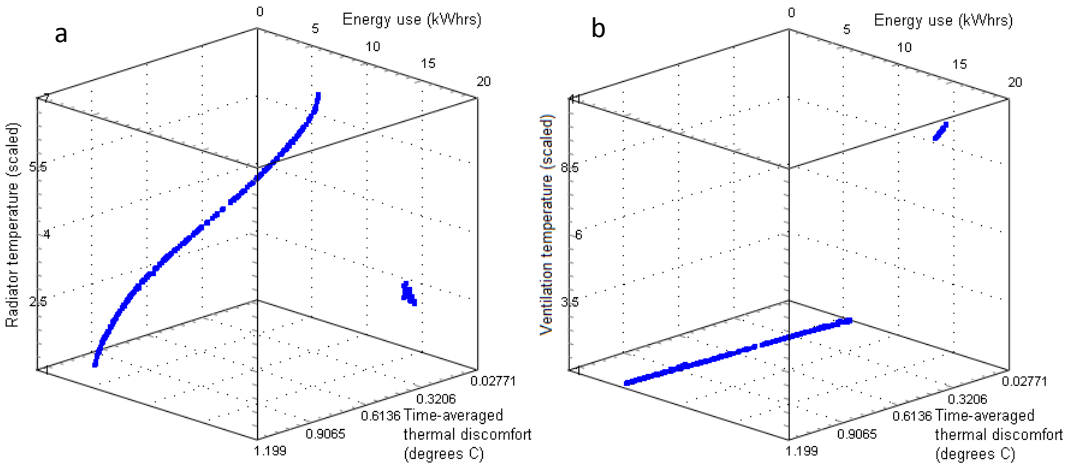

Figure 5.6: Pareto front for the mid-season case, SW radiator position, plotted against radiator temperature (a) and ventilation temperature (b) design variables.

The minimum ventilation temperature of 1 has an actual value of 15 °C, so in winter the air will always be being heated, whereas for the mid-season case during the warmest part of the day the air will actually need to be cooled to reach 15 °C.

It is interesting to note that in the mid-season case, radiator temperature does not reach its maximum before solutions start using the ventilation system to heat the room. This suggests that above a radiator temperature of approximately 6 (actual value 30 °C) the radiator temperature cannot be raised further without overheating the room and increasing thermal discomfort as a result. This further suggests that using the ventilation system to provide heating has the property of reducing temperature swing in the room (ie. reducing the amplitude of peak temperatures), allowing solutions to come closer to 0 °C discomfort than is possible using only the radiator.

5.3.2 Spatial variation

To examine the effects of spatial variation on the results, optimisations were performed for a winter case with the radiator on the NE face of the room opposite the window, at a further two locations; position 7 close to the radiator and position 12 close to the window (as shown in Fig. 5.1). Fig. 5.7 shows the Pareto fronts at all three locations with the solutions with ventilation temperature greater than minimum removed in order to show spatial variation more clearly. It can be seen that, as would be expected it generally takes less energy to maintain the same level of comfort close to the radiator, and more energy close to the window. This is a result of the differences in local operative temperature at these locations. However when the most optimal solution was picked from each Pareto front using the criteria of lowest discomfort whilst maintaining HVAC supply at minimum, only a 2.7% variation was seen in optimum radiator temperature between the positions closest to the radiator and window.

For the purposes of comparison, optimisations were also run for an identical set of cases, except that the radiator was on the SW face under the window. Fig. 5.8 shows the Pareto fronts at positions 7, 9 and 12 (as shown in Fig. 5.1) for this case. Comparing this to Fig. 5.7, it can be seen that there is generally less spatial variation, as would be expected given that the radiator and window are on the same wall in this case, as opposed to on opposite walls. It is also interesting to note that generally the fronts all exists at a slightly higher energy use. This is likely due to that fact that when the radiator is on the external wall, heat loss through the building fabric is explicitly modelled and will result in a small amount of extra energy loss from the radiator.

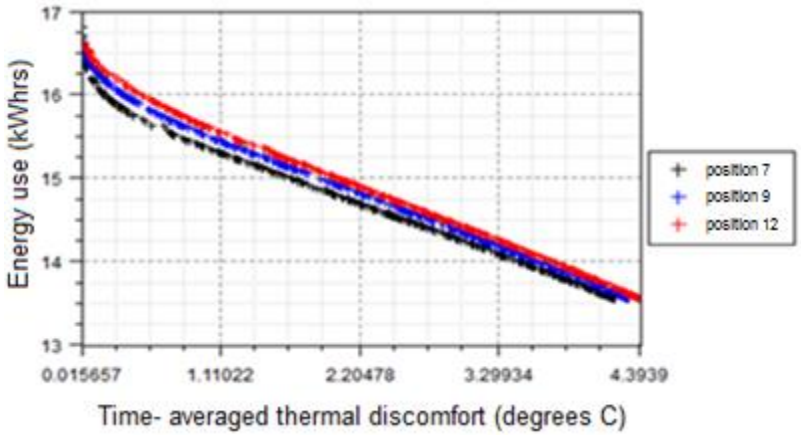

Figure 5.7: Pareto fronts at 3 different positions, for the winter case, with the radiator on the NE face, for the full 24 hour period.

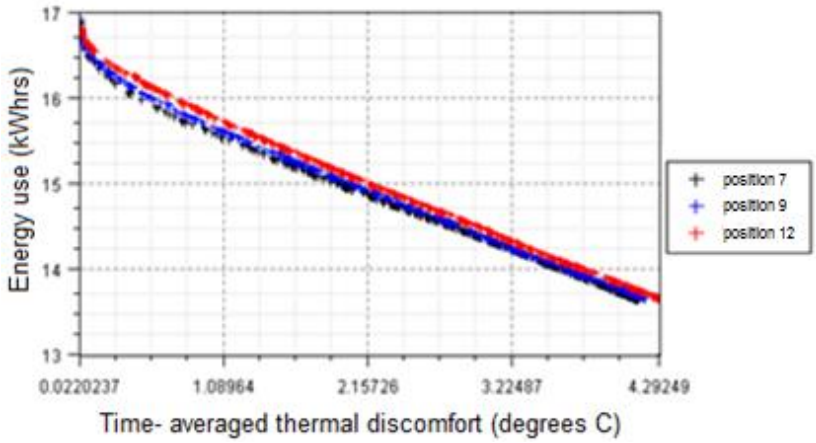

Figure 5.8: Pareto fronts at 3 different positions, for the winter case, with the radiator on the SW face, for the full 24 hour period.

Whilst the difference between the various positions was small, the results showed that spatial variation does affect optimum conditions, but this variation appears to peter out when approaching zero thermal discomfort. It was also evident in the test model results that there was very little spatial variation in air temperature, the vast majority of the variation came from the radiant environment.

5.3.3 Variation of time period

Optimisations were performed separating the 24 hour period into two different sets of periods; 6 hour periods and 3 hour periods. These were evaluated at position 9 the centre of the room for the winter case, with the radiator on the SW face of the room. Pareto fronts for the 6 hour periods are shown in Fig. 5.9; it can be seen that the energy requirement for comfort is much higher during the night-time periods, as would be expected.

The radiator temperatures for the optimum solutions (using the criteria of 0 °C discomfort) of these periods are shown by the blue bars in figure 5.9; all optimum solutions maintained ventilation temperatures at minimum. Clearly, the only deviation occurs during the 13-18 hour period. This can be attributed to solar gains, as the window is on the SW face and so the room will only experience directly impinging solar radiation principally during this period. However, it is interesting that despite the significantly lower radiator temperature during this period, the total energy use is not reduced substantially in comparison to the 7-12 hour period. This is due to the slight offset in the sinusoidal temperature data; peak temperature occurs at 12'noon, which means that the 13-18 hour period will implicitly have a slightly higher convective energy load than the 7-12 hour period. It is merely coincidence that this difference effectively cancels out the energy reduction of the lower radiator temperature. The offset also explains the difference in energy use between the 1-6 and 19-24 hour periods.

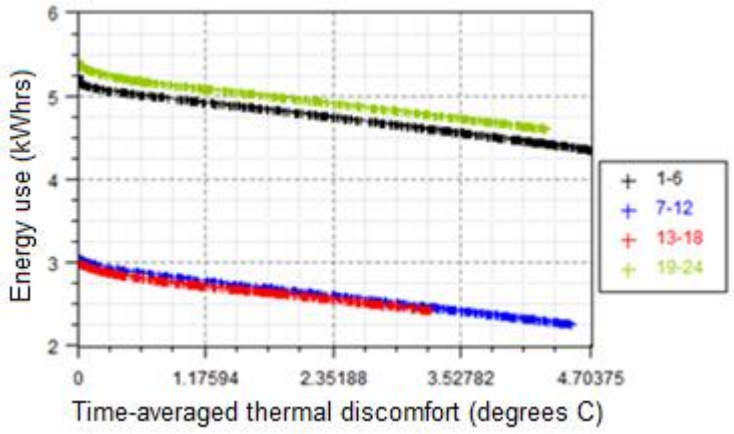

Figure 5.9: Pareto fronts for 6 hour periods, for the winter case, at a position in the centre of the room, with the radiator on the SW face.

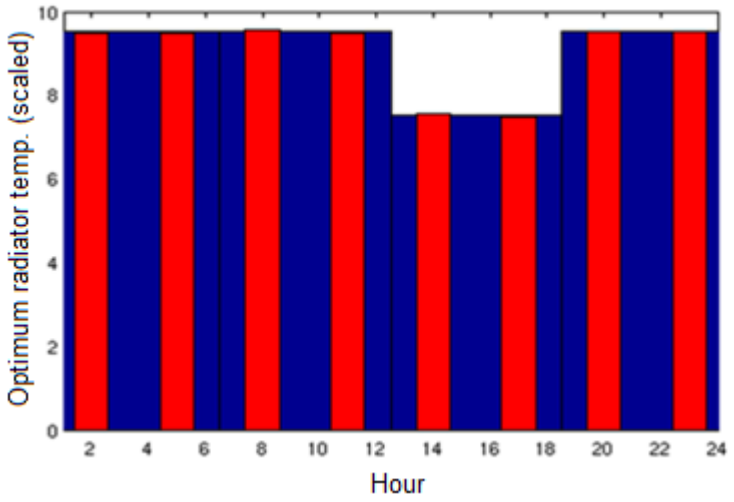

Figure 5.10: Optimum radiator temperatures for 6 hour periods (blue bars) and 3 hour periods (red bars)

Pareto fronts for the 3 hour periods were not plotted similarly to Fig. 5.9 because they were found to give no more information than the 6 hour periods. This is evident from Fig. 5.10; the red bars show optimum conditions for the 3 hour periods and it is clear that no additional information is gained from this additional temporal resolution. In Fig. 5.10, the lower values between 13:00 – 18:00 are approximately 7.5 which represent a radiator temperature of 35 °C, and the higher values elsewhere are approximately 9.6 which represent a radiator temperature of 39.2 °C.

However, the 6 hour optima do exhibit improvements over the 24 hour optima. The aggregate energy use for the optimum solutions for the 6 hour time periods was 16.65 kWh, and this achieved 0 °C discomfort for the full 24 hour period. However for the 24 hour solutions, 0.014 °c discomfort was achieved at an energy use of 22.77 kWh and 0.034 °C discomfort was achieved at 16.74 kWh energy use.

5.4 Summary

In general the study presented in this chapter suggested that T-BOT is broadly fit-for-purpose, able to successfully identify optimum solutions from the design space, insofar as the metamodelled design space captured optimum solutions. It is worth noting that the optimality of solutions in absolute terms (i.e. consideration of metamodel fidelity) was not considered here, and was left for a more comprehensive study presented in Chapter 7. However, a number of issues with the methodology were identified.

In the present study the summer results were found not to be informative, as the room was not able to achieve thermal comfort. This was likely due to the somewhat high values of the sinusoidal dry bulb temperature profile (15-25 °C) and the limited provision of cooling; the only means by which this could be achieved was by ventilation. However the minimum ventilation supply temperature was 15 °C, which is likely to be able to achieve sufficient cooling in practice, so this is a somewhat counter-intuitive outcome. A similar outcome was evident in the winter results; without a radiator the room was also unable to achieve thermal comfort as demonstrated in Fig. 5.2. It was concluded that the modelling of the ventilation system (i.e. mechanical supply on the ceiling and extract on the SE wall as shown in Fig. 4.2), along with the lack of other excitations of airflow in the room, likely resulted in “short-circuiting” of the ventilation; air mixing was not sufficient for the ventilation to provide sufficient temperature control. In response to this the CFD modelling in the subsequent case study model was

improved to model excitations due to the thermal plume of an occupant, as detailed in section 6.4.3, to induce better air mixing.

Furthermore, this highlighted a key factor of the methodology. T-BOT is only able to provide solutions within the confines of the model and the design variables. If the model is unable to satisfy conditions that T-BOT is trying to achieve, in this case thermal comfort, then T-BOT cannot give viable solutions, as these will lie outside the design space. It is critical that the building model has sufficient flexibility to achieve desired conditions within the design space, in order for the results to be useful. In response to this the model used for the later case study was based on an actual building, and the systems therein were modelled to have greater flexibility in operation as detailed in section 6.2.1.

The results presented in section 5.3.2 suggest that spatial variation does influence the results of BTO, though the magnitude of this influence in these results is small. The improved CFD modelling mentioned above, as well as improving air mixing due to the thermal plume, may also induce greater variation of air temperature due to the local characterisation of occupant casual gains; i.e. the convective component of occupant casual gains was explicitly injected into a specific set of CFD grid cells. In the present model, casual gains were injected by the building-side DTM domain, and hence the convective component thereof was evenly distributed in the room in accordance with the bulk air assumption as detailed in section 3.2.3.

On the other hand, results in section 5.3.3 suggest that variation in optimisation time period (synonymous with control periods due to constant set-points) can have a more significant effect on results. Results of four separate optimisations of 6 hours each, provided a clear benefit in terms of both thermal comfort and energy use, over a single optimisation of 24 hours. However, doubling the resolution to eight optimisations of 3 hours each did not provide any additional benefit over the 6 hour optimisations. This suggests that there is a limit to the benefit that can be achieved from shorter control periods in this case, and furthermore highlights a key ability of T-BOT; to explore control periods with relative ease. This is discussed in greater detail in section 8.2.2. Based on the results of this study, optimisation periods of 24 and 6 hours were considered in the subsequent case study.

Finally, it is worth noting that the conclusions from results of this study broadly conform to what would be expected based on perennial building performance wisdom. For example, achieving thermal comfort for the mid-season case requires significantly less energy than for the winter case. The conclusions from the 6 hour optimisation results also match what would be expected in terms of solar incidence as detailed in section 5.3.3. The comparison between different radiator positions shown is Figs. 5.2 and 5.3 exhibit consistent fronts for all cases with

radiators, even though the sample, metamodels and optimisations for each radiator position were entirely independent of one another. These observations provide some measure of qualitative validation of the methodology. T-BOT does not “know” what is reasonable and should be expected, it can only draw conclusions based on the simulation results. The fact that these conclusions generally conform to what is expected suggests that the methodology is able to broadly capture variation of conditions in a sensible manner.

Chapter 6: Case Study Model Development

6.1 Chapter Overview

This chapter details the development of the model representing the case study site, Bradford Royal Infirmary. First, the case study site is described and the individual location selected to be modelled are specified. Following this, the model is described and compared with experimental results gathered at the case study site. Also the selection of climate conditions for the model is discussed. Finally, the implementation of the model in T-BOT is detailed.

A portion of this work was presented in a conference paper written by the present author, with support from the project supervisors [Cowie et al., 2015]; specifically elements of sections 6.3 and 6.4.3.

6.2 Case Study Site – Bradford Royal Infirmary

6.2.1 General site description

Figure 6.1: Site map of Bradford Royal Infirmary, showing the boundary of the site and the location of the building selected for modelling [Google, 2014].

Bradford Royal Infirmary (BRI) is a large NHS teaching hospital, situated on a single site in the north-west of Bradford. Fig. 6.1 shows a photographic map of the site, highlighting the maternity ward. This was selected as it is a tower block with single hospital bedrooms on a middle floor; ideal for case study in the present work. Fig. 6.2 (taken from Lomas et al. [2012])

shows the dates of construction for the buildings on the site, and it can be seen that the maternity ward was built in the 1945-1967 period.

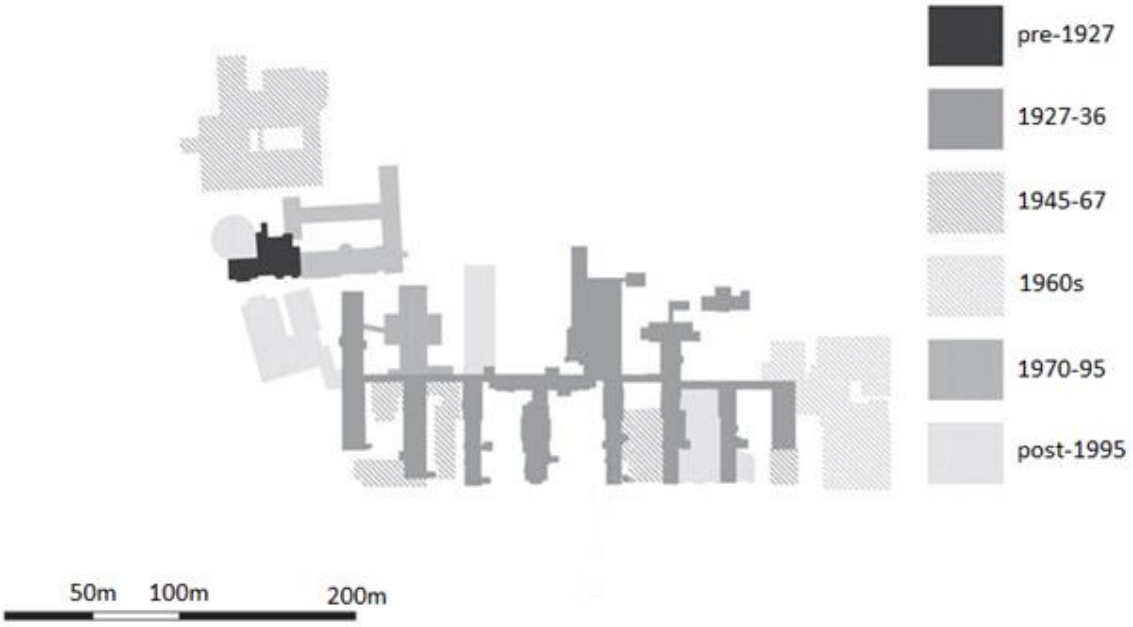

Figure 6.2: Showing eras of construction for the buildings of Bradford Royal Infirmary [Lomas et al., 2012].

6.2.2 Maternity ward description

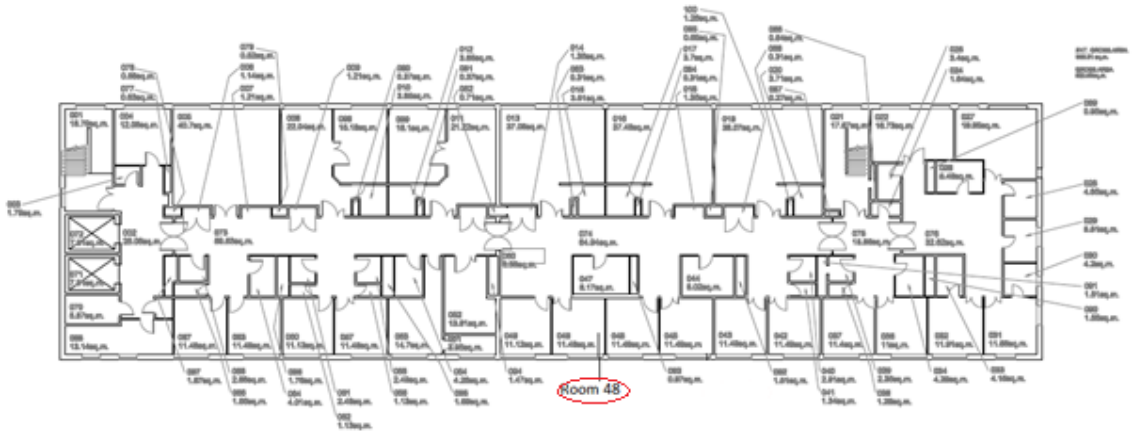

Figure 6.3: Schematic of the 2nd floor of the maternity tower block, highlighting modelled room.

This building consists of a large ground floor complex at the foot of a 5-storey tower block. The locations of interest for the present project were the single-occupancy bedrooms on the second floor of the tower block. Temperature monitoring data was available for room 48 on this floor as marked on Fig. 6.3 (the temperature monitoring data is discussed in detail in

section 6.2.4). The position of the external windows in this room is also shown on a photographic view of the building in Fig. 6.4. The model is described in section 6.3.

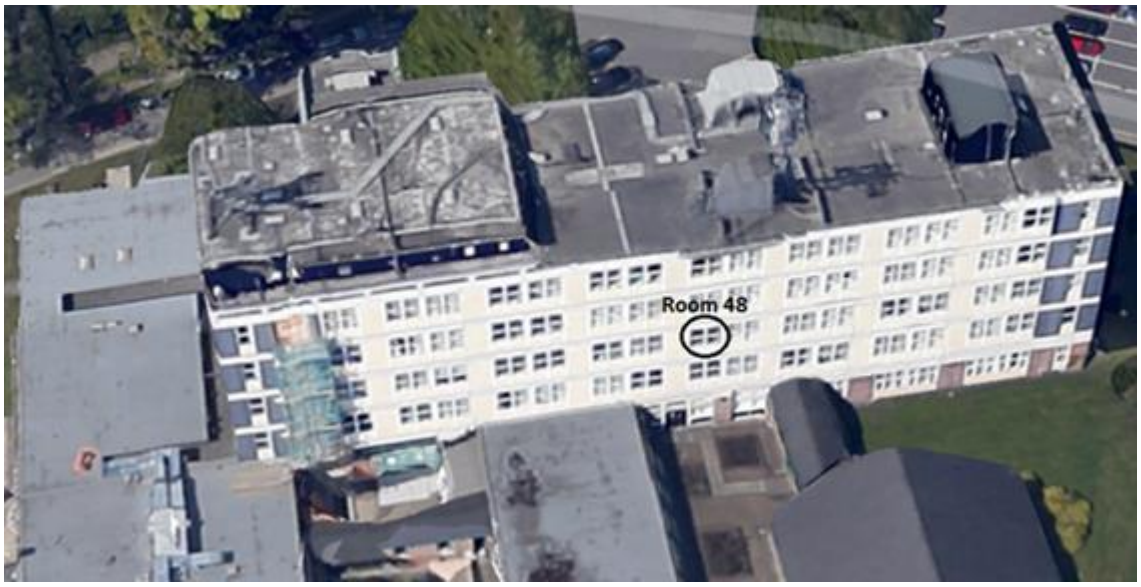

Figure 6.4: Photographic view of the maternity tower block, showing the positions of the external windows in the modelled room [Google, 2014].

Available information on the construction of this building was limited; only a rough description could be obtained. The tower block is a reinforced concrete construction, with insulated cladding panels fitted on the façade. Exact dimensions and material properties of the construction had to be assumed; these are given in section 6.3.2.

The building is mechanically ventilated by supply only in the axial corridor running the length of the building. Ventilation in the room is largely accomplished by operable windows in each, resulting in a mixed-mode system. This reliance on natural ventilation rather complicates the modelling process, as wind induced air flow through the windows becomes integral to the performance of the room and hence must be taken into account. This was accomplished in the case study model by introducing another modelling domain, an air flow network (AFN). AFNs are discussed in detail in section 6.3.3.

Instead of wall-mounted metal radiators, this building is equipped with heated ceilings. This was another factor in the selection of the buildings; a large radiant panel attached to the ceiling removes the need to consider different locations for radiators. In practice the system is generally only used for heating, as cooling may introduce issues such as condensation risk. However given that the summer results of the proof of concept study (detailed in Chapter 5) were found to be uninformative due to lack of provision for cooling, the ceiling panel in the case study model was configured to provide a wider range of temperatures to provide cooling

as well as heating. Whilst this is unrealistic in terms of the modelled space, the greater flexibility allowed evaluation of both summer and winter conditions with the same model.

As the name suggests, this ward caters to maternity-related patients. It was therefore assumed that the most likely residents of the ward bedrooms would be pregnant women and new-born children, with a fair chance of them being in ill-health in the case of prolonged residents. It was therefore decided to maintain the most stringent category of comfort criteria (ie. category I in Table 2.1).

6.2.3 Temperature monitoring data

The temperature monitoring data that was used in exploring the performance of the case study model was not gathered as part of the present project. The data was gathered as part of a much wider project termed DeDeRHECC (Design and Delivery of Robust Hospital Environments in a Changing Climate). The data for BRI that is used in the present project is only a sub-set of the large volume of data gathered across multiple sites. This data has been used in a number of studies under the remit of DeDeRHECC, including Lomas et al. [2012], Lomas and Giridharan [2012] and Short et al. [2012].

The data was gathered using a variety of Hobo U12 (+-0.35 °C), Hobo pendant (+-0.53 °C) and Tiny Tag (+-0.5 °) loggers [Lomas et al., 2012]; accuracy within the range 0 – 50 °C is shown in brackets. All loggers were calibrated prior to use, resulting in a difference between them of less than 0.2°C for spaces of the same temperature [Lomas et al., 2012]. The data provided took the form of hourly temperature readings in a variety of rooms. As well as the monitored room marked in Figs. 6.3 and 6.4, data was also available for the corridor. This proved useful for evaluating model performance, as described in section 6.3.4.

6.3 Case Study Model

6.3.1 Geometry

This model was intended to represent a single occupancy bedroom in the maternity tower block, marked in Figs. 6.3 and 6.4 as room 48. General geometry of the room was taken from schematic drawings such as that reproduced in Fig. 6.3. Fig. 6.5 shows a diagram of the model and its key features.

Whilst the model is set up to simulate the room only, the model includes three distinct zones. The first, labelled in Fig. 6.5 as “mat_sin1_rm”, is the zone representing the room itself. The

zone “frengC” represents the cavity of the radiant ceiling, and is necessary for the thin-zone approach of modelling the system, as described in section 4.3.1. The final zone labelled “corridor” is a notional boundary zone only, necessary in order to accommodate the AFN as described in section 6.3.3. The geometry, volume and contents of the corridor zone are entirely arbitrary, as no DTM calculations are needed for it. It is present only to provide ESP-r with an entity to represent the other side of air flow through the door. As such, it may be summarily ignored in examining the model.

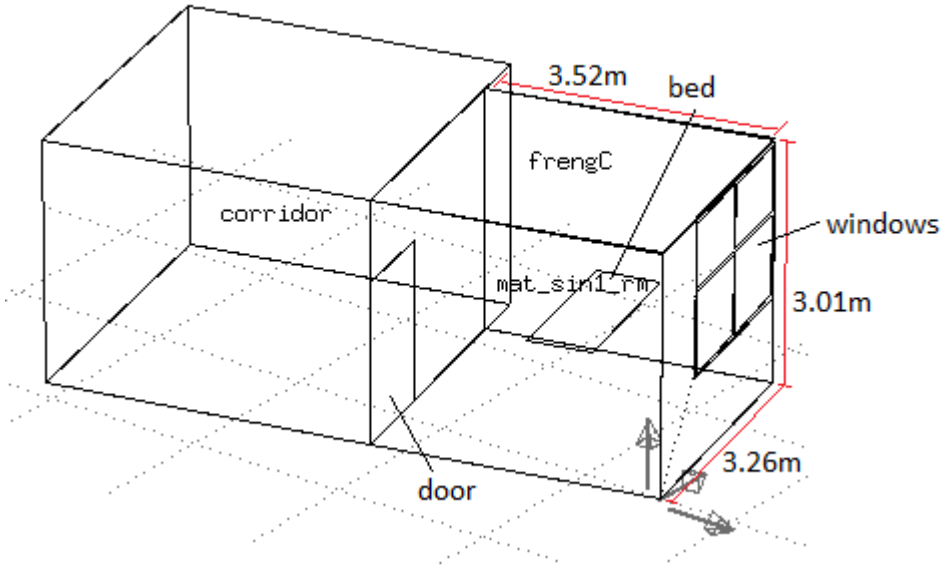

Figure 6.5: Geometry of the maternity single occupancy room model. Dimensions of the room and features of it are marked. The other labels refer to zones within the ESP-r model; mat_sin1_rm = maternity single occupancy room model, frengC = radiant ceiling.

Surface	**Boundary Condition**	**Face***	**Area**
Floor	Solid surface	Bottom	11.48
Ceiling	Solid surface	Top	11.48
South wall	Solid surface	South	6.11
East wall	Solid surface	East	10.60
North wall	Solid surface	North	8.56
West wall	Solid surface	West	10.60
Window	Solid surface	South	3.84
Window opening	Velocity	South	0.65
Door	Velocity	North	2.04
* The window is on the south face and the bed is attached to the east face.			

Table 6.1: CFD boundary conditions in the case study model.

Construction	Material	Thickness (mm)	Thermal Conduct. (W/mK)	Density (kg/m^3)	Specific Heat Capacity (J/kgK)	IR Emiss--ivity	Solar Emiss--ivity	U Value (W/m^{2}K)
Internal walls*	dense plaster	12.5	0.5	1300	1000	0.91	0.5	3.401
	cast concrete	100	1.35	2000	1000	-	-	
	dense plaster	12.5	0.5	1300	1000	0.91	0.5	
Ceiling**	carpet	6	0.06	186	1360	0.9	0.6	1.57
	chipboard	19	0.15	800	2093	-	-	
	air gap	50	n/a****	0	0	-	-	
	heavy mix concrete	140	1.4	2100	653	-	-	
	steel	4	50	7800	502	0.12	0.2	
Floor*	steel	4	50	7800	502	0.12	0.2	1.321
	heavy mix concrete	140	1.4	2100	653	-	-	
	air gap	50	n/a****	0	0	-	-	
	chipboard	19	0.15	800	2093	-	-	
	carpet	6	0.06	186	1360	0.9	0.6	
External wall***	approximation	326	0.265	1166	1166	0.9	0.5	0.714
Door*	wood	25	0.19	700	2390	0.9	0.65	3.316
Windows	glass	6	0.76	2710	837	0.83	0.05	2.811
	air gap	12	n/a****	n/a	n/a	-	-	
	glass	6	0.76	2710	837	0.83	0.05	
Window frame	grey aluminium	4	210	2700	880	0.82	0.72	0.461
	glass fibre quilting	80	0.04	12	840	-	-	
	grey aluminium	4	210	2700	880	0.82	0.72	
Radiant panel*****	white aluminium	4	210	2700	880	0.82	0.32	n/a

* These surfaces were considered semi-adiabatic (see below) and hence the properties given here are largely moot.

** The radiant panel was modelled as being attached to the soffit of this construction; ie. the soffit of the panel forms the actual ceiling of the room.

*** This surface was an approximation of an assumed construction, see below for details.

**** All air gaps assumed to have R value = 0.17 m^{2}K/W

***** This was modelled as a distinct zone. Front face and edges were as shown above, back face was modelled as panel -> ceiling, see ** comment.

Table 6.2: Material properties of surfaces in the case study model.

As can be seen in Fig. 6.5, the model includes a 4-pane window with an explicitly modelled frame. The bottom two panes of this window were considered operable, and all panes were assumed to be double glazed. The bed was modelled essentially as an obstruction to airflow, but was also considered a surface in the radiant heat exchange calculations for the room. It also serves to give a clear indication of where the occupants of the room are likely to be, ie. lying on it or standing beside it. The model also includes a door, through which air may flow between the room and the corridor. Of course, this is dependent on how far open the door is

at the time, and this was taken into account in the modelling. Boundary conditions for the model are shown in table 6.1.

6.3.2 Constructions, casual gains and controls

Construction details for all surfaces are given in Table 6.2. The external wall construction was an approximation based on a typical concrete building, taking into account the limited information available on the building in question. Properties of this assumed construction are shown in Table 6.3. As with the test model, all surfaces other than the external face were considered to have conditions on the other side identical to those in the room. This effectively results in semi-adiabatic behaviour; the only difference being that heat storage within these constructions is modelled.The casual gains in the model are given in Table 6.4. These were constant, and based on the assumptions of one occupant in the room, and some electrical equipment in operation at all times such as a television or medical equipment. Lighting gains were based on typical values recommended in the help text of ESP-r.

Material	Thickness (mm)	Conductivity (W/mK)	Density (kg/m^3)	Specific Heat Capacity (J/kgK)	IR Emissivity	Solar Emissivity	U Value (W/m^{2}K)
plasterboard	13	0.21	900	1000	0.91	0.7	
fibreboard insulation	50	0.06	300	1000	-	-	
air gap	50	n/a*	n/a	n/a	-	-	0.71
cast concrete	200	1.35	2000	1000	-	-	
dense plaster	12.5	0.5	1300	1000	0.91	0.5	

* Air gap assumed to have R value = 0.17 m^{2}K/W

Table 6.3: Material properties of the assumed construction of the maternity tower block, representing a concrete-framed building with insulated cladding panels.

Source	Sensible (W)	Latent (W)	Radiant/Convective
Occupants	75	20	0.5/0.5
Lights	40	0	0.3/0.7
Equipment	50	0	0.5/0.5

Table 6.4: Casual gains assumed in the case study model.

There were several controls necessary to consider in the case study model; control of the window opening area, control of the radiant panel and control of the corridor temperature.

The radiant panel and corridor temperatures were maintained at constant values by injecting heat flux into the zone air-point at each time-step. For the window opening area, a far more complex control was used to attempt to encapsulate human behaviour.

Rather than a constant opening area, a multi-stage control algorithm was used with 4 distinct control regions. This necessitated 3 temperature set-points and 4 specified opening areas, one for each control region, as illustrated in Fig. 6.6. Opening areas in each control region were calculated by a nominal opening area multiplied by ratios:

- In the low control region, opening area = nominal area * ratio 1
- In the mid-low control region, opening area = nominal area
- In the mid-high control region, opening area = nominal area * ratio 2
- In the high control region, opening area = nominal area * ratio 3

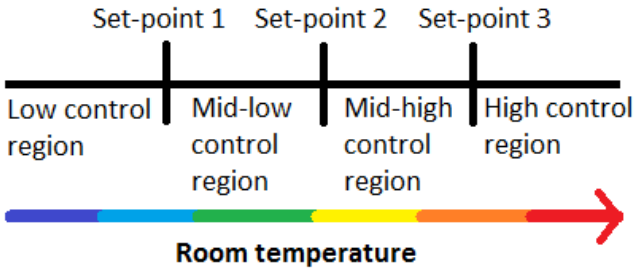

Figure 6.6: Illustrating the window control algorithm used for the case study model.

The benefit of this control was that it allowed the model some functionality to encapsulate human adaptive actions; if conditions in the room got too warm, the window would likely be opened (or opened wider). If conditions in the room got too cold, the window would likely be closed.

6.3.3 Additional ventilation modelling

Node	**Representing**	**Type**	**Temperature**
1	Outside	Boundary, wind-driven pressure	Climate
2	Room	Internal, unknown pressure	Unknown
3	Corridor	Internal, unknown pressure	Controlled
4	Notional boundary	Internal, fixed pressure	Equal to node 3
Component	**Representing**	**Type**	
1	Window opening	Controlled area	
2	Door opening	Fixed area and controlled discharge factor	
3	Pressure venting	Fixed area	

Connection	From node	To node	Via component
1	3 (outside)	1 (room)	1 (window opening)
2	2 (corridor)	1 (room)	2 (door opening)
3	2 (corridor)	4 (boundary)	3 (pressure venting)

Table 6.5: Nodes, components and connections setup for the AFN used in the case study model.

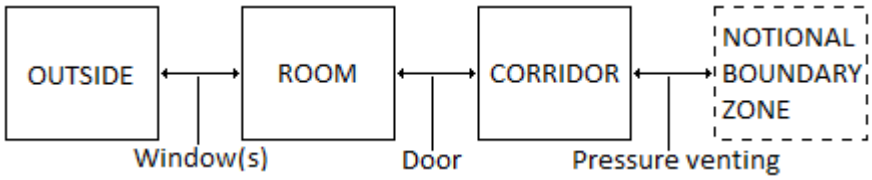

Figure 6.7: Diagram of the AFN described in Table 6.5.

In the test model, the CFD was directly conflated to the DTM with constant flow rate in- and out- lets for the HVAC system. However for the case study model, the room was not individually mechanically ventilated. This meant that variable flows through windows and doors had to be considered instead. This introduced further complexity in the ventilation modelling, as wind pressures and inter-zonal pressure differences are then required to calculate flow rates. In practice, within the context of ESP-r models this can be achieved by adding a further modelling domain. This is generally termed an air flow network (AFN); essentially it serves to simulate inter-zonal air flow whereas the CFD calculates intra-zonal air flow. The AFN was then conflated to the CFD domain, providing dynamic flow rates for openings on a time-step basis.

In ESP-r an AFN consists of nodes, representing boundary conditions and zones; components which represent flow openings; and connections which govern how the nodes are connected to one another via the components. The AFN for the case study model consisted of wind-driven flow through one or more operable windows, and flow through a door to a notional corridor zone. The setup of nodes, components and connections that was used to simulate this in the case study model is specified in Table 6.6 and illustrated in Fig. 6.7.

The extra notional boundary zone (node 4 in Table 6.7) was necessary to allow pressure venting from the corridor zone. This was needed because in the actual buildings, the corridors were connected to many other zones than the room under consideration. To simulate this entire system explicitly would require a far more complex AFN and hence increase the computation time. As such, it was decided to simplify the modelling of the corridor zone; to represent any other connections that exist between the corridor and other zones or boundary pressures, a single extra node was defined as a fixed neutral boundary pressure (relative

pressure = 0) with temperature forced equal to the corridor zone at each time-step. This is tantamount to assuming that given the typically large volume of corridor spaces and the large amount of openings likely to exist, the various flows in and out of the corridor though HVAC in/out-lets, window openings, and connections to other areas of the building cancel one another out to a resulting relative pressure of approximately 0. This assumption ensures that a significantly positive or negative pressure in the room zone, arising from wind pressure for example, stimulates air flow between the room and corridor; the boundary node allows the corridor zone to vent surplus or deficit air necessary to accommodate this inter-zonal flow. As the temperature of the corridor zone was controlled by a design variable in the case studies, it was convenient to force the temperature of any such make-up air into the corridor zone to the ambient temperature of the corridor.

6.3.4 Model performance

To ensure model performance was broadly representative of the actual building, the simulated temperature profile of the room was compared to measured data as detailed in section 6.2.3. For the purposes of these simulations, the measured temperature data for the corridor was imposed directly on the corridor zone. Information on door discharge factor and radiant panel temperature were not available, so these values were calibrated to achieve the best fit to the measured data; a door discharge factor of 0.8 and radiant panel temperature of 25 °C were found to be most appropriate. The simulations were performed for November and December in 2010, a period of 61 days. This was chosen as climate monitoring data was also available for the site during this period, also gathered under the remit of the DeDeRHECC project. The CFD was de-activated for the purposes of these simulations, as the large time period of 2 months would have made the computation time unfeasible were CFD included.

Whilst the climate monitoring data for the site did include several metrics of solar radiation, they were gathered from two distinct sources and were found to be somewhat inconsistent between themselves, as well as being incompatible with the metrics required by ESP-r. As such, to construct viable climate data for the simulated period, solar radiation data was taken from CIBSE climate data from a weather station in Bingley, approximately 12 miles from the case study site. This data was in the form of direct normal and global diffuse solar radiation, which are the metrics required for ESP-r climate data files. Other climate metrics including dry bulb temp and humidity for the monitoring data at the site and the CIBSE data for Bingley were compared, and found to be in good agreement, so it was assumed that the CIBSE solar data was sufficiently representative of the climate at the site.

Data on the occupation and use of the room during the simulation period was not available, so two separate simulations were performed; one with assumed casual gains that would be associated with general usage of the room, and one without. It was assumed that the actual temperature profile of the room would switch between these two profiles, though it was not known when exactly these switches would take place.

A comparison of the measured data and the simulation results is shown in Fig. 6.8. It can be seen that the model predictions generally follow the general trends of the measured data. The main differences between the measured and predicted results are the high and low peaks in the measured data, which are generally not captured in the predicted results. These are likely a result of further unknowns in the operation of the room; for example the use of space heaters is not uncommon in this building during winter, which may account for the spikes of higher temperature, and the window might be open at times which may account for the periods of low temperature. Neither of these factors were included in the modelling, as no information was available on their operation in practice. Furthermore, it is unlikely that both the radiant panel temperature and the amount the door was open were constant throughout the entire 2 month period, and variation of these parameters was also not taken into account in the modelling for the same reason.

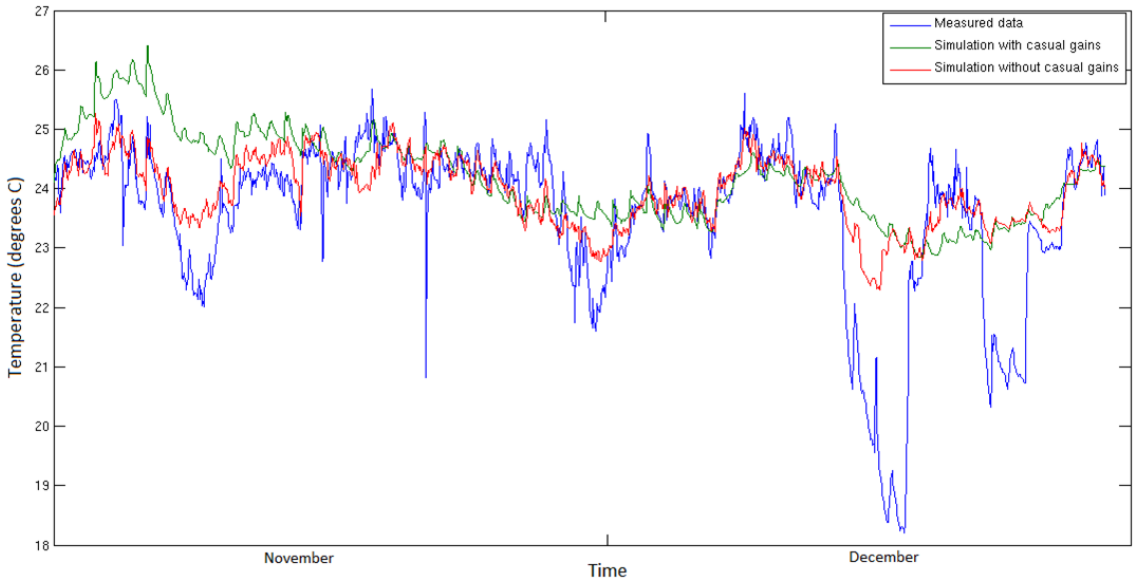

Figure 6.8: Showing a comparison of measured temperature data in the room and simulated results.

6.3.5 Climate conditions for case studies

Two distinct climate cases were selected to be analysed for both summer and winter conditions; extreme conditions and average conditions. This resulted in four broad categories of climate cases:

- Extreme summer conditions,
- Average summer conditions,
- Extreme winter conditions,
- Average winter conditions.

No viable climate data was gathered at the case study site during the summer period, so in lieu of this data the 2010 CIBSE data from the aforementioned weather station in Bingley was used. For winter data, the climate data gathered at the site in November and December was assumed to be representative. From analysis of the 2010 Bingley weather data, December was found to have the lowest average dry bulb temperature and direct normal solar radiation, whilst November had the 4th lowest average temperature and the 3rd lowest direct normal solar radiation.

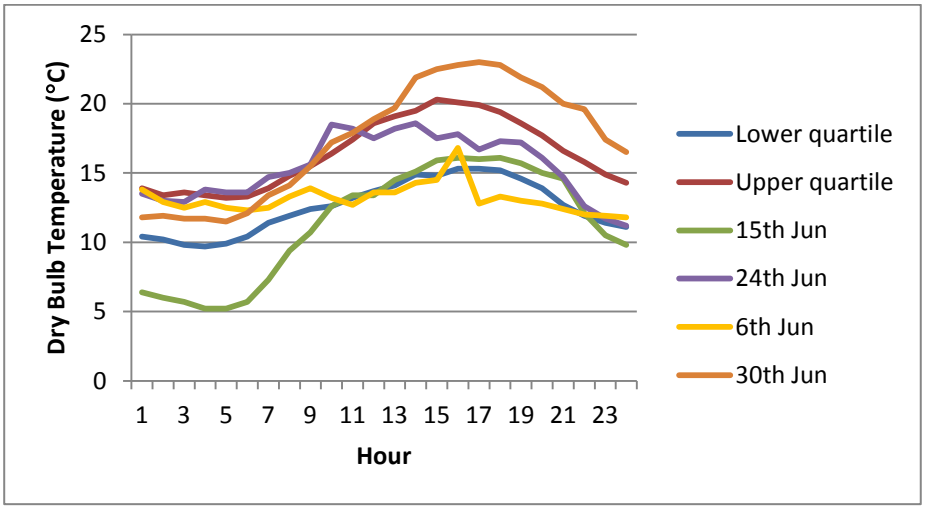

Figure 6.9: Hourly dry bulb temperature profiles for the four days with average dry bulb temperature closest to the average dry bulb temperature of the representative 2 month summer period, plotted alongside the hourly upper and lower quartile dry bulb temperature profiles for the same period.

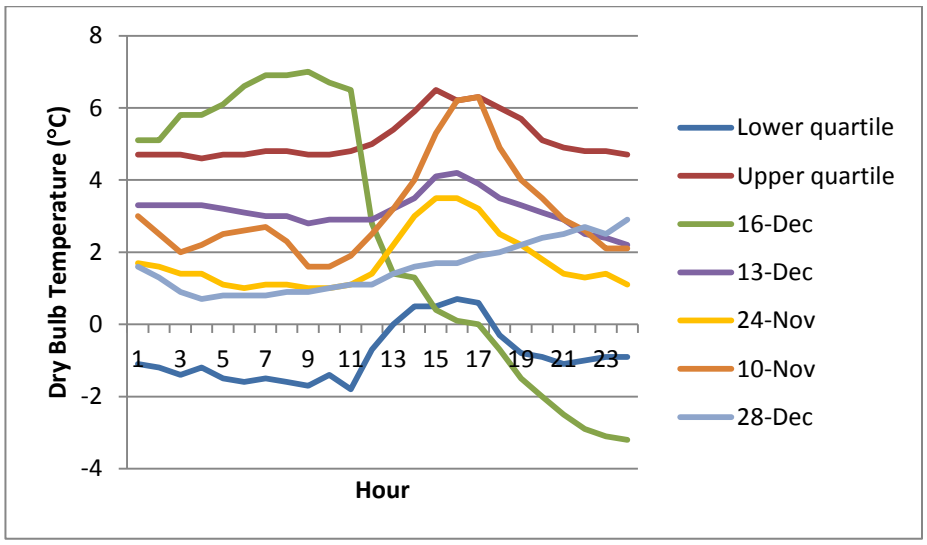

Figure 6.10: Hourly dry bulb temperature profiles for the five days with average dry bulb temperature closest to the average dry bulb temperature of the representative 2 month winter period, plotted alongside the hourly upper and lower quartile dry bulb temperature profiles for the same period.

For extreme summer conditions, the single warmest day in the entire 2010 Bingley weather data was taken. This was found by taking the 5 days in the 12 month period with the highest average dry bulb temperature, and simulating model performance under these climate conditions, as solar gains may have a significant effect on model performance in addition to dry bulb temperature. The model set-up was simplified for these simulations, not including the CFD domain and having assumed constant values of radiant panel temperature, corridor temperature and door discharge factor. Since the purpose of these simulations was to assess comparative performance, local air temperatures and exploration of design variable variation were not required; it was the comparison between the model predictions for the 5 different days that was important, not the predictions themselves. The 23rd of May was found to have the highest average predicted room temperature, so this was taken as the representative extreme summer day.

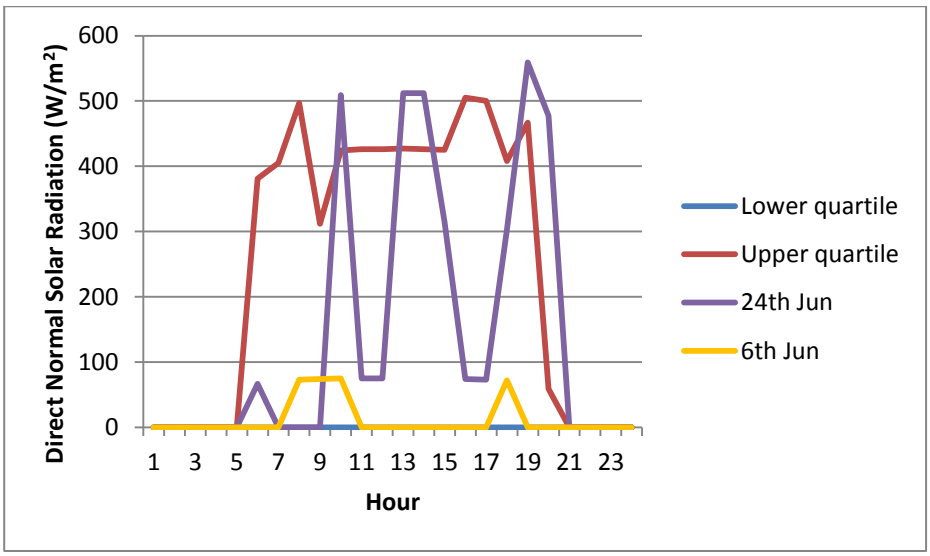

Figure 6.11: Hourly direct normal solar radiation profiles for the two days with the most average dry bulb temperature profiles, plotted alongside the hourly upper and lower quartile direct normal solar radiation profiles for the representative 2 month summer period.

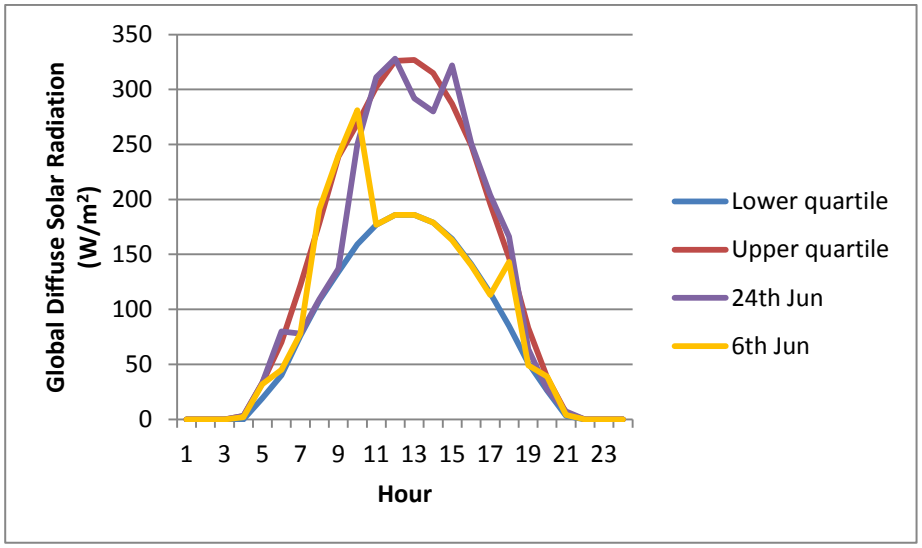

Figure 6.12: Hourly global diffuse solar radiation profiles for the two days with the most average dry bulb temperature profiles, plotted alongside the hourly upper and lower quartile global diffuse solar radiation profiles for the representative 2 month summer period.

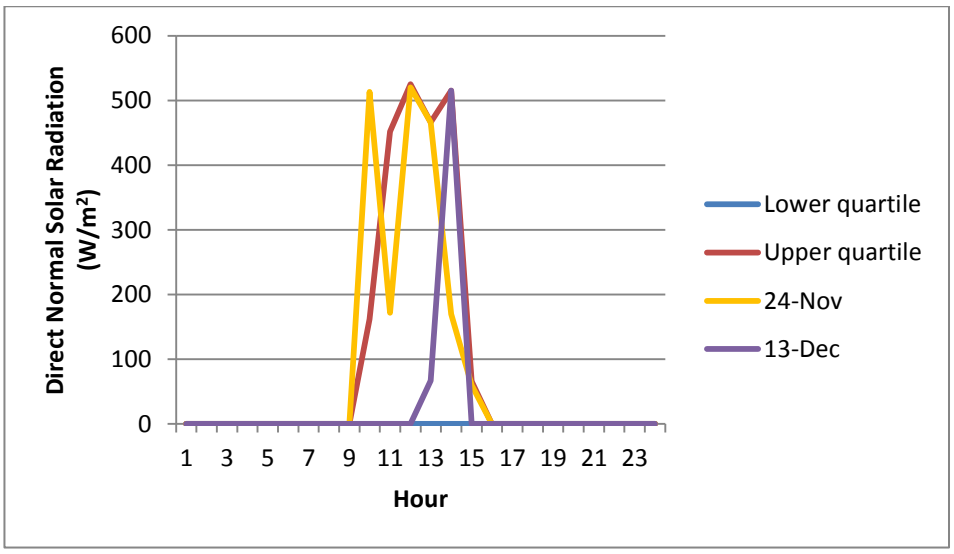

Figure 6.13: Hourly direct normal solar radiation profiles for the two days with the most average dry bulb temperature profiles, plotted alongside the hourly upper and lower quartile direct normal solar radiation profiles for the representative 2 month winter period.

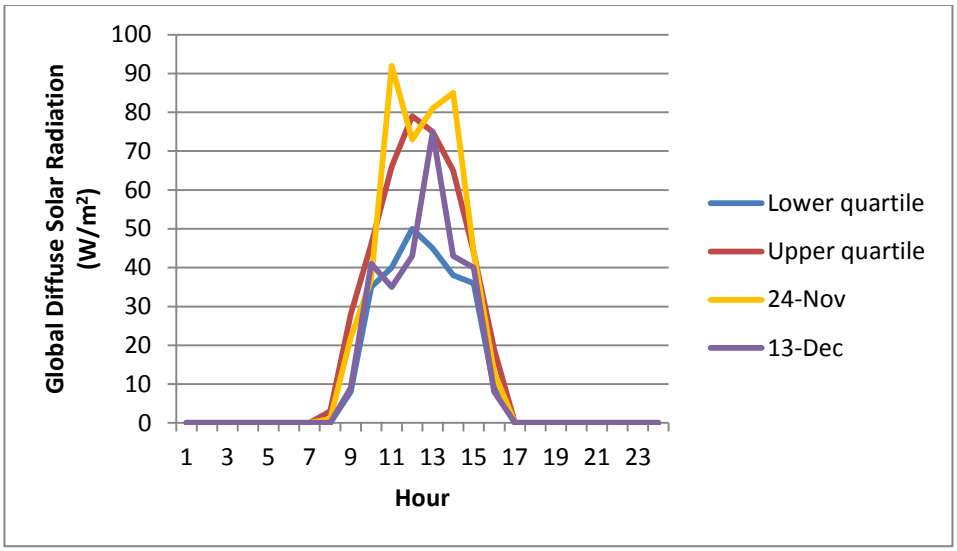

Figure 6.14: Hourly global diffuse solar radiation profiles for the two days with the most average dry bulb temperature profiles, plotted alongside the hourly upper and lower quartile global diffuse solar radiation profiles for the representative 2 month summer period.

For extreme winter conditions, a synonymous process was followed; the 5 days in the representative 2 month winter period with, on average, the coldest dry bulb temperatures were simulated with the model. The 19th of December was found to have the coldest predicted room temperature, so this was taken as the representative extreme winter day. The process may be summarised by the following steps:

1. Find the 5 days with the highest/lowest dry bulb temperature.
2. Simulate model performance with these days' climate.

3. Day resulting in highest/lowest average room temperature was selected to represent extreme conditions.

Climate case	Day	Average dry bulb temperature (°C)
Extreme summer	23rd May	19.09
Average summer day 1	6th Jun	14.28
Average summer day 2	24th Jun	14.46
Extreme winter	19th Dec	-4.42
Average winter day 1	13th Dec	3.16
Average winter day 2	24th Nov	1.75

Table 6.6: Showing selected days for each climate case to be used in the case studies.

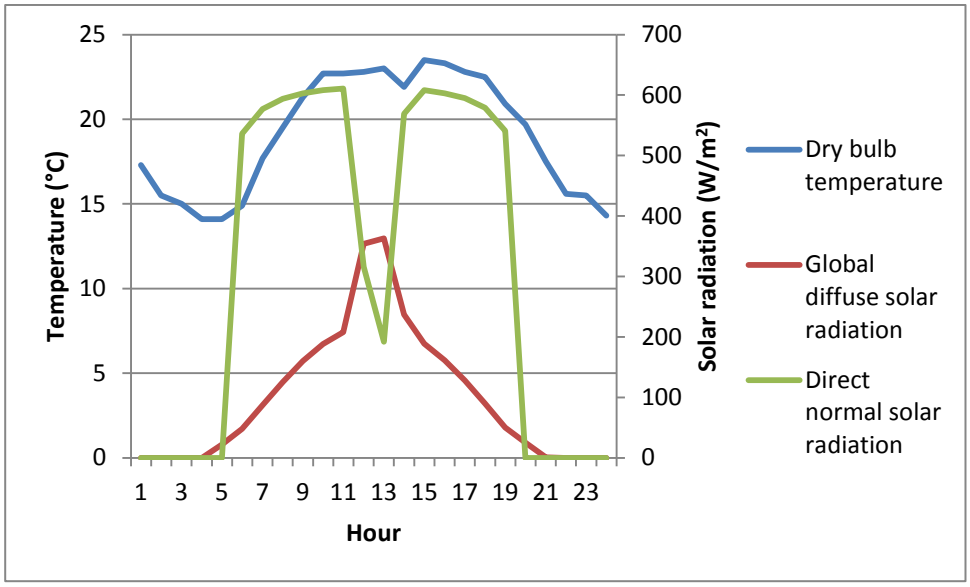

Figure 6.15: Dry bulb temperature and solar radiation for extreme summer day (23rd May).

For average conditions, the process was not quite so simple for a number of reasons. Firstly, for average summer conditions at least, a representative period had to be selected from the Bingley climate data. The two months with the highest average dry bulb temperature were selected for this; June and July. Secondly, the three most important climate metrics (dry bulb temperature, direct normal solar radiation and global diffuse solar radiation) might not necessarily agree with one another when determining the average day. Therefore, it was decided that two representative days would be selected for both summer and winter average conditions.

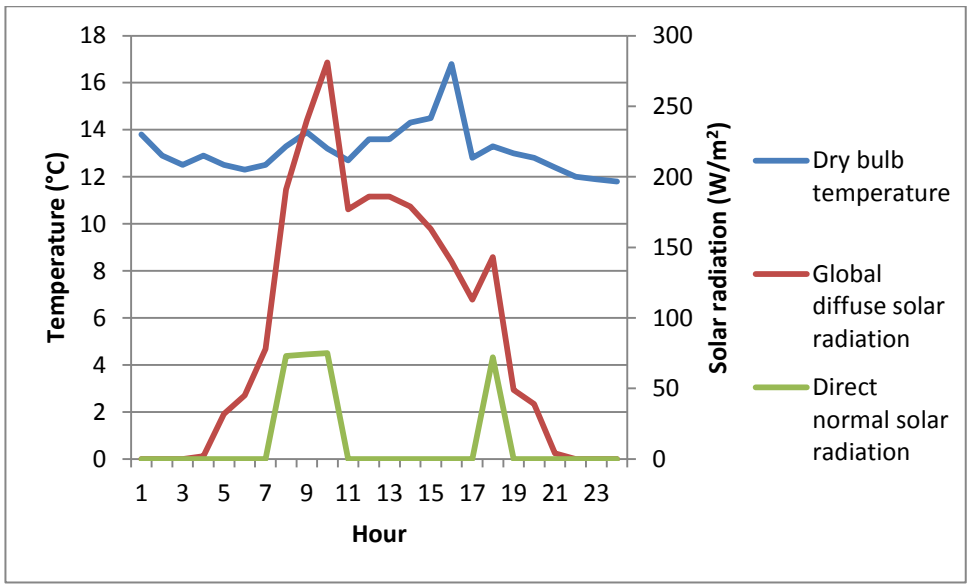

Figure 6.16: Dry bulb temperature and solar radiation for average summer day 1 (6th June).

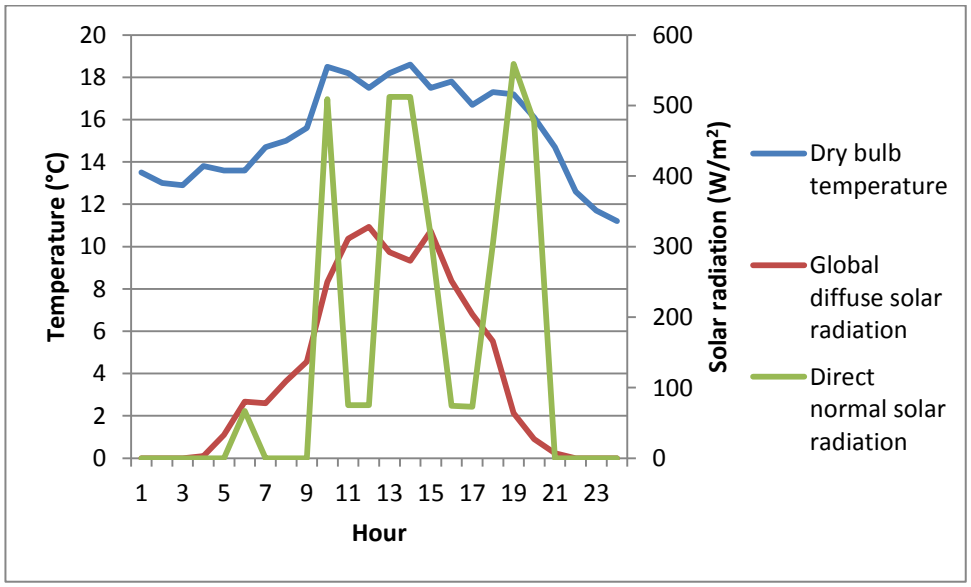

Figure 6.17: Dry bulb temperature and solar radiation for average summer day 2 (24th June).

To determine average conditions, the following process was followed:

1. For each hour in the 24 hour period, the upper and lower quartile of dry bulb temperature was found for the representative 2 month period.
2. These upper and lower quartile daily profiles were plotted on a graph.
3. Then the temperature profiles for the four or five days with average dry bulb temperature close to the average dry bulb temperature for the representative 2 month period, were plotted alongside the quartile profiles.
4. Using this graph, the two days that gave the best representation of the quartile profiles were selected.

5. Then the upper and lower quartile profiles for direct normal solar radiation and global diffuse solar radiation were found and plotted on graphs.
6. Direct normal solar radiation and global diffuse solar radiation for the two days selected in step 4 were then plotted alongside the quartile profiles, to ensure that they gave a broadly representative view of these two variables.

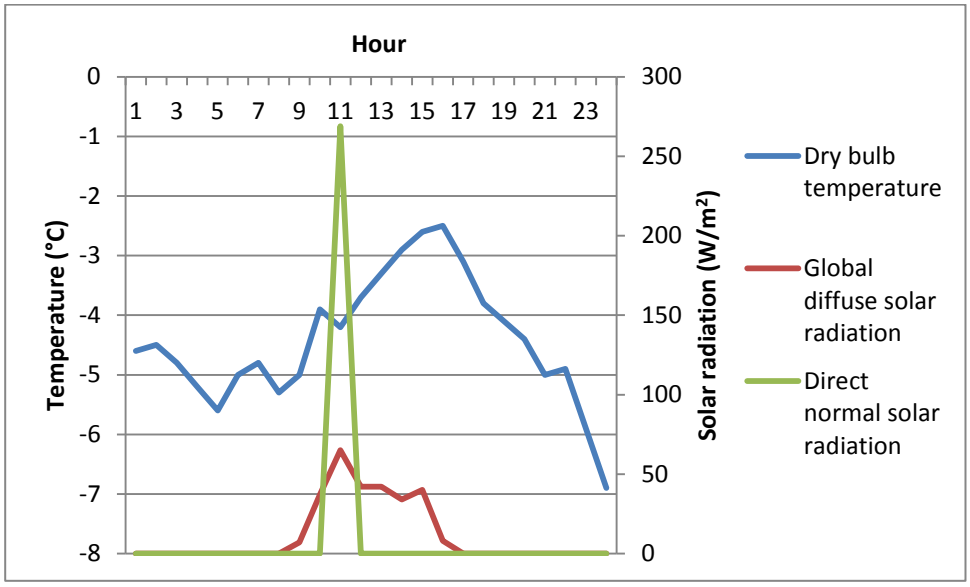

Figure 6.18: Dry bulb temperature and solar radiation for extreme winter day (19th December).

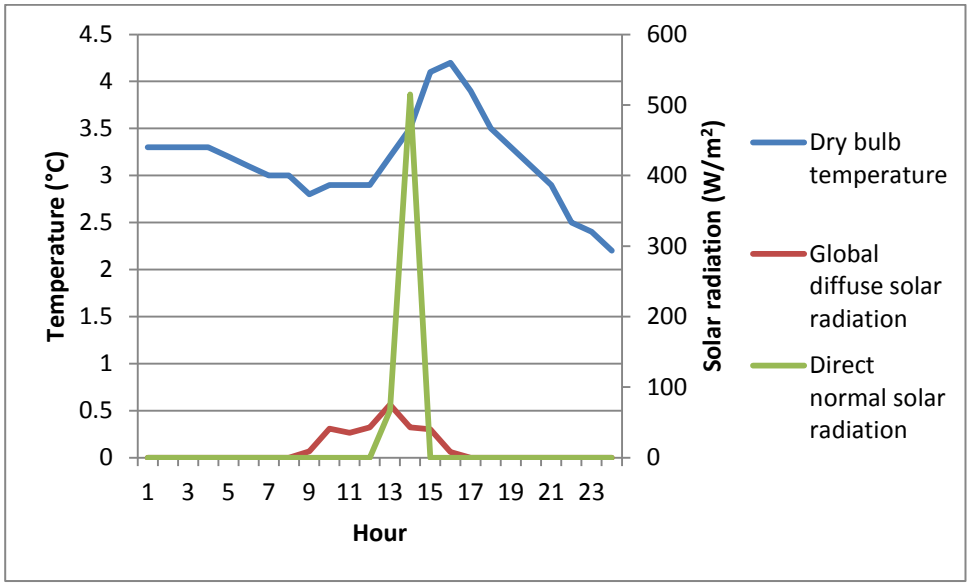

Figure 6.19: Dry bulb temperature and solar radiation for average winter day 1 (13th December).

Fig. 6.9 and 6.10 show the graphs resulting from step 3 for the summer and winter periods respectively. For the summer period (Fig. 6.9), the 24th and the 6th of June were selected as being most representative of average conditions. For the winter period (Fig. 6.10), the 24th of

November and the 13th of December were selected. Figs. 6.11 and 6.12 show the graphs resulting from step 6 for the summer period, and Figs. 6.13 and 6.14 show the same graphs for the winter period.

From Figs. 6.11 and 6.12, it can be seen that the 6th of June and the 24th June are broadly representative of average summer conditions in terms of solar gains. From Figs. 6.13 and 6.14, it can be seen that the 24th of November and the 13th of December are also broadly representative of average winter conditions in terms of solar gains. Table 6.6 gives a summary of the selected days for each climate case. Dry bulb temperature and solar gains for each selected day are shown in Figs. 6.15-6.20.

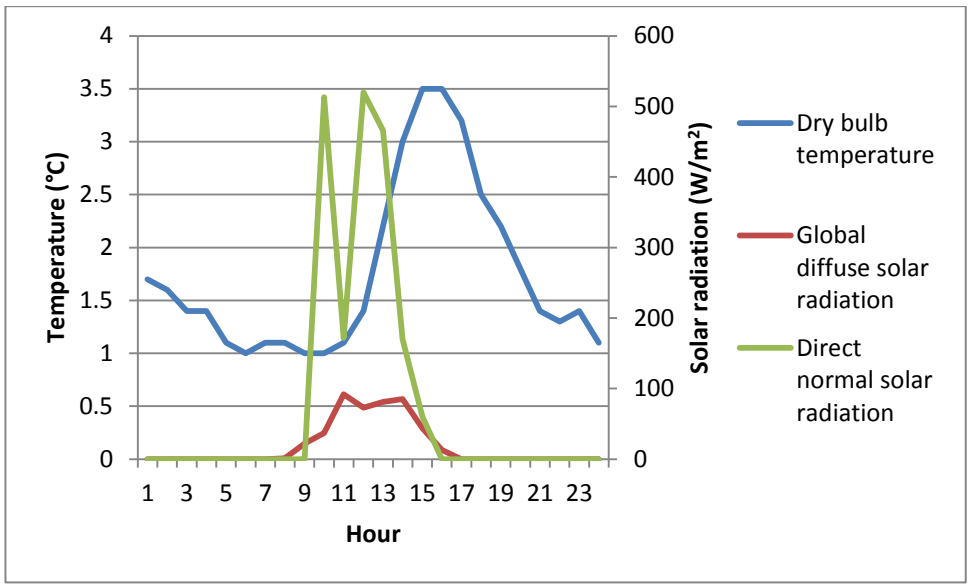

Figure 6.20: Dry bulb temperature and solar radiation for average winter day 2 (24th November).

6.4 Application of Case Study Model in T-BOT

6.4.1 Design variables

A total of 12 design variables were made programmed into T-BOT for the case study model. These were:

1. External wall thermal conductivity
2. External wall heat storage properties
3. Window nominal opening area
4. Window control opening area ratio 1
5. Window control opening area ratio 2

6. Window control opening area ratio 3
7. Window control temperature set-point 1
8. Window control temperature set-point 2
9. Window control temperature set-point 3
10. Radiant panel temperature
11. Door discharge factor
12. Corridor temperature

1. External wall thermal conductivity

This was identical the synonymous design variable in the test model, described in section 4.6.2. Values were linearly scaled to a range of 1 to 11 for the purposes of optimisation.

2. External wall heat storage properties

This also was identical to the synonymous design variable in the test model, controlling both density and specific heat capacity. A description of this design variable may also be found in section 4.6.2. Values were linearly scaled to a range of 1 to 11 for the purposes of optimisation.

3. Window nominal opening area

This design variable controlled the nominal window opening area, used in the control algorithm for the window as described in section 6.3.2. Values were linearly scaled to a range of 1 to 11 for the purposes of optimisation.

4, 5 & 6. Window control opening area ratios

These design variables controlled the opening area ratios 1, 2 and 3 respectively, also used in the window control as described in section 6.3.2. Values of all three were linearly scaled to a range of 1 to 11 for the purposes of optimisation.

7, 8 & 9. Window control temperature set-points

These design variables controlled the temperature set-points 1, 2 and 3 respectively, which dictated the operation of the window control algorithm as described in section 6.3.2. Values of all three were linearly scaled to a range of 1 to 11 for the purposes of optimisation.

However, because these variables needed to be sequential (ie. *set-point 1 < set-point 2 < set-point 3*), a distinct additional scaling system had to implemented. It was decided to use a sequential dynamic scaling process whereby the aggregate range of all three of the design variables was effectively rescaled to the range of one of them (note that all 3 of these design

variables were required to have the same limits). This was done by applying the following process before the samples were simulated:

- First the value of set-point 2 was linearly rescaled such that the minimum limit of the design variable was the value of set-point 1; its maximum was unchanged.
- Then the value of set-point 3 was linearly rescaled such that it's minimum value was the rescaled value of set-point 2; again its maximum was unchanged.

This ensured that no matter what the sample coordinates in the design space in terms of these three design variables, the relationship *set-point 1 < set-point 2 < set-point* 3 would always be satisfied. This was necessary as the DOE could not be constrained to force this relationship.

Before the samples were then passed to the optimisation program, these rescaled design variable values were replaced with the original unscaled values, and then linearly scaled to values of 1 to 11 as normal. This allowed the optimisation to be independent of the dynamic scaling process described above, ie. the optimisation did not need to be constrained to force the relationship *set-point 1 < set-point 2 < set-point 3*; the dynamic scaling was already inherent in the sample results and was hence encapsulated by the metamodels. Finally, after rescaling the design variable values from the optimised solutions back to actual values, the dynamic scaling process was applied as above, which then gave a true representation of how the design variable values applied to the model.

10. Radiant panel temperature

This design variable controlled the internal temperature of the radiant panel on the ceiling. Values were linearly scaled to a range of 1 to 11 for the purposes of optimisation.

11. Door discharge factor

This design variable controlled the discharge factor of the door between the corridor and the room, effectively dictating how far open the door was. This variable was implemented in the AFN, as described in Table 6.5. Values were linearly scaled to a range of 1 to 11 for the purposes of optimisation.

12. Corridor temperature

This design variable controlled the temperature of the air in the corridor. Values were linearly scaled to a range of 1 to 11 for the purposes of optimisation.

6.4.2 Objective functions

As with the test model, two objective functions were included in the optimisations for the case studies; time average deviation of operative temperature from a defined comfort range, and energy use. Whilst these objective functions were generally as described in section 4.6.1, a few notable adjustments were made for the purposes of the case studies; these are described here.

PMV formulation of thermal discomfort objective

For the case study model, PMV [Fanger, 1972] was implemented for use in the thermal discomfort objective, as the literature review (sub-section 2.1.4) revealed that a thermally comfortable environment for patients does not necessarily create a comfortable environment for staff as well. These differences in comfort between patients and staff could be taken into account through use of the PMV model by varying input values of clothing and metabolic rate. However, due to the necessity of simplified ventilation modelling in the corridor zones adjacent to the actual room, humidity was not controlled in the model. This meant that PMV values extracted directly from ESP-r (or calculated solely from values extracted from ESP-r) could not be relied upon. As such, an alternate approach was taken whereby limits of operative temperature were calculated from limits of PMV using assumed values including humidity (given in Table 6.7). Limits of PMV were taken from Fig. 2.1. Inspiration for this technique was taken from BS EN ISO 7730 [BSi, 2005], which gives in Annex E tables of operative temperature limits for given limits of PMV; these are synonymous with the criteria used to convert the objective function to terms of PMV in the present work.

In the calculation of PMV air temperature and MRT have different influences. Because of this, at a constant operative temperature the value of PMV can vary. As such, some assumption must be made on the maximum variation between air temperature and MRT in order to develop criteria for operative temperature from limits of PMV. As in BS EN ISO 7730 [BSi, 2005] a maximum variation of 5°C was assumed; it is reported that this ensures an accuracy of at least 0.1. The limits of operative temperature that were imposed were those at which the values of PMV were at most +-0.2 for both extremes of the 5°C maximum variation, ensuring the criteria erred on the side of caution in terms of this variation. This is explained mathematically below.

The formula for calculation of PMV, taken from BS EN ISO 7730:2005 [BSi, 2005] is given in Equation 6.2:

$$PMV = \left[0.303e^{(-0.036M)} + 0.028\right] \cdot \left\{ \begin{array}{c} (M-W) - 3.05 \cdot 10^{-3}[5773 - 6.99(M-W) - P_a] \\ -0.42[(M-W) - 58.15] - 1.7 \cdot 10^{-5} M(5867 - P_a) \\ -1.4 \cdot 10^{-3} M(34 - \theta_a) - 3.96 \cdot 10^{-8} f_{cl}[(\theta_{cl} + 273)^4 - (\theta_r + 273)^4] \\ -f_{cl} h_c(\theta_{cl} - \theta_a) \end{array} \right\} \quad (6.2)$$

Where:

- M is metabolic rate (W/m^2);
- W is effective mechanical power of work being performed (W/m^2);
- I_{cl} is clothing insulation (m^{2}K/W);
- f_{cl} is clothing surface area factor, calculated according to Equation 6.3;

$$f_{cl} = \begin{cases} 1.00 + 1.29 I_{cl}, & I_{cl} \leq 0.078 \\ 1.05 + 0.645 I_{cl}, & I_{cl} > 0.078 \end{cases} \quad (6.3)$$

- θ_a is air temperature (°C);
- θ_r is mean radiant temperature (°C);
- v_{ar} is relative air velocity (m/s);
- P_a is water vapour partial pressure (Pa);
- h_c is convective heat transfer coefficient (W/m^{2}K), calculated according to Equation 6.4;

$$h_c = \begin{cases} 2.38|\theta_{cl} - \theta_a|^{0.25}, & 2.38|\theta_{cl} - \theta_a|^{0.25} > 12.1\sqrt{v_{ar}} \\ 12.1\sqrt{v_{ar}}, & 2.38|\theta_{cl} - \theta_a|^{0.25} < 12.1\sqrt{v_{ar}} \end{cases} \quad (6.4)$$

- θ_{cl} is clothing surface temperature (°C), calculated iteratively according to Equation 6.5;

$$\theta_{cl} = 35.7 - 0.028(M-W) - I_{cl}\{0.0000000396 f_{cl}[(\theta_{cl} + 273)^4 - (\theta_r + 273)^4] + f_{cl} h_c(\theta_{cl} - \theta_a)\} \quad (6.5)$$

Along with this equation, BS EN ISO 7730 [BSi, 2005] also gives a computer code implementation of it that will calculate PMV from given input variables. A functionally identical program, implemented in MATLAB, was used in the present work to calculate PMV (notated as $PMV(\ldots)$ in formulas herein); this code is given in appendix A.

Limits of operative temperature were calculated with simple optimisation problems; a separate optimisation problem was required for each of the two limits. The objective function

of these small scale optimisation problems was variation of PMV from the limit under consideration:

$$F\left(D_{\theta_o}\right) = \left|PMV_{limit} - PMV_r\left(D_{\theta_o}\right)\right| \tag{6.6}$$

Where $F(\theta_o)$ is the objective function value subject to the operative temperature design variable D_{θ_o}, PMV_{limit} is the limit of PMV under consideration, and $PMV_r\left(D_{\theta_o}\right)$ is a representative value of PMV calculated from the design variable D_{θ_o} according to equation 6.7:

$$PMV_r\left(D_{\theta_o}\right) = \sqrt{max\begin{Bmatrix} PMV^2([\theta_a, \theta_r]_{ex1}) \\ PMV^2([\theta_a, \theta_r]_{ex2}) \end{Bmatrix}} \tag{6.7}$$

Where $PMV^2(\dots)$ is shorthand for $\left(PMV(\dots)\right)^2$, and $[\theta_a, \theta_r]_{ex1}$ and $[\theta_a, \theta_r]_{ex2}$ are the two most extreme permutations of air temperature θ_a and MRT θ_r that result in θ_o (calculated according to equation 4.12 or 4.13 depending on air speed), considering the maximum variation of 5°C (eg. assuming air speed below 0.1 m/s, if $\theta_o = 20$ then $[\theta_a, \theta_r]_{ex1} = [17.5, 22.5]$ and $[\theta_a, \theta_r]_{ex2} = [22.5, 17.5]$). The optimum values of the design variable for these problems are hence the limiting operative temperatures within which PMV does not stray outside the defined limits, effectively converting the thermal discomfort objective function to terms of PMV when these operative temperature limits are imposed as $\theta_{o,lb}$ and $\theta_{o,ub}$ in equation 4.15.

Case	PMV limits	Clothing (clo)	Metabolic rate (met)	Air speed (m/s)	Humidity (%)	Operative temperature limits (°C)
Patient	+-0.2**	0.7*	1*	0.1	50	24.511 - 25.408
Visitor	+-0.5***	1	1	0.1	50	21.690 - 24.906
Staff	+-0.5***	0.8*	1.2*	0.1	50	21.103 - 24.650
* Source: [de Giuli et al., 2013].						
** Category I in Table 2.1, particularly thermally sensitive.						
*** Category II in Table 2.1, normal thermal sensitivity.						
All other values assumed.						

Table 6.7: Details for different cases examined using the PMV formulation of thermal discomfort objective.

There were three distinct cases of PMV input variables intended to represent different classes of occupant. The input variable values and resulting limits of operative temperature for the three cases are given in Table 6.7.

Formulation of energy objective

In the case study model, characterisation of ventilation energy use had to be revised in order to accommodate the additional ventilation modelling described in section 6.3.3. Radiant energy loads were extracted directly from ESP-r as described in section 4.6.1.

As detailed in section 6.3.3, controlled ventilation in the case study model consisted of flow through the door from the corridor; there was no direct active air supply to the room. This made characterisation of ventilation energy use difficult, as the HVAC in the corridor would need to be modelled in full in order to obtain accurate values of energy loads. As a result, it was decided to examine energy contributions to the room from the corridor; this was extracted directly from ESP-r in the form of net hourly energy flow due to air flowing from the corridor into the room. Air flowing the other way, from the room to the corridor, was not considered. This was because it was warmth (or coolth) delivered from the corridor that was of interest in the objective; if the airflow through the door is only flowing from the room to the corridor then no energy is being delivered to the room by the HVAC system. In this situation it is likely that the radiant system in the room would need to increase their energy use significantly to maintain comfort; hence the energy objective (ie. the sum of the ventilation and radiative components) should adjust itself accordingly provided the values are appropriately scaled.

The radiant and ventilation components were both in units of kW/h. Nevertheless, they measured theoretically different quantities; whereas the radiant component measured energy delivered to the radiant panel to maintain the temperature, the ventilation component directly measured energy contribution to the room from the system. However, the radiant panel was modelled as being highly insulated on all sides but that exposed to the room. It was therefore assumed that energy lost from the radiant panel that is not delivered to the room was negligible, and hence the net energy loss from the panel was an acceptable approximation for the energy delivered to the room. To maintain the operational temperature, energy supplied to the panel must equal net energy loss. The two components were therefore considered to be comparable with equal weighting.

6.4.3 Further developments to the methodology

In view of the results of the proof of concept study detailed in Chapter 5, a number of additional developments were made to the general methodology to attempt to address various issues.

Categorical design variables

As described in section 4.6.2, in the test model was included a discrete categorical design variable. This effectively increased the number of simulations necessary by a factor of the number of possible values of the design variable. The case study model was more complex than the test model, and took far longer to simulate (this is discussed in detail in section 5.3.2). As such, discrete states were considered as entirely separate cases as opposed to a design variable. This simplified the process of writing design variable writer scripts (described in section 4.7) for the test model significantly, as the model manipulation scripts necessary for categorical design variables were far more complex to write than value substitution scripts.

Improvements to the CFD

As detailed in section 5.2.2, it was found that there was very little spatial variation in air temperature in the test model. As such, when developing the case study model, it was decided to address this by assigning casual gains from occupants to specific nodes within the CFD grid. Further to this, buoyancy modelling was included in the CFD calculations in the form of full density calculations by the ideal gas law:

$$PV = nG\theta \tag{6.8}$$

Where P is the pressure (Pa), V is the volume (m^3), n is the amount of gas (moles), G is the ideal gas constant, and θ is the temperature.

The Boussinesq approximation was considered as an alternative approach however this requires an arbitrary reference temperature to be set based on the temperature range in the room. As the process of running simulations needed to be fully automated by code, it was difficult to set an appropriate reference temperature for each simulation individually, so the ideal gas law approach was taken instead. This was not found to compromise the simulations in any way; simulation time was increased by only 2.4% and convergence remained unaffected.

To clearly show the effect of these improvements, two simulations were performed with the case study model; one with the buoyancy modelling and local heat source and one without. The casual gains were assumed to be associated with a patient lying on the bed. Since the

climate was arbitrary for these simulations, they were performed with the default ESP-r test weather data.

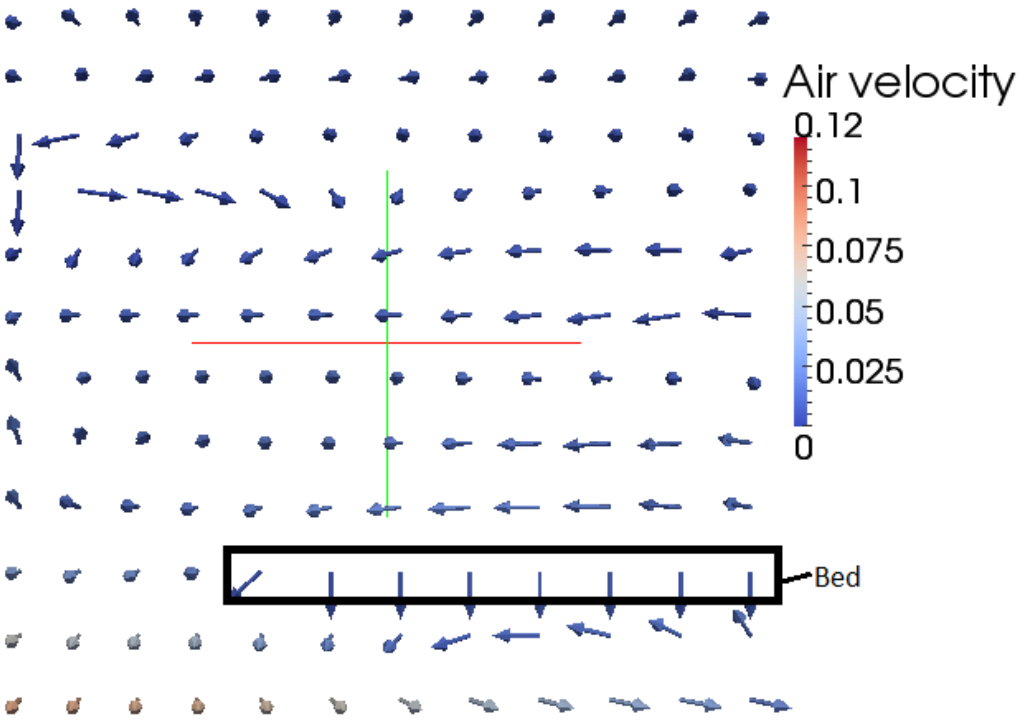

Figure 6.21: Air flow results without buoyancy modelling and locally modelled patient, coloured by air velocity.

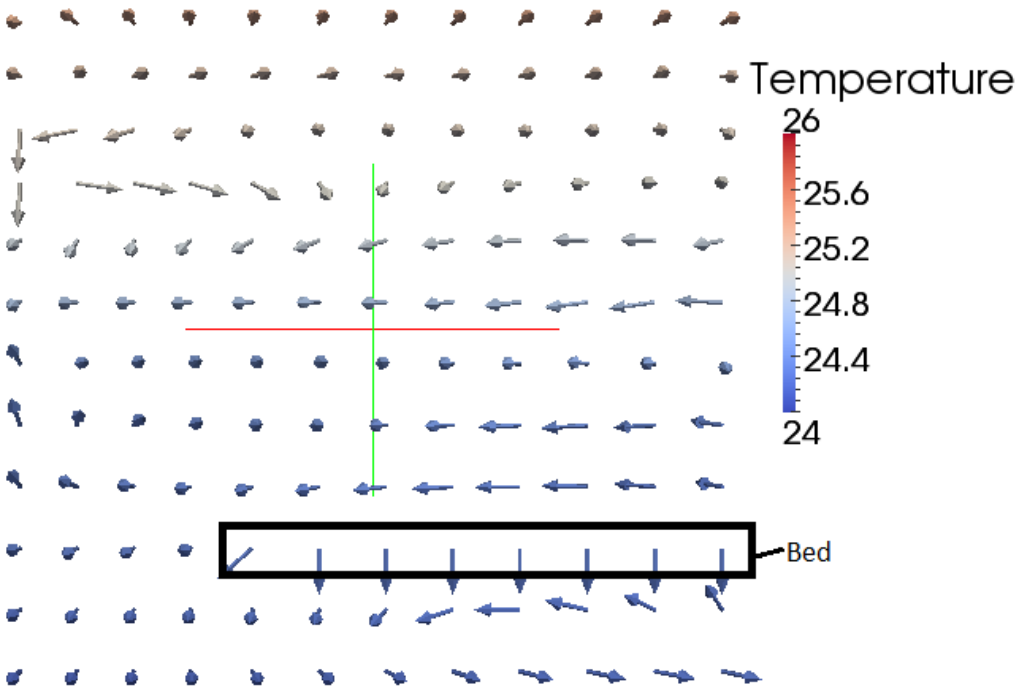

Figure 6.22: Air flow results without buoyancy modelling and locally modelled patient, coloured by temperature.

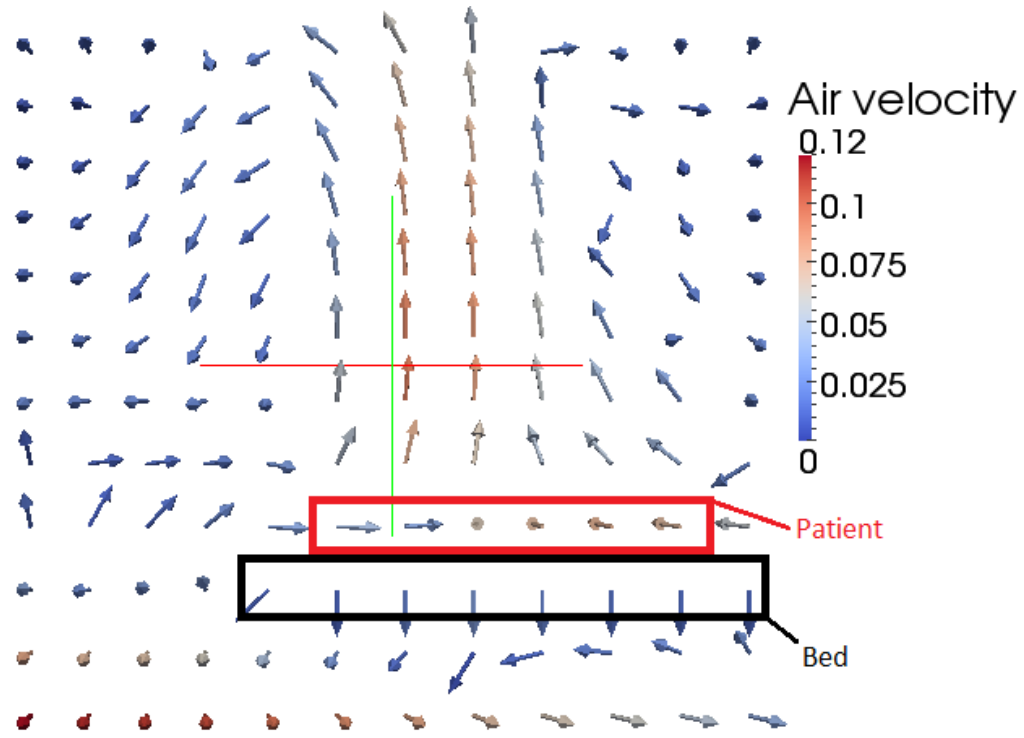

Figure 6.13: Air flow results with buoyancy modelling and locally modelled patient, coloured by air velocity.

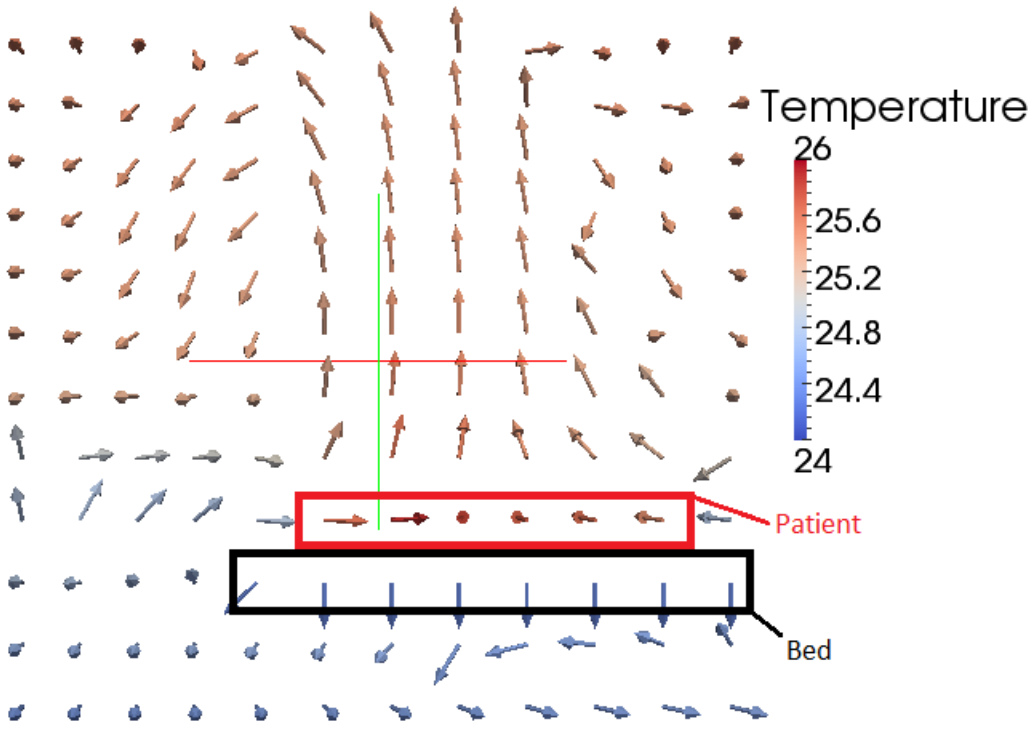

Figure 6.24: Air flow results with buoyancy modelling and locally modelled patient, coloured by temperature.

Figs. 6.21 and 6.22 show airflow results on a vertical plane that encapsulates the bed and patient (these are marked on the figures) for the case without buoyancy modelling and local heat gains, coloured according to air velocity and temperature respectively. Figs. 6.23 and 6.24 show synonymous results for the case with these improvements. From comparison of the two sets, it is clear that the additional modelling in the CFD domain results in a thermal plume from the patient, stimulating vertical mixing in the air flow. It also results in greater variation of both air velocity and temperature in the room, which supports the notion of evaluating comfort locally.

Improvements to the DOE

The DOE described in section 4.5 is relatively basic, with no consideration given to sample distribution of the combined sample set. To try to reduce the tendency of MLSR to over-model the responses, an improved DOE was implemented for the case studies.

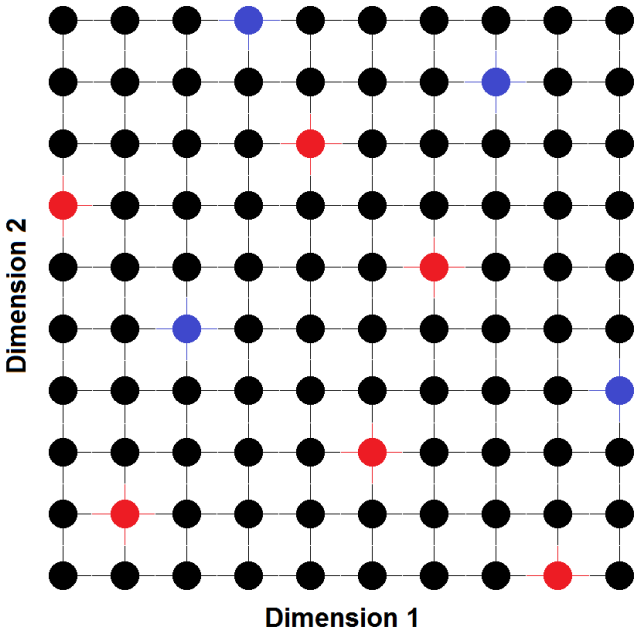

Figure 6.25: Illustration of nested build (red) and validate (blue) samples.

Essentially, the new approach provided nested sample distributions; that is to say the two sample sets (ie. “build” and “validation” sets) do not have independent sample distribution, but rather are optimised to provide a well distributed sample in terms of the combined set, as well as individually. This concept is illustrated in Fig. 6.25 using a simple case with 2 dimensions. A build sample of 6 and a validate sample of 4 is chosen, so there are 10 sample points in total. Each dimension is then discretised to 10 levels as shown. It can be seen that the build sample (red points), the validate sample (blue points) and the combined sample all satisfy the criteria of a latin hypercube design; no point has a value in any dimension equal to

that of any other point. There are several benefits to the nested approach. Firstly, the improved distribution of the combined sample ensures that validation sample points are as far from build sample points as possible, which ensures they provide a better estimation of metamodel performance. Secondly a better distributed combined sample should result in improved fidelity from the resulting metamodels.

The new approach introduced optimisation into the process to simultaneously optimise the distribution of the build, validate and combined samples. This was formulated with three equally weighted objective components (metrics of the two individual sample distributions and the distribution of the combined set) and design variables representing sample coordinates. Solution of this optimisation problem was performed with a GA by a 3rd party program written specifically for the task, termed permGA [Bates et al, 2004]. This GA was iterated 10,000 times.

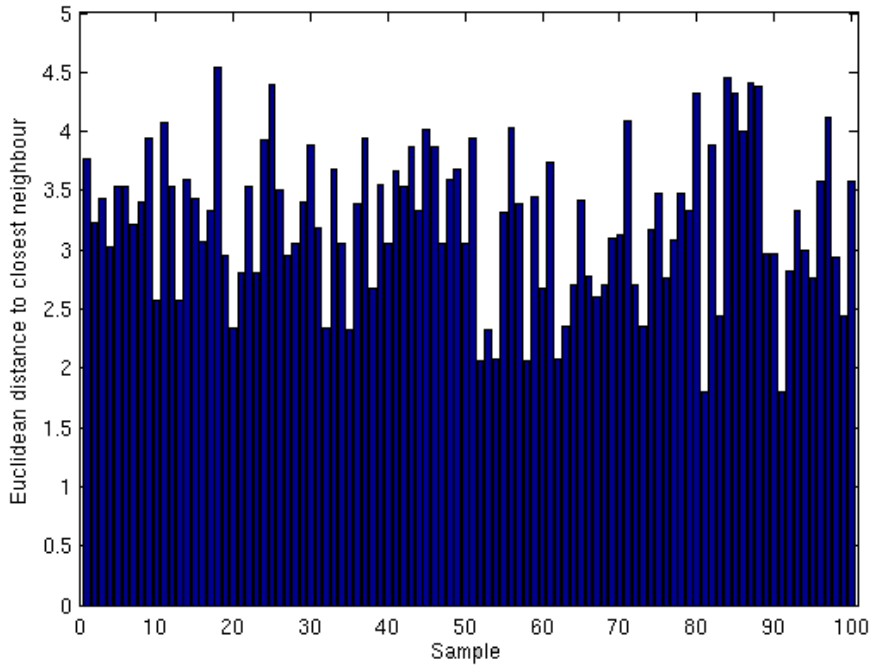

Figure 6.26: Sample dstribution bar-chart for a 100-point DOE of 5 design variables, constructed with the nested DOE technique.

Separate build and validate sets is particularly useful for MLSR, as the validate sample allows independent assessment of fit from the build sample. Predicted values of objective components at points in the validate set are compared to actual values from the building simulation, and these residuals can be used to form the objective of an optimisation problem to determine the optimum value of the "closeness of fit" parameter in MLSR. After the validate set is used to determine an optimum closeness of fit for the build set, the two sets are combined and the metamodel is reconstructed. However in the present work metamodel

responses were examined and if necessary the closeness of fit parameter was reduced to prevent over-fitting, as detailed in section 4.4.1.

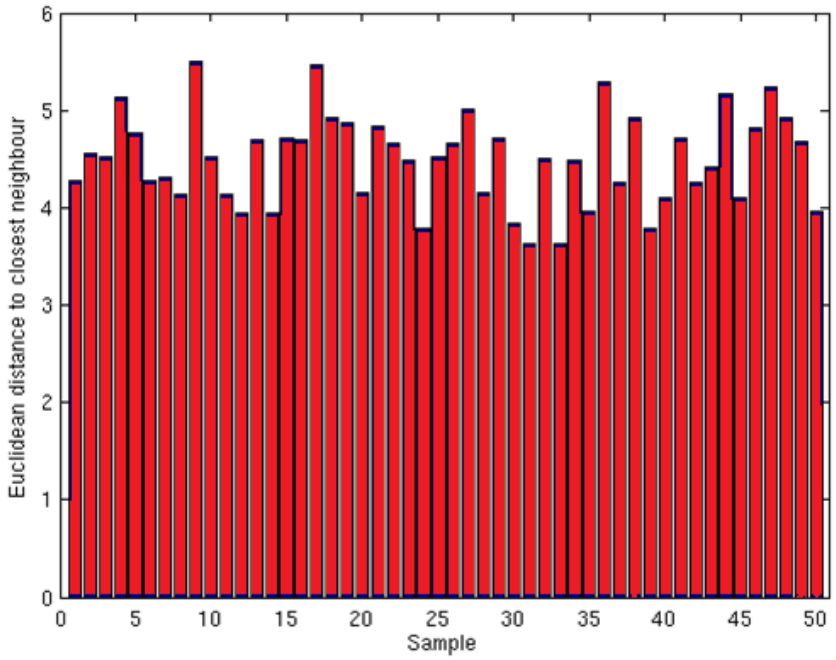

Figure 6.27: Sample dstribution bar-chart for a 50-point DOE of 5 design variables, constructed with the nested DOE technique.

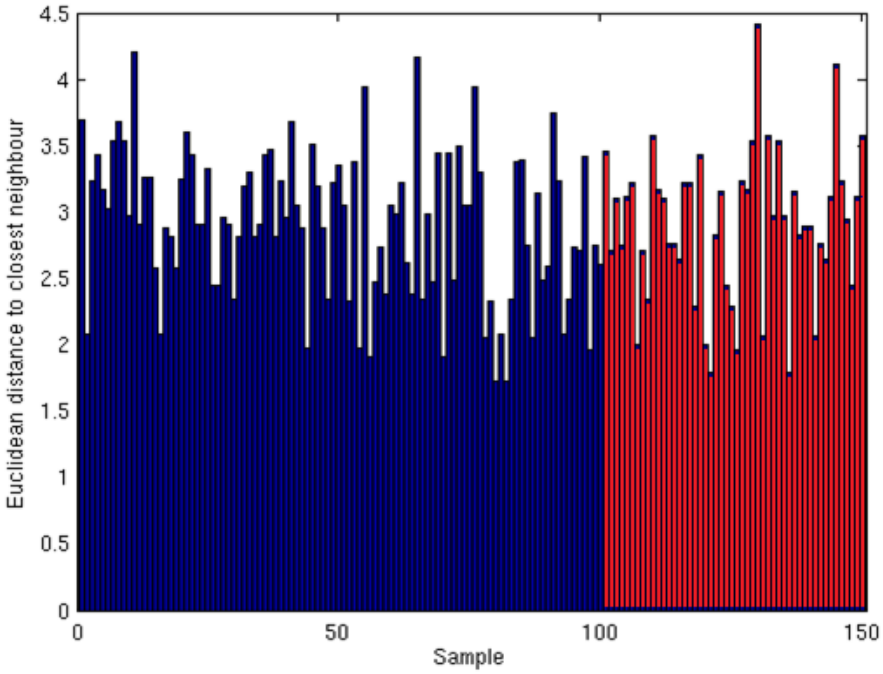

Figure 6.28: Sample distribution bar-chart for the combined 150-point DOE of 5 design variables constructed using the nested DOE technique; the build set is shown in blue and the validate set in red.

Sample distribution bar charts for the build sample set, the validate sample set, and the combined set, generated with the nested DOE are shown in Figs. 6.26 – 6.28 respectively. These samples were constructed with identical parameters to the sets reported in section 4.5; design variables 1 (external wall thermal conductivity), 2 (external wall heat storage properties), 10 (radiant panel temperature), 11 (door discharge factor) and 12 (corridor temperature) as reported in section 6.4.1, each having a range values of 1-11. To directly compare the standard and nested DOE techniques, the data for Figs. 4.15 – 4.17 reported in section 4.5 are recreated in Table 6.8 and contrasted against synonymous data for Figs. 6.26 – 6.28.

DOE	Sample set	Maximum variation in distance to closest neighbour, as percentage of maximum distance to closest neighbour	Mean of distance to closest neighbour	Standard deviation of distance to closest neighbour	Standard deviation as percentage of mean
Standard DOE	build	62%	3.15	0.621	20%
	validation	51%	3.86	0.505	13%
	combined	79%	2.81	0.51	18%
Nested DOE	build	60%	3.25	0.637	20%
	validation	34%	4.5	0.468	10%
	combined	61%	2.9	0.559	19%

Table 6.8: Contrasting sample distribution data for the old and new DOEs.

It can be seen from Table 6.8 that the nested DOE results in generally superior sample distributions. In the case of the 100 sample “build” set the nested DOE has slightly greater mean and standard deviation, and the maximum variation percentage is slightly lower. In this case then the sample distribution is better spread through the design space with more consistent distances between points; however this improvement is only marginal. For the 50 sample “validation” set, the mean is significantly higher for the nested DOE whilst the standard deviation is a little lower and the maximum variation percentage is greatly reduced. This suggests that this sample set shows a far more marked improvement over the standard DOE. Whilst the reduced standard deviation would on its own normally imply a sample with marginally poorer spread, the greatly increased mean and greatly reduced maximum variation percentage belie a far better distributed and more consistent sample set. In the case of the combined set, similarly to the “build” set both the mean and standard deviation are slightly

higher, however similarly to the "validation" set the maximum variation percentage is significantly lower. This implies that whilst the spread of the combined sample is only marginally increased, the consistency of the distribution is significantly improved with the new DOE. From these results, it was concluded that the nested DOE technique provided superior sample distributions across the board, and the resulting DOEs were therefore adopted for all case studies.

Developments to parallelism framework

Conflating an AFN to the CFD domain results in a co-dependency problem; the AFN flow rates depend on the pressure evaluated by the CFD, and the CFD solution depends on the flow rates calculated by the AFN. The two domains are therefore calculated iteratively until mutual convergence is achieved. It was found that simulations that converged had done so by the 11th iteration in almost all cases, so this was set as the maximum number of permitted iterations to limit computation time. Nevertheless an average of 10 iterations for each building simulation, each requiring a separate CFD simulation, still resulted in an increase in simulation time of approximately a factor of 10. This was found to be unsustainable within the parallelism framework used for the test model, as described in section 4.5. As such, two distinct developments were made to the parallelism framework; firstly the automation scripts were developed to allow splitting of samples over multiple nodes, and secondly T-BOT was implemented on a new HPC facility that became available during the course of the project.

In the early stages of development of T-BOT, due to the fact that the vast majority of controlling scripts were written in MATLAB, it was convenient to simply make use of the MATLAB Parallel Computation Toolbox to accomplish sample parallelisation. However, this was later found to be bound by a number of limitations. Firstly, the parallelism was constrained to shared memory parallelism; this essentially necessitates that all threads share common memory, and hence on a HPC cluster limits the computations to a single node. Secondly, as a result of this the maximum number of cores that could be used to simulate a sample was limited to the number of cores on the node minus one (this extra one must be set aside for the control process). The increase in computation power necessary to simulate samples for the case study model was not practicable within the confines of these limitations. As a result, the automation scripts were developed to allow samples to be split amongst multiple nodes of the HPC cluster. This was inspired by the realisation that shared memory parallelism was not strictly necessary for T-BOT; each individual simulation could be entirely independent of its peers. Therefore, a sample set may be arbitrarily split into any number of sub-sets, and each of these sub-sets simulated simultaneously on a different node. For example, say one needs to simulate a sample of 100 simulations, each taking 6 hours. Such a

sample would not be able to run on a single node, as 100 sims multiplied by 6 hours = 600 computation hours, divided by 7 threads = 86 hours run-time, which is greater than the 48 hour maximum run-time imposed on the HPC facility. However, splitting this sample over 4 nodes results in 600 computation hours divided by 7*4 threads = 21 hours run-time, which is feasible.

As well as these improvements, toward the end of the present project a brand new HPC facility came into service at the University of Leeds. This HPC facility was superior to the previous one, in that each node had 16 cores instead of 8. There were also other improvements, including better processors and quicker access to hard disk storage from the computational nodes. It was found that the MATLAB Parallel Computation Toolbox was further limited to 12 worker threads per instance, so only 13 of those 16 cores could be used per node. Nonetheless, the new HPC cluster was found to provide markedly improved performance over the old one, resulting in an overall improvement in simulation time of approximately 50%. Along with the developments to the parallelism framework described above, it was then found to be feasible to run practically any sized samples of the case study model on the new HPC cluster.

Chapter 7: Case Study Results

7.1 Chapter Overview

Taking on board lessons from the initial proof-of-concept study presented in Chapter 5, a case study was then performed to explore the performance of T-BOT in more detail. Changes to the methodology and the building model have been reported in detail in Chapter 6. In this chapter case study results are presented and their implications noted; in-depth analysis and discussion of these results and performance of the model in a wider context is left to Chapter 8. Section 7.2 summarises conditions of the study, sections 7.3 – 7.6 present results to demonstrate key outcomes, and section 7.7 analyses metamodel fidelity.

A portion of this work was presented in a conference paper written by the present author, with support from the project supervisors [Cowie et al., 2015]; specifically elements of sections 7.3 and 7.4.

7.2 Conditions of Study

A wide variety of cases were simulated, exploring differences between season, spatial location, optimisation period, design variables and thermal comfort criterion. Due to the volume of results, instead of reporting every case, sections 7.3 – 7.6 of this chapter present selected results exploring key outcomes. Full results from the case study are presented in Appendix B. Section 7.7 then analyses metamodel fidelity for a variety of cases shown in previous sections. Table 7.1 is a list of all simulated cases, and highlights the cases that are reported here, as well as indicates what they are used to demonstrate by the column in which they are highlighted. Figure 7.1 shows comfort evaluation locations associated with location indices. Table 7.1 also introduces case references used to identify cases for the purposes of brevity. Such case references are written in *italics*, and concisely define values of the contextual parameters reported in Table 7.1. More detailed guidance on the case referencing system is given in Appendix B.1

I can be seen from Table 7.1 that two different sets of design variables have been used in the case study. A different set of design variables necessitates a new sample. The first of these samples, with 5 design variables, operated under the assumption that the window remained closed at all times. This sample was of the default size; 100 build samples and 50 validation samples. The design variables included were:

- External wall thermal conductivity (0.1-1.5 W/mK)
- External wall thermal storage variables (500-2000 kg/m3 & J/kgK)
- Radiant panel temperature (5-30 °C for summer cases, 20-40 °C for winter cases).
- Door discharge factor (0.1-0.9)
- Corridor air temperature (18-26 °C for summer cases, 20-28 °C for winter cases)

Climate	Locations*	Periods	Design variables	Comfort metric	Case Reference
Extreme summer	15	24 hours	5	operative temp.	*SumHot1_15_1-24_dv5_ot*
			5	PMV - patient	*SumHot1_15_1-24_dv5_pmvP*
			9	operative temp.	*SumHot1_15_1-24_dv9_ot*
			9	PMV - patient	*SumHot1_15_1-24_dv9_pmvP*
		4 x 6 hours	5	operative temp.	*SumHot1_15_1-6,7-12,13-18,19-24_5_ot*
			9		*SumHot1_15_1-6,7-12,13-18,19-24_9_ot*
	13	24 hours	5	PMV - staff	*SumHot1_13_1-24_dv5_pmvS*
				PMV - visitor	*SumHot1_13_1-24_dv5_pmvV*
			9	PMV-staff	*SumHot1_13_1-24_dv9_pmvS*
			9	PMV - visitor	*SumHot1_13_1-24_dv9_pmvV*
Average summer day 1	9	24 hours	5	operative temp.	*SumAve1_9_1-24_dv5_ot*
	9	4 x 6 hours			*SumAve1_9_1-6,7-12,13-18,19-24_dv5_ot*
	13	24 hours			*SumAve1_13_1-24_dv5_ot*
	13		9		*SumAve1_13_1-24_dv9_ot*
Average summer day 2	9	24 hours	5	operative temp.	*SumAve2_9_1-24_dv5_ot*
	9	4 x 6 hours			*SumAve2_9_1-6,7-12,13-18,19-24_dv5_ot*
	13	24 hours			*SumAve2_13_1-24_dv5_ot*
	13		9		*SumAve2_13_1-24_dv9_ot*
Extreme winter	9	24 hours	5	operative temp.	*WinCol1_9_1-24_dv5_ot*
		4 x 6 hours			*WinCol1_9_1-6,7-12,13-18,19-24_dv5_ot*
	15	24 hours	5	operative temp.	*WinCol1_15_1-24_dv5_ot*
				PMV - patient	*WinCol1_15_1-24_dv5_pmvP*
Average winter day 1	9	24 hours	5	operative temp.	*WinAve1_9_1-24_dv5_ot*
	9	4 x 6 hours			*WinAve1_9_1-6,7-12,13-18,19-24_dv5_ot*
	17	24 hours			*WinAve1_17_1-24_dv5_ot*
Average winter day 2	9	24 hours	5	operative temp.	*WinAve2_9_1-24_dv5_ot*
		4 x 6 hours			*WinAve2_9_1-6,7-12,13-18,19-24_dv5_ot*
	17	24 hours			*WinAve2_17_1-24_dv5_ot*

* Location indices relate to Figure 7.1.

Table 7.1: List of all simulated cases, highlighting cases reported in this chapter.

Design variables 1 and 2, pertaining to external wall material properties, were found not to give useful information in the majority of the initial run of SumHot1 cases. As such, in subsequent cases using the 5 design variable sample (i.e. any case not under extreme summer

conditions) these two design variables were frozen at default values when constructing the metamodels (as shown in Table 6.1), and were not considered in the optimisation.

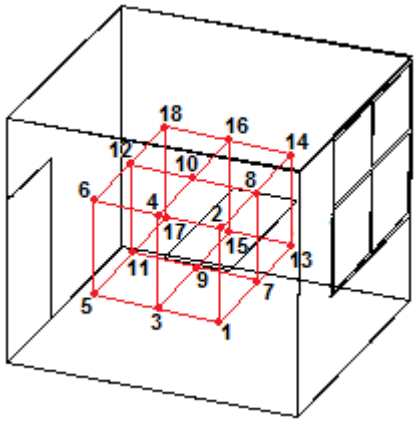

Figure 7.1: Comfort evaluation points in the model.

The second sample with 9 design variables investigated the effect of allowing the window to be open. Due the greater dimensionality of the problem, a larger sample was needed, so the number of samples in this set was doubled; 200 build samples and 100 validation samples. The design variables included were:

- Window control opening area ratios (3 variables) (0.1-11)
- Window control temperature set-points (3 variables) (18-28 °C)
- Radiant panel temperature (5-30 °C for summer cases, 20-40 °C for winter)
- Door discharge factor (0.1-0.9)
- Corridor air temperature (18-26 °C for summer cases, 20-28 °C for winter cases)

Results are presented in terms of Pareto fronts; these represent the predicted optimum trade-off between the two objectives. The theoretical basis of Pareto fronts is discussed in section 2.4.3, and practical guidance on interpreting them is given in section 5.3.1. In some cases, 3-D graphs of the Pareto fronts against design variables are given. Guidance on interpreting these outputs is also given in section 5.3.1, but is reiterated for convenience upon their first appearance in this chapter.

All design variables were scaled to ranges of 1-11 as described in section 6.4.1. Where particular design variable values are reported in the text, scaled values are first given, followed be actual values in parentheses. For example, a reported value of “9.530 (26.33 °C)” means the scaled value of that design variable for that point on the Pareto front is 9.530, which equates to a value of 26.33 °C for that design variable in the building model (this example is radiant panel temperature).

7.3 Spatial Variation of Optima

A key postulate of the present thesis is that as thermal comfort is known to vary spatially even within rooms, optimum conditions may also vary spatially. In other words, boundary conditions (heat input etc.) that achieve optimum thermal comfort at one location in a room may not necessarily achieve similarly optimum comfort in another location. In this section results are presented to explore differences in optimal conditions depending on where in the room the objective functions are evaluated.

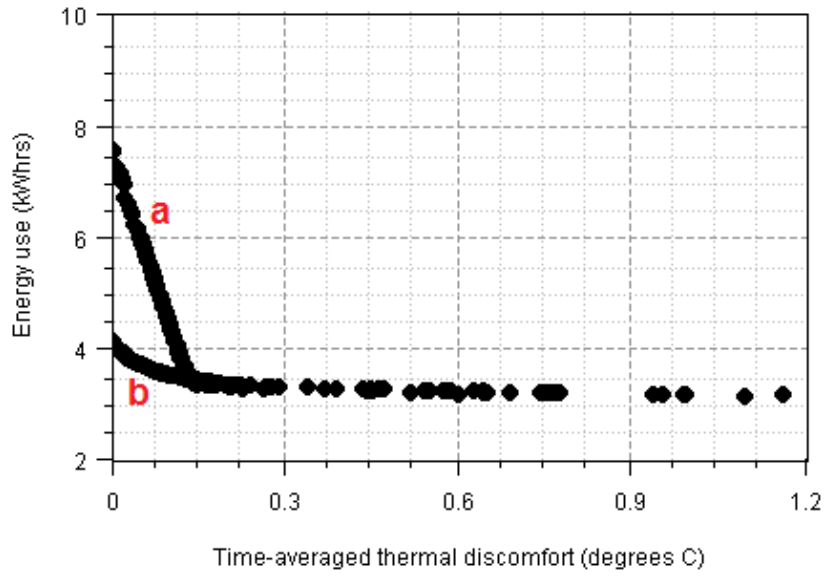

Figure 7.2: Pareto fronts for cases at two different positions in the room, one in the middle (a) and another nearer the window (b) (case references *SumAve1_9_1-24_dv5_ot* and *SumAve1_13_1-24_dv5_ot* respectively).

Position	Time-averaged thermal discomfort (°C)	Energy use (kWh)	Radiant panel temp. (°C)	Door discharge factor	Corridor temp. (°C)
9	0	7.539	26.33	0.112	19.53
13	0	4.146	24.83	0.1	21.54
9	0.146	3.353	22.88	0.1	21.86
13	0.153	3.424	23.18	0.1	22.13

Table 7.2: Showing solutions at 0 °C and 0.15 °C discomfort for two cases at different positions in the room, one in the middle (position 9) and another nearer the window (position 13) (case references *SumAve1_9_1-24_dv5_ot* and *SumAve1_13_1-24_dv5_ot* respectively).

Fig. 7.2 shows two Pareto fronts; one for a position in the centre of the room (position 9 in Fig. 7.1, “a” in Fig. 7.2) and one for a position near the corner of the room, in between the bed and

the window (position 13 on Fig. 7.1, “b” in Fig. 7.2). Both fronts are for cases with average summer weather (day 1), optimizing over the entire day, with a 5 dimensional design space (though design variables of external wall thermal conductivity and thermal storage capacity were fixed at constant values to remove them from the optimisation) using operative temperature comfort criteria. It is evident that the fronts have similar shapes with thermal discomfort above approximately 0.15, but differ greatly at values below. Values of energy use and the design variables at solutions giving 0 °C and 0.15 °C discomfort are shown in Table 7.2.These results suggest that a fully thermally optimum solution is quite different for the two positions. The solution at the centre of the room (position 9) has a far higher energy use than that of the solution near the window (position 13). This seems fairly counterintuitive as it would be expected that the temperature swing nearer the window would be wider. However results for a synonymous case on a different average summer day (day 2) show a similar trend, as shown in Fig. 7.3. The results in Table 7.2 suggest that the greatly increased energy use is due to an oppositional use of the two available HVAC systems; the corridor air is cooled whilst the radiant panel heats the room. This behaviour is typically associated with reducing temperature variation over the optimisation period, and it can therefore be surmised that temperature varies more in the centre of the room than next to the window. This is perhaps due to the increased influence of thermal mass in the external wall, or the prevailing air flow patterns at the time. More in-depth analysis would be needed to establish this, however such a counterintuitive outcome is nonetheless a good demonstration of the potential value of using T-BOT.

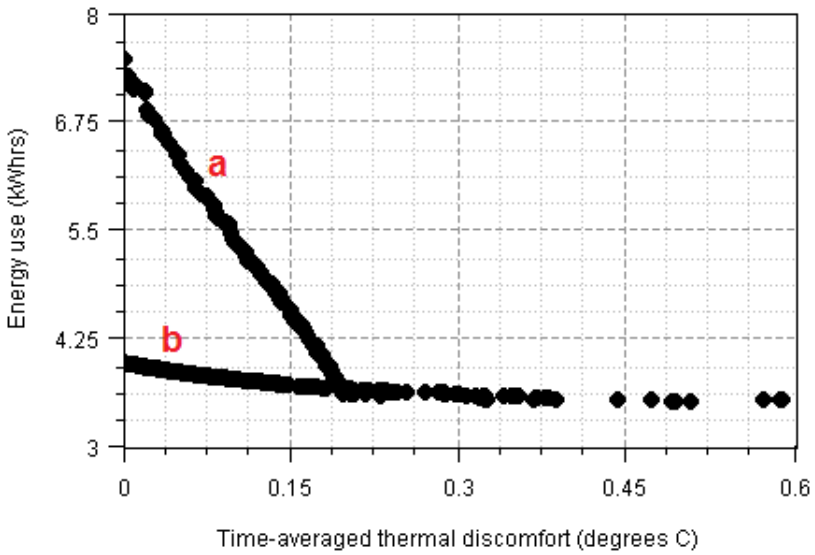

Figure 7.3: Pareto fronts for cases at two different positions in the room, one in the middle (a) and another nearer the window (b) (case references *SumAve2_9_1-24_dv5_ot* and *SumAve2_13_1-24_dv5_ot* respectively).

In practical terms, these results suggest that achieving optimum thermal comfort for position 13 is easier than for position 9. However the solutions at 0.15 °C thermal discomfort have very similar design variable values, indicating a solution that is more spatially consistent while allowing only relatively minor discomfort.

Fig 7.4 shows a similar example for cases with average winter weather (day 1); 7.4(a) is for a position in the centre of the room (position 9 in Fig. 7.1) and 7.4(b) is for a position near the corner next to the bed, opposite the window (position 17 in Fig. 7.1). Values of energy use and the design variables at 0 °C discomfort are shown in Table 7.3.

Position	Time-averaged thermal discomfort (°C)	Energy use (kWh)	Radiant panel temp. (°C)	Door discharge factor	Corridor temp. (°C)
9	0	2.809	23.06	0.310	23.92
17	0	2.871	23.65	0.36	23.96

Table 7.3: Showing solutions at 0 °C discomfort for cases *WinAve1_9_1-24_dv5_ot* and *WinAve1_17_1-24_dv5_ot*.

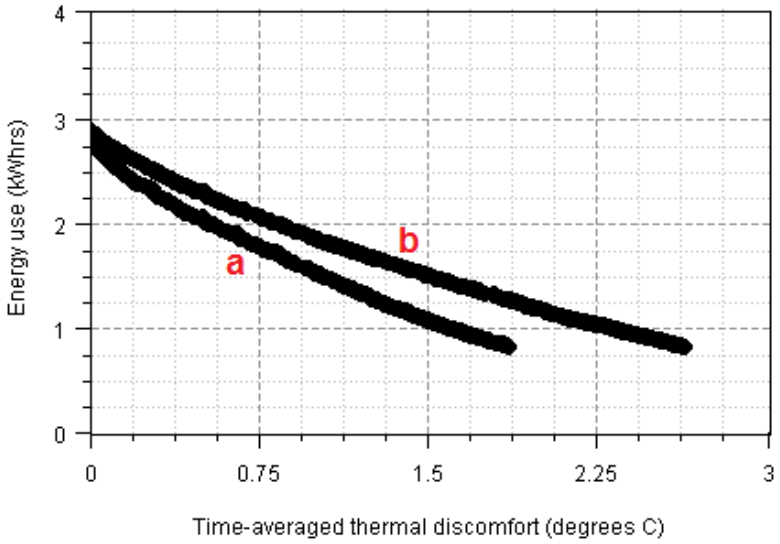

Figure 7.4: Pareto fronts for cases at two different positions in the room, one in the middle (a) and another next to the bed (b) (case references *WinAve1_9_1-24_dv5_ot* and *WinAve1_17_1-24_dv5_ot* respectively).

Comparing Figs. 7.4 and 7.2 it is evident that solutions in winter generally have lower energy use than those in summer. This could be a result of solar gains, as shading and blinds were not modelled, but it is likely that the greater oppositional use of the radiant and convective systems evident in summer results contributes significantly to this.

It is evident in Fig. 7.4 that the two fronts are more similar shapes than those of the average summer conditions results in Figs. 7.2 and 7.3. Solutions achieving 0 °C discomfort are very similar, but solutions with higher thermal discomfort take increasingly more energy to achieve in the corner of the room than in the centre. In contrast to the summer results, these suggest that a single solution may be able to achieve optimum thermal comfort at both locations.

7.4 Sub-daily Control Periods

This section demonstrates the ability of T-BOT to optimise over different time periods with a minimum of computational overhead, and explores the effect on results. This is done by comparing results for a whole 24 hour period and synonymous results for the same period broken up into four separate 6 hour periods.

Fig. 7.5 shows the Pareto front for a whole-day optimisation with average summer weather (day 1), at a position in the middle of the room (position 9 in Fig. 7.1), using the 5 design variable sample as described in section 7.2, with operative temperature comfort criteria. Fig. 7.6 shows Pareto fronts at each corresponding 6 hour period, and values of energy use and the design variables for various solutions are summarised in Table 7.4.

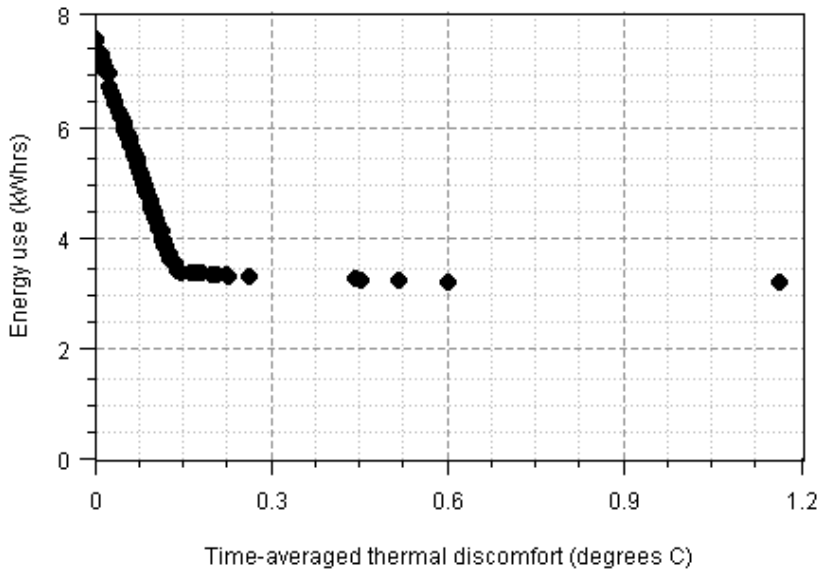

Figure 7.5: Pareto front for the 24 hour period (case reference *SumAve1_9_1-24_dv5_ot*).

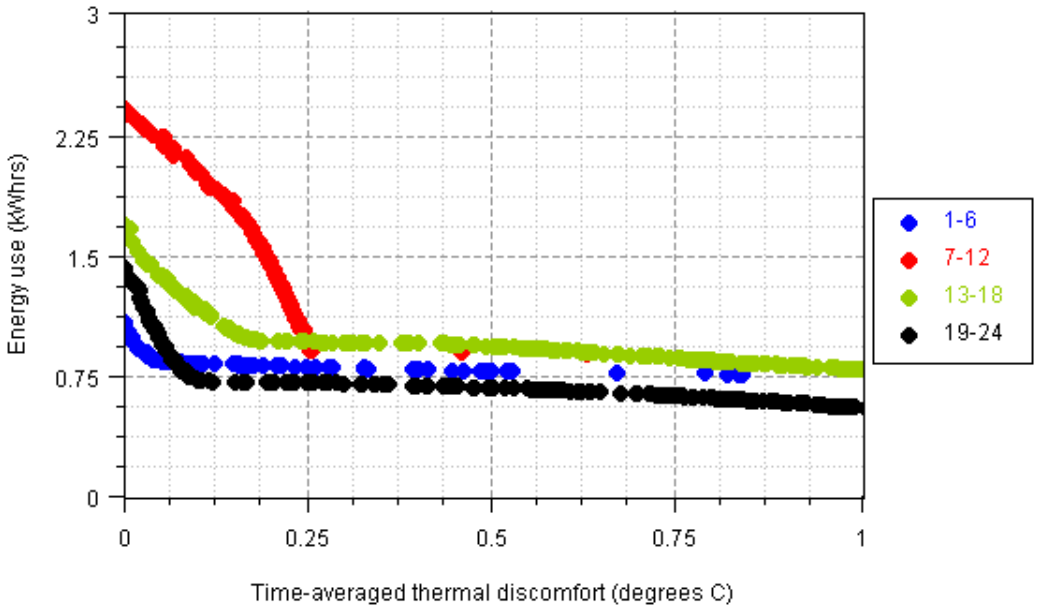

Figure 7.6: Pareto fronts for the 6 hour periods (case reference *SumAve1_9_1-6,7-12,13-18,19-24_dv5_ot*).

It is evident that of the curves on Figs. 7.5 and 7.6, that for the 7-12 period has a shape that is very different to the others'. This is unusual, and may be related to the relatively high solar gains during this period compared to others. Looking back at Fig. 6.11, it can be seen that direct normal solar intensity peaks in the period 6:00am – 12:pm on this day. From Table 7.4 it can be seen that solutions at 0 discomfort exhibit oppositional use of the radiant and convective systems, which is typically associated with attenuating peaks in room temperature. Such peaks could easily be caused by unshaded solar gains, particularly with the large east facing window on a sunny morning. This theory is supported by the oppositional use of the systems even in the 7 -12 solution at 0.2 discomfort, which suggests it is more difficult to retain a stable temperature during this period.

It may be calculated from Table 7.4 that the aggregate energy use for solutions of the four 6 hour periods at 0 °C discomfort is approximately 13% lower than for the 1-24 hour period, suggesting that the finer control allows for more optimal energy use in this case. Other results for summer conditions also exhibit this feature in varying magnitude, more so for extreme summer weather. However, the aggregate energy use of solutions at 0.2 °C discomfort for the 6 hour periods is approximately 26% higher than that for the 1-24 hour solution. Other summer results are not consistent in this respect.

In general, results suggest that under summer conditions, control periods (i.e. periods of constant set-points) of 6 hours provide more energy efficient operation than 24 hour control periods, likely due to the ability to adapt to dynamic prevailing conditions through the day,

rather than attempt to maintain a comfortable temperature all day with the same setpoints. However results suggest that this may not be the case if some discomfort is allowed; again this is fairly intuitive as a small amount of time-averaged discomfort is likely to accommodate fleeting peaks in temperature, and hence reduce or eliminate the need for the oppositional system use.

Period	Time-averaged thermal discomfort (°C)	Energy use (kWh)	Radiant panel temp. (°C)	Door discharge factor	Corridor temp. (°C)
1-6 (0:00-6:00am)	0	1.087	24.53	0.112	21.2
	0.208	0.815	22.66	0.101	21.29
7-12 (6:00am-12:00pm)	0	2.404	26.61	0.101	18.2
	0.2	1.426	24.1	0.2	20.65
13-18 (12:00-18:00pm)	0	1.688	26.32	0.103	19.37
	0.209	0.967	22.68	0.1	22.38
19-24 (18:00-24:00pm	0	1.414	25.92	0.121	20.68
	0.209	0.969	22.95	0.1	22.3
1-24 (0:00am-24:00pm)	0	7.539	26.33	0.112	19.53
	0.202	3.321	22.8	0.1	22.54

Table 7.4: Energy use and design variable values for solutions at certain values of thermal discomfort for average summer (day 1) 24 hour and 6 hour cases (case references *SumAve1_9_1-24_dv5_ot* and *SumAve1_9_1-6,7-12,13-18,19-24_dv5_ot* respectively).

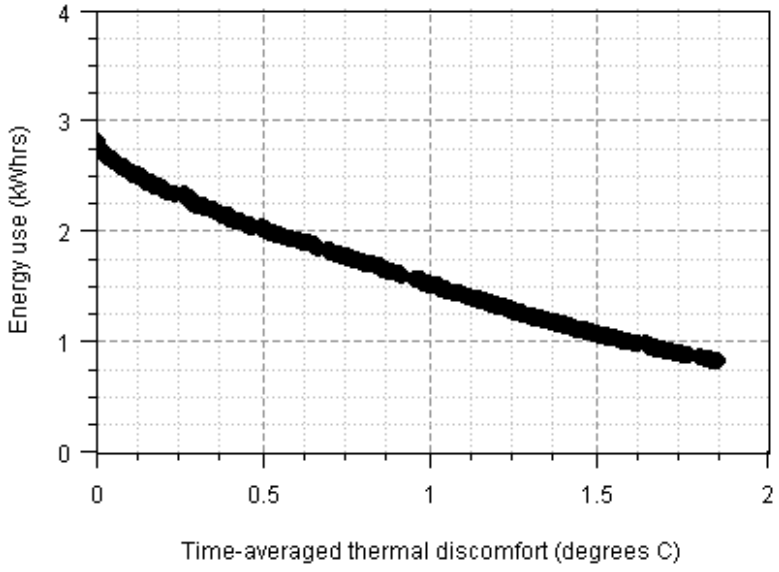

Figure 7.7: Pareto front for 24 hour period (case reference *WinAve1_9_1-24_dv5_ot*).

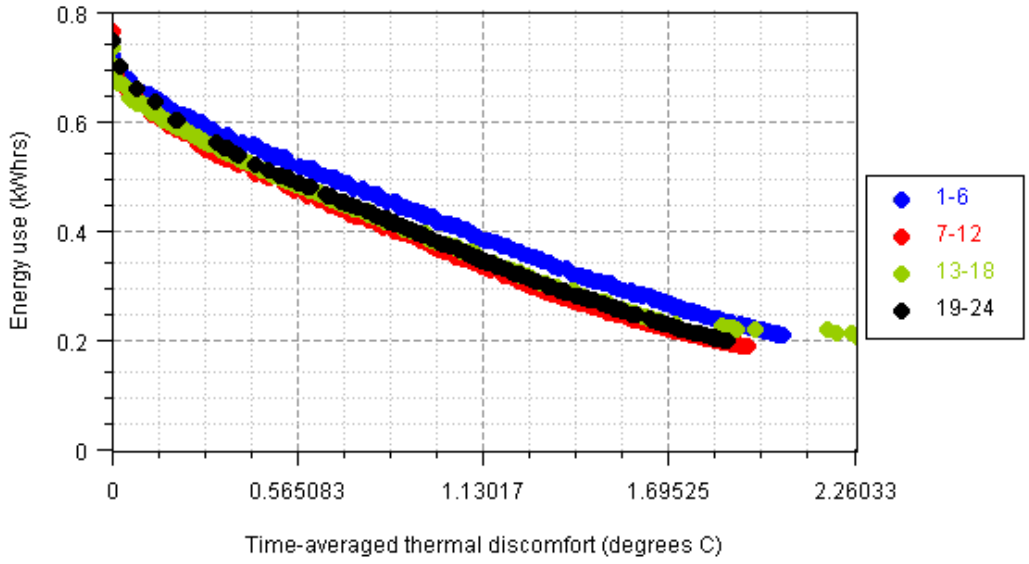

Figure 7.8: Pareto fronts for 6 hour periods (case reference *WinAve1_9_1-6,7-12,13-18,19-24_dv5_ot*).

Period	Time-averaged thermal discomfort (°C)	Energy use (kWh)	Radiant panel temp. (°C)	Door discharge factor	Corridor temp. (°C)
1-6	0.007	0.697	23.46	0.318	23.61
	0.204	0.613	22.46	0.178	22.92
7-12	0.005	0.688	23.18	0.318	23.80
	0.193	0.586	22.33	0.158	22.97
13-18	0.003	0.704	22.67	0.299	24.08
	0.200	0.590	22.30	0.238	22.92
19-24	0.003	0.715	23.34	0.318	23.76
	0.202	0.602	22.39	0.170	22.96
1-24	0	2.809	23.06	0.310	23.92
	0.199	2.384	22.30	0.155	22.98

Table 7.5: Energy use and design variable values for solutions at certain values of thermal discomfort for average winter (day 1) 24 hour and 6 hour cases (case references *WinAve1_9_1-24_dv5_ot* and *WinAve1_9_1-6,7-12,13-18,19-24_dv5_ot* respectively).

Figs. 7.7 and 7.8 show Pareto fronts for a similar case with average winter weather, and Table 7.5 summarises results in the same manner. It can be seen that the shapes of the Pareto fronts are more consistent in this case, and contrary to summer results the aggregate energy use of the solutions close to 0 °C discomfort for the 6 hour periods is 0.2 % lower than that for the 24 hour period, and the aggregate energy use of solutions at around 0.2 °C discomfort is

0.3 % higher. Clearly these differences are far less significant than in summer results, suggesting that under winter conditions finer control periods do not provide significant energy savings. This conclusion is supported by other winter simulation results.

7.5 Thermal Comfort Criteria

Within the current framework of T-BOT, limits of thermal comfort can be defined in terms of two different metrics. The default is operative temperature, but limits may also be derived from limiting values of PMV, by the method detailed in section 6.4.2. In this section results are presented to explore the effect of using these different thermal comfort criteria.

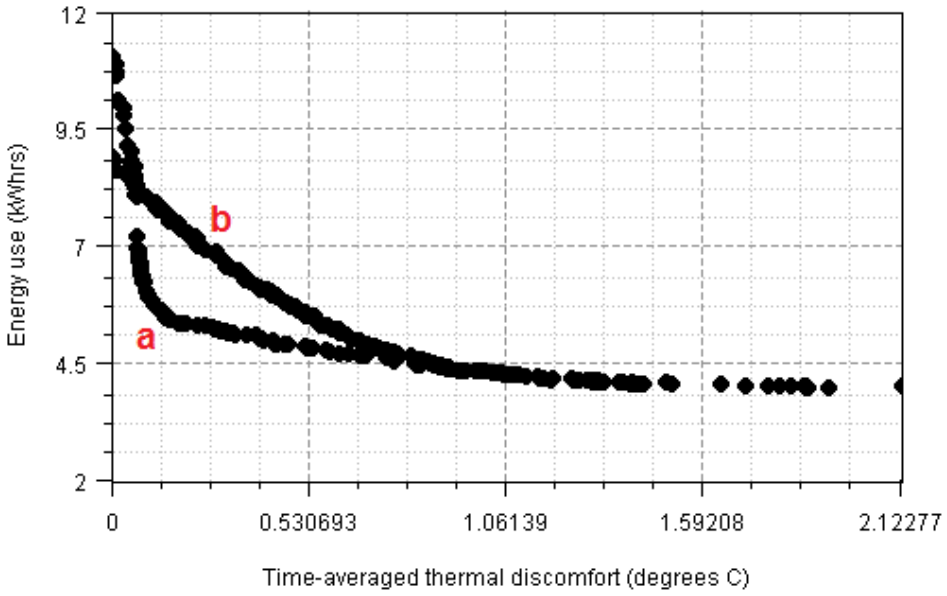

Figure 7.9: Pareto fronts for cases using operative temperature (a) and patient PMV (b) (case references *SumHot1_15_1-24_dv5_ot* and *SumHot1_15_1-24_dv5_pmvP* respectively).

Fig. 7.9(a) shows the Pareto front for a case with extreme summer weather, at position 15 (patient in the bed), optimising over the whole day, with 5 design variables, using operative temperature comfort criteria. Fig. 7.9(b) shows the Pareto front for an identical case, except that patient PMV comfort criteria are used instead. Values of 24 hour energy use and the five design variables for solutions at 0 and 0.5 °C discomfort are given in Table 7.6.

The results at 0 discomfort are somewhat counterintuitive, as the higher energy use of the operative temperature case suggests that thermal optimality using operative temperature criteria is harder to achieve than using PMV criteria. From Table 6.6 in section 6.4.2 it can be seen that the patient PMV criteria result in more stringent operative temperature criteria, with less than 1 °C comfort range rather than 2 °C as with operative temperature criteria. It would

be expected that more energy would be required to maintain the temperature within tighter limits, but these results suggest the opposite. However, results at 0.5 discomfort show the opposite trend, with the operative temperature solution having lower energy use than the PMV one. This could suggest that achieving very good comfort in the operative temperature case requires a greater expense of energy, but again this seems unlikely given the narrower comfort range of the PMV case. The discrepancy could be due to metamodel infidelity, considered in greater detail in section 7.7.

Comfort criteria	Thermal discomfort (°C)	Energy use (kWh)	Wall therm. conduct. (W/mK)	Wall storage props. (kg/m^3 & J/kgK)	Radiant panel temp. (°C)	Door discharge factor	Corridor temp. (°C)
Operative temp.	0	11.05	0.253	2000	27.23	0.898	19.12
Operative temp.	0.519	4.87	0.478	1991	24.41	0.9	21.949
Patient PMV	0	8.9	1.42	1987	24.63	0.9	21.9
Patient PMV	0.5	5.68	0.859	1997	24.65	0.9	22.55

Table 7.6: Showing solutions at 0 °C discomfort for cases with operative temperature and patient PMV representation of thermal comfort (cases *SumHot1_15_1-24_dv5_ot* and *SumHot1_15_1-24_dv5_pmvP* respectively).

Fig. 7.10 shows Pareto fronts for position 13, located between the bed and the window, with PMV calculated using the staff (a) and visitor (b) comfort criteria, defined in section 6.4.2. Values of energy use and design variables for solutions at 0 °C discomfort are shown in Table 7.7.

Comfort criteria	Thermal discomfort (°C)	Energy use (kWh)	Wall therm. conduct. (W/mK)	Wall storage props. (kg/m^3 & J/kgK)	Radiant panel temp. (°C)	Door discharge factor	Corridor temp. (°C)
Staff PMV	0	6.3	0.34	1990	20.7	0.9	20.3
Visitor PMV	0	8.38	1.45	1645	21.8	0.11	21.0

Table 7.7: Showing solutions at 0 °C discomfort for cases with staff and visitor PMV thermal comfort (cases *SumHot1 _13_1-24_dv5_pmvS* and *SumHot1 _13_1-24_dv5_pmvV* respectively).

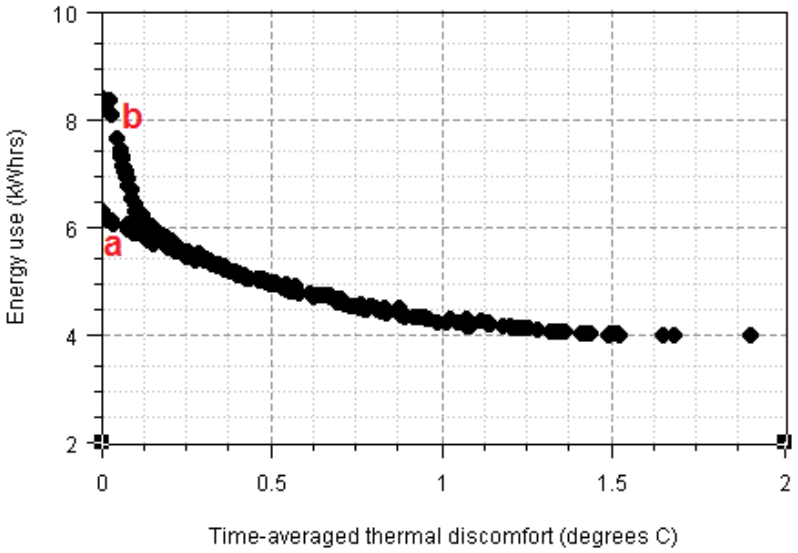

Figure 7.10: Pareto front for cases with staff PMV (a) and visitor PMV (b) thermal comfort (cases *SumHot1 _13_1-24_dv5_pmvS* and *SumHot1_13_1-24_dv5_pmvV* respectively).

From Table 6.6 it is evident that these two sets of comfort criteria are very similar; the visitor comfort range is approximately half a degree higher than that of staff due to the assumption that they will be wearing more clothes (coats, etc.). Nevertheless the results suggest that solutions achieving optimal thermal comfort in this case are quite different. A thermally optimum solution with staff criteria is predicted with significantly less energy use than with visitor criteria.

Finally, Fig. 7.11 shows Pareto fronts with extreme winter weather, at a position lying on the bed, optimising over the whole day, with 5 design variables (though two were fixed at constant values as detailed in section 7.1)), using operative temperature (a) and patient PMV (b) criteria. Values of energy use and design variables for the solution at 0 °C discomfort are shown in Table 7.8.

Comfort criteria	**Thermal discomfort (°C)**	**Energy use (kWh)**	**Radiant panel temp. (°C)**	**Door discharge factor**	**Corridor temp. (°C)**
Operative temp.	0	4.179	23.74	0.361	24.09
Patient PMV	0	8.808	21.47	0.755	27.3

Table 7.8: Showing solutions at 0 °C discomfort for cases with operative temperature and patient PMV comfort criteria (case references *WinCol1 _15_1-24_dv5_ot* and *WinCol1 _15_1-24_dv5_pmvP* respectively).

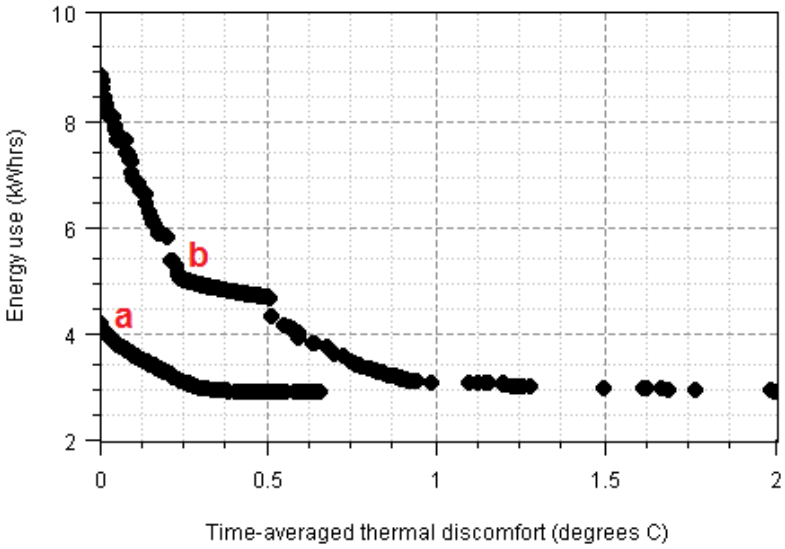

Figure 7.11: Pareto fronts for cases with operative temperature and patient PMV thermal comfort (case references *WinCol1_15_1-24_dv5_ot* and *WinCol1_15_1-24_dv5_pmvP* respectively).

Conversely to the summer results, the solution with PMV criteria has a much higher energy use than that with operative temperature criteria. This is more intuitive; the narrower comfort range of the PMV criteria would be expected to increase energy requirement. However the front for the PMV case has clear discontinuities, which can be an indication of poor metamodel fidelity (see section 7.7).

7.6 Design Variable Performance

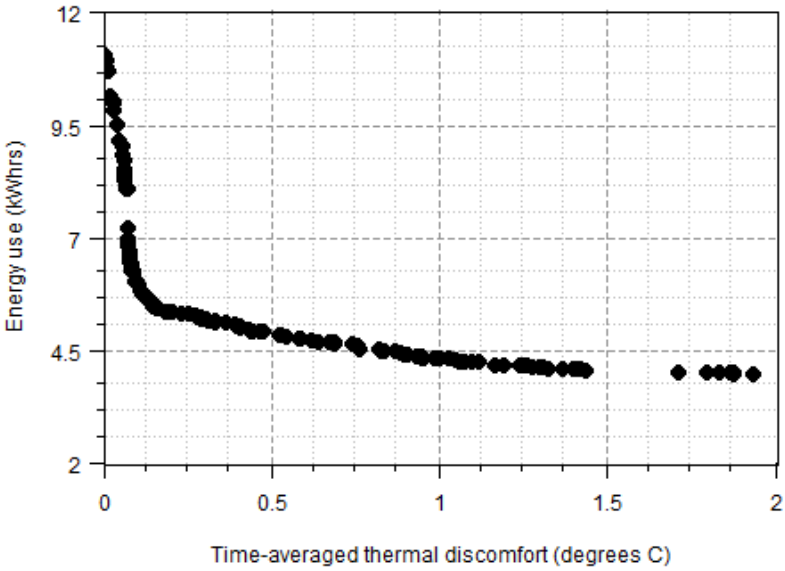

Figure 7.12: Pareto front for case *SumHot1_15_1-24_dx5_ot*.

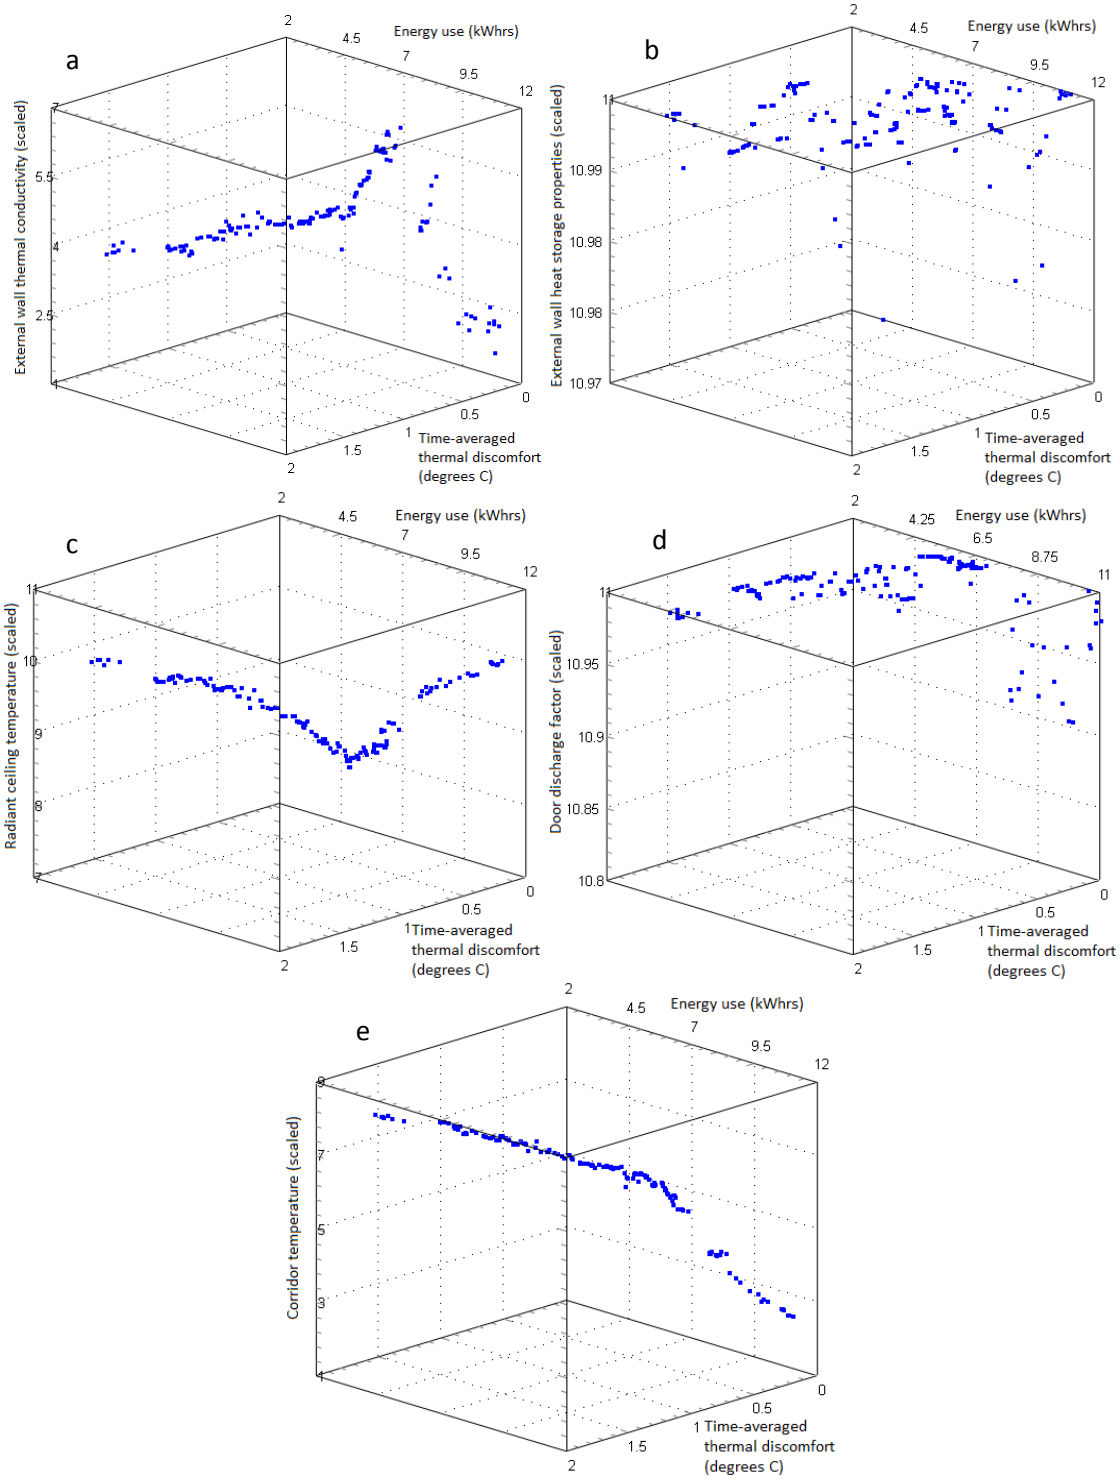

Figure 7.13: Pareto front (Fig. 7.12) plotted against each of 5 design variables.

It was evident from results for cases using the 5 design variable samples that the external fabric heat storage properties and door discharge factor design variables were generally less informative than others in the 5 design variable sample. In the majority of optimisations where both these design variables were considered, both converged to their limiting values, i.e. the entire Pareto front would lie on the edge of the design space in the axial direction of the design variable. This is illustrated in Figs. 7.12 and 7.13. Fig. 7.12 shows the Pareto front

for an optimisation for a patient lying on the bed, under extreme summer conditions, over a whole day, using the operative temperature thermal discomfort criterion. Fig. 7.13 shows the same Pareto front plotted against each of the five design variables. Guidance on interpreting these graphs is shown in Fig. 7.14; essentially the Pareto front should be projected onto the top of the cube, and coordinates can then be located in the 3D space by tracing the vertical distance.

Looking at the vertical axis scales in Fig. 7.13, it can be seen that those of Figs. 7.13(b) and 7.13(d) are both effectively zero-length, with values at the upper limiting value. In physical terms, this suggests that a thermally massive structure and as much air flow from the corridor as possible is beneficial in this case. High thermal mass can be desirable for passive cooling potential, and increased airflow improves the ability of the corridor to control the room (and possibly passive cooling from infiltration); both reasonable outcomes for an extreme summer climate case. However in terms of the optimisation procedure, having design variables converge to their limits is not desirable; this is discussed in detail in section 7.7.

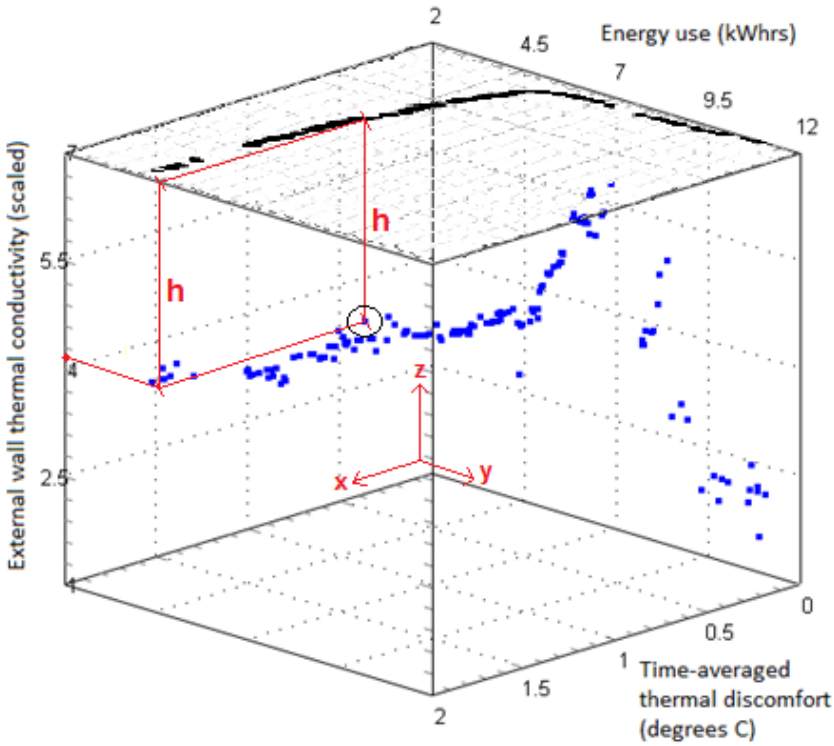

Figure 7.14: Three dimensional Pareto front vs design variable graph with Pareto front projected onto top face, showing guidance on reading z coordinates of points.

Also, whilst less inclined to converge to its limits over the entire Pareto front, the external wall thermal conductivity design variable very often converged near one of its limiting values for solutions close to 0 °C discomfort, as is evident in Fig. 7.13(a). In this case solutions are close to the minimum value of external wall thermal conductivity, indicating a poorly insulated

envelope. Examining Figs. 7.13(c) and 7.13(e), there is evidence of oppositional use of the two systems, the radiant panel with a high temperature and the corridor with a low temperature. This makes is somewhat difficult to assess the "reasonable-ness" of a poorly insulated envelope, as it is not immediately clear if net heating or cooling is being supplied, but it is generally accepted that well insulated envelopes are preferable in a UK climate.

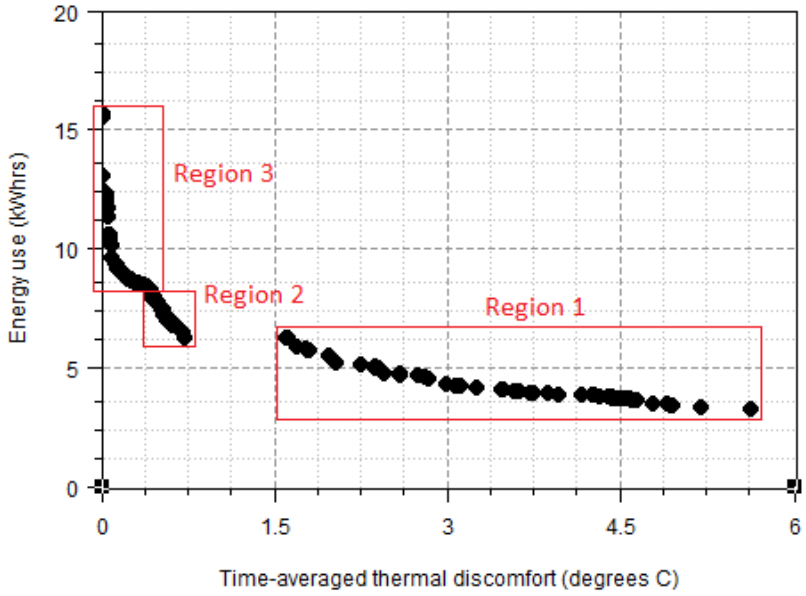

Figure 7.15: Pareto front for case *SumHot1_15_1-24_dv9_ot*, with regions marked corresponding to Fig. 7.16(a).

Cases with 9 design variables were simulated for some summer scenarios, as these included design variables representing a complex 4-stage window opening control algorithm. This was an attempt to intelligently optimise natural ventilation. However, the majority of these results suggested that the three-stage window control algorithm and 6 associated design variables (as detailed in section 6.3.2) generally did not perform as intended. Fig. 7.15 shows the Pareto front for a case identical to that shown above in Fig. 7.12, except 9 design variables were considered instead of 5, replacing the two external wall design variables with the six window control design variables. Fig. 7.16 shows this Pareto front plotted against each design variable. Corresponding regions are marked on Figs. 7.15 and 7.16(a).

Many of the window control design variables have values at one of their limits for the majority of solutions, and for most solutions the 3-stage window control algorithm is not used, i.e. the temperature lies entirely within control region at a given solution. The 6 design variables representing the algorithm then become invariant and add unnecessary dimensionality to the design space. This is reflected by the other three design variables shows in Figs. 7.16(g) – 7.16(i), which are generally fairly scattered over the Pareto front and exhibit clear discontinuities; this was typically a sign of poor metamodel fidelity as discussed in section 7.7.

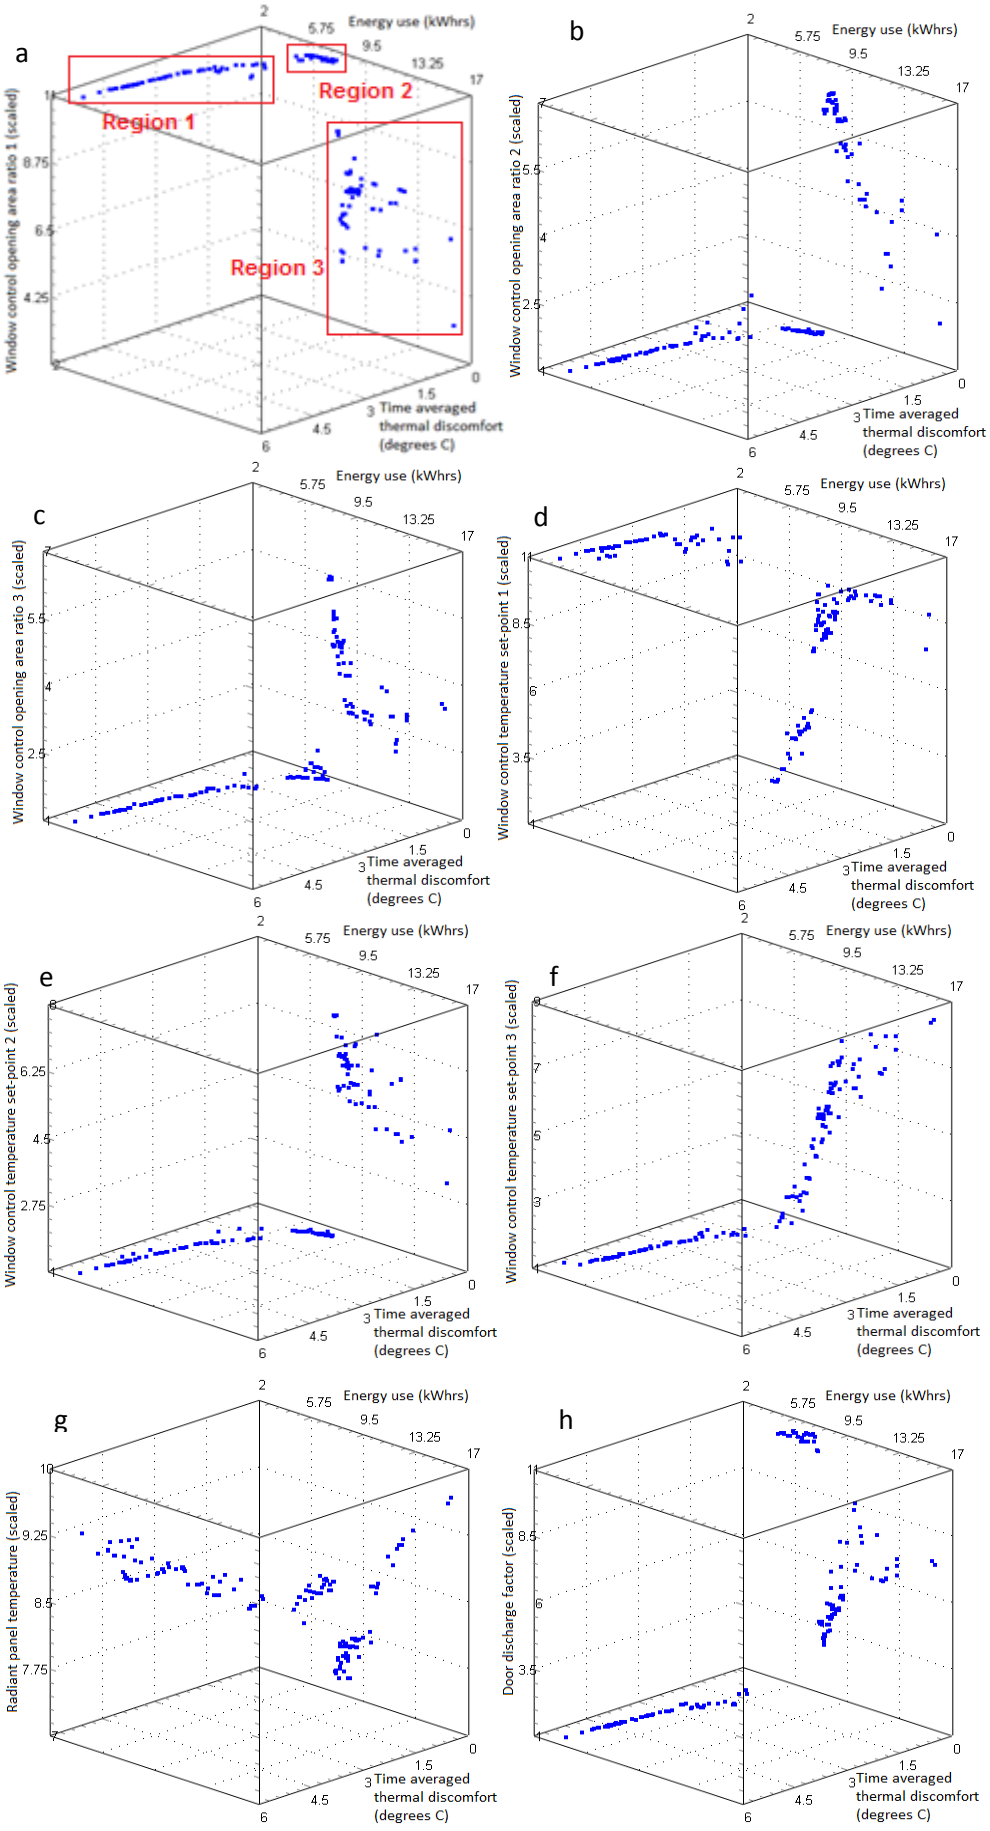

Figure 7.16, continued overleaf

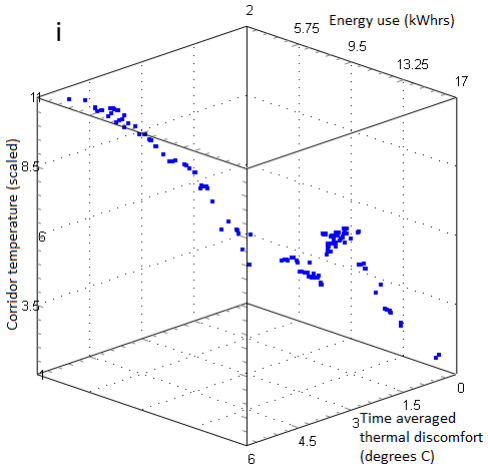

Figure 7.16: Pareto from for case *SumHot1_15_1-24_dv9_ot* (Fig. 7.15), plotted against each design variable.

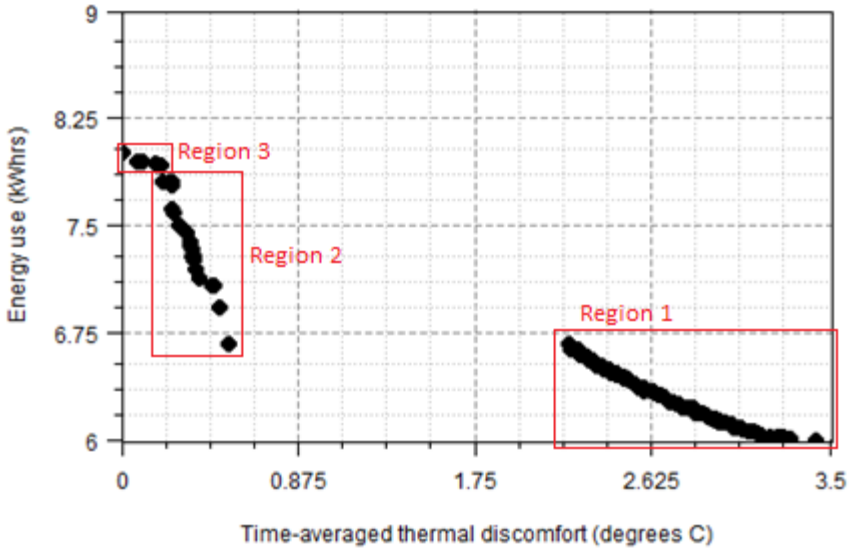

Figure 7.17: Pareto front for case *SumHot1_13_1-24_dv9_pmvV*.

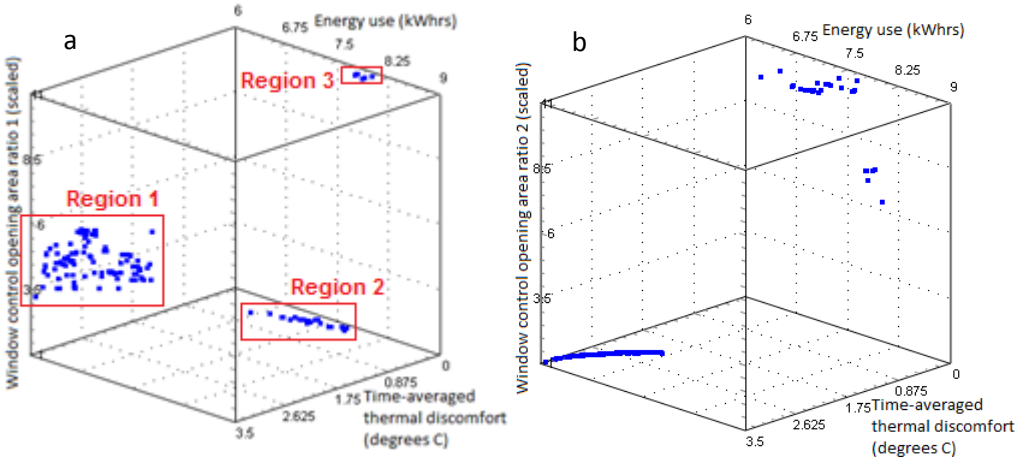

Figure 7.18, continued overleaf

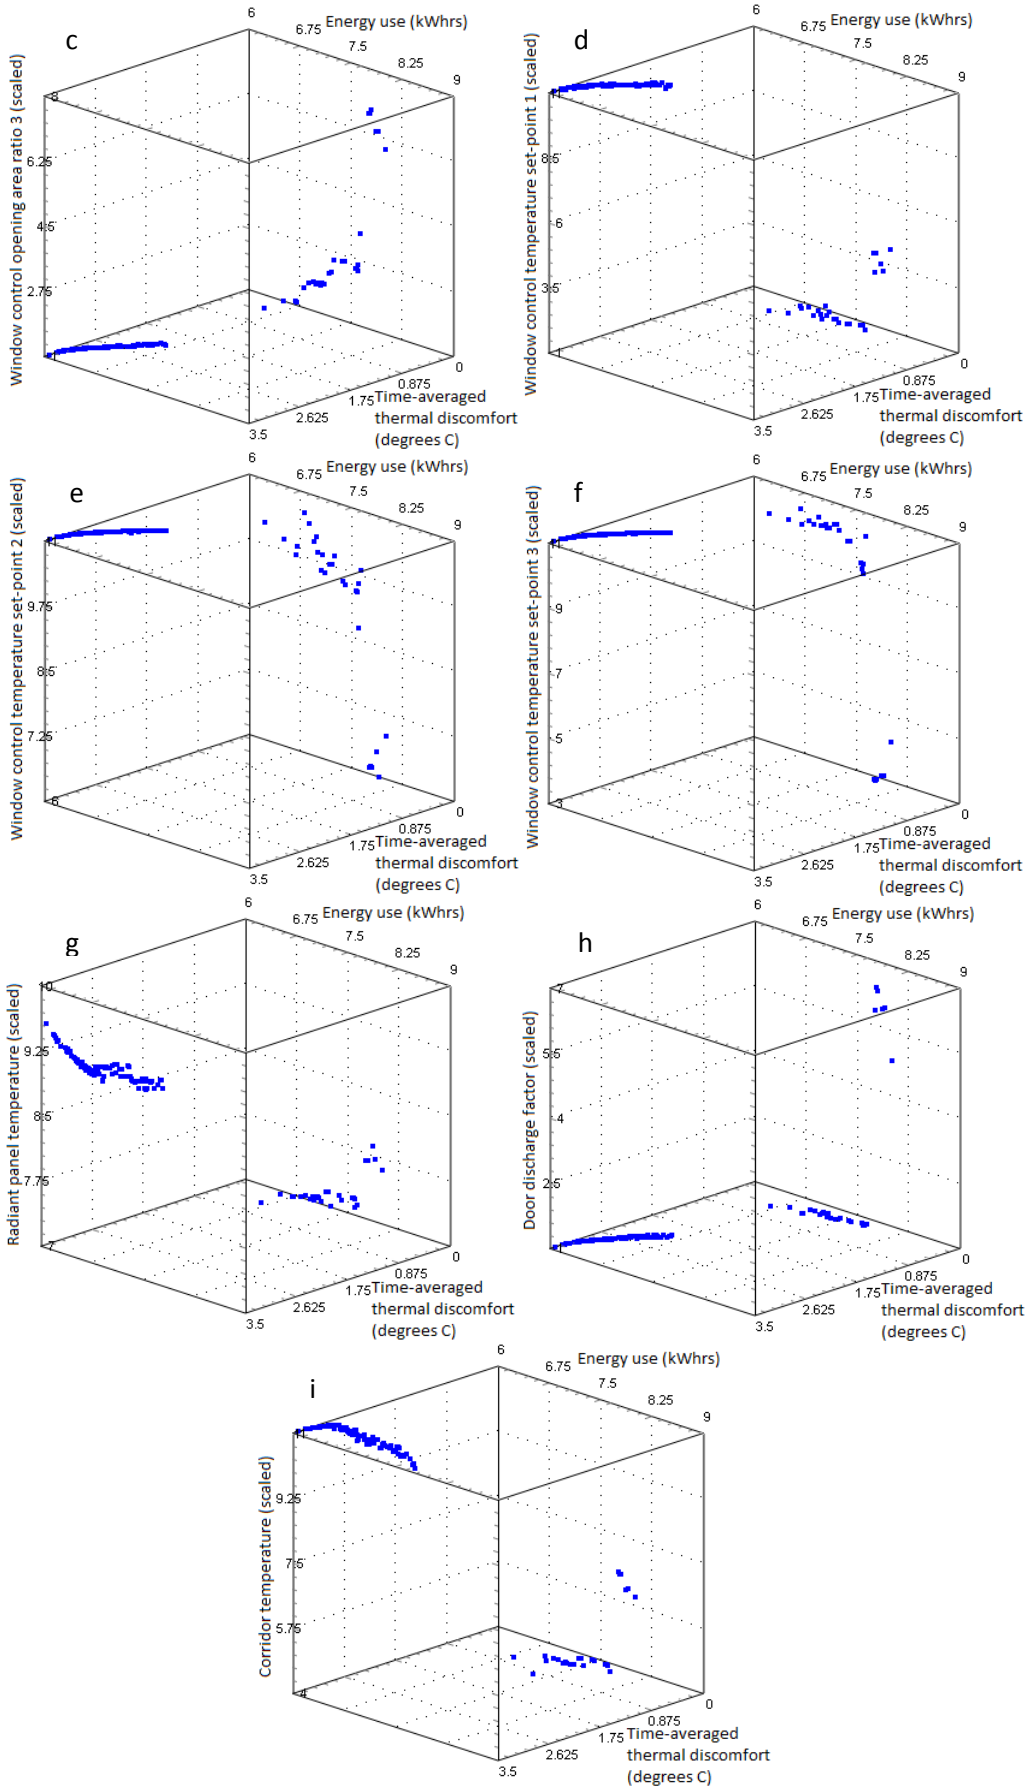

Figure 7.18: Pareto front for case *SumHot1_13_1-24_dv9_pmvV* (Fig. 7.17, plotted against each design variable.

Another example using PMV comfort criteria is shown in Figs. 7.17 and 7.18. This front exhibits even more extreme discontinuities , and solutions are similarly scattered in terms of all design variables. Table 7.9 shows a summary of the window control design variables in this case, corresponding to the regions shown in Figs. 7.17 and 7.18(a). It can be seen that in all regions the control system is not performing effectively, as the setpoints are such that not all control regions are used. Only in region 3 are the setpoint design variables spread in a manner that would be consistent with an effective multi-stage control algorithm, however as these solutions are those with the lowest discomfort, temperatures are unlikely to stray much below the low setpoint or above the mid setpoint; temperatures at low discomfort solutions do not stray significantly from the comfort range,in this case 21.7 – 24.9 °C as shown in Table 6.6. Whilst this could be an indication of the control system attempting to limit discomfort by changing window opening near the boundaries of the comfort region, the significant discontinuities of the Pareto front and the poor performance of the control system in the majority of cases make it difficult to give the benefit of the doubt in these situations..

	Variable	**Low**	**Mid**	**High**
Region 1	Temperature set-points	11 (28°C)	moot	moot
	Opening area ratios	widely spread in range 2-5 (1.19-4.46)	moot	moot
Region 2	Temperature set-points	1 (18°C)	10.8-8.7 (27.8-26.7°C); somewhat spread	moot
	Opening area ratios	moot	moot	moot
Region 3	Temperature set-points	spread in range 3-4 (21-22°C)	spread in range 6-7 (24.5-25.6°C)	spread in range 3-4.2 (25.2-26.4°C)
	Opening area ratios	moot	spread in range 6.3-7.4 (5.88-7.08)	moot
All ranges reported in order of decreasing thermal discomfort (unless otherwise specified), actual values shown in parenthesis.				

Table 7.9: Window control design variable performance for an extreme summer case using PMV comfort criteria.

The complex control behaviour encapsulated by this algorithm results in a relatively high number of heavily inter-dependant design variables. These results suggest that the additional

complexity this introduces into the design space, which has implications both on the initial sample (larger sample needed) and the optimisation (often superfluous dimensionality), is not a good price for the additional conclusions given.

7.7 Metamodel Fidelity

To explore metamodel fidelity, 20 solutions were selected from the Pareto fronts presented above and simulated with ESP-r. Solutions at 0 discomfort and another some way up the Pareto front were examined from each case. Direct comparison of actual results from the simulation with results predicted by the metamodels, provides a quantitative assessment of how well the metamodels perform in each case. Selected solutions are listed in Table 7.10, and comparison results are listed in Table 7.11 and shown in Figs. 7.19 and 7.20 for thermal discomfort and energy use objectives components respectively.

Case Reference	Figure	Solutions (X axis, Y axis)*	ID**
SumAve1_9_1-24_dv5_ot	7.2(a)	0.000, 7.54	1
		0.146, 3.35	2
SumAve1_13_1-24_dv5_ot	7.2(b)	0.000, 4.15	3
		0.153, 3.42	4
WinAve1_9_1-24_dv5_ot	7.4(a)	0.000, 2.81	5
		0.504, 2.01	6
WinAve1_17_1-24_dv5_ot	7.4(b)	0.000, 2.87	7
		0.502, 2.29	8
SumAve1_9_7-12_dv5_ot	7.6	0.000, 2.40	9
		0.200, 1.43	10
WinAve1_9_13-18_dv5_ot	7.8	0.000, 0.74	11
		0.200, 0.59	12
SumHot1_15_1-24_dv5_ot	7.9(a)	0.000, 11.05	13
		0.201, 5.35	14
SumHot1_15_1-24_dv5_pmvP	7.9(b)	0.000, 8.90	15
		0.206, 7.32	16
SumHot1_15_1-24_dv9_ot	7.15	0.000, 15.74	17
		0.507, 7.65	18
SumHot1_13_1-24_dv9_pmvV	7.16	0.000, 8.01	19
		0.482, 6.93	20
* Solutions are identified by their predicted objective function values, as shown in associated Figure.			
** Referenced in Tables 7.10 and 7.12.			

Table 7.10: Predicted solutions simulated to examine metamodel fidelity.

ID	Results	Cold thermal discomfort	Warm thermal discomfort	Total thermal discomfort	Radiant energy use	Convective energy use	Total energy use
1	Simulated	0.000	0.000	0.000	1.39	6.20	7.59
	Predicted	0.000	0.000	0.000	1.62	5.92	7.54
	Difference	0.000	0.000	0.000	0.23	-0.28	-0.05
	% Difference	-	-	-	16.2	-4.5	-0.7
2	Simulated	0.091	0.000	0.091	3.06	0.84	3.90
	Predicted	0.146	0.000	0.146	3.33	0.02	3.35
	Difference	0.055	0.000	0.055	0.27	-0.82	-0.55
	% Difference	60.4	-	60.4	8.9	-97.6	-14.0
3	Simulated	0.000	0.000	0.000	2.91	1.18	4.09
	Predicted	0.000	0.000	0.000	2.64	1.51	4.15
	Difference	0.000	0.000	0.000	-0.27	0.33	0.06
	% Difference	-	-	-	-9.3	28.0	1.5
4	Simulated	0.089	0.000	0.089	3.49	0.59	4.08
	Predicted	0.130	0.022	0.152	3.27	0.16	3.42
	Difference	0.041	0.022	0.063	-0.23	-0.43	-0.66
	% Difference	46.1	-	70.8	-6.4	-73.1	-16.1
5	Simulated	0.000	0.000	0.000	2.01	0.85	2.86
	Predicted	0.000	0.000	0.000	1.74	1.07	2.81
	Difference	0.000	0.000	0.000	-0.27	0.22	-0.05
	% Difference	-	-	-	-13.4	25.9	-1.7
6	Simulated	0.436	0.000	0.436	1.58	0.99	2.57
	Predicted	0.504	0.000	0.504	1.23	0.78	2.01
	Difference	0.068	0.000	0.068	-0.35	-0.21	-0.56
	% Difference	15.6	-	15.6	-22.2	-21.2	-21.8
7	Simulated	0.000	0.000	0.000	2.10	0.65	2.75
	Predicted	0.000	0.000	0.000	1.93	0.94	2.87
	Difference	0.000	0.000	0.000	-0.17	0.29	0.12
	% Difference	-	-	-	-8.1	44.6	4.4
8	Simulated	0.466	0.000	0.466	1.11	0.74	1.85
	Predicted	0.502	0.000	0.502	1.36	0.93	2.29
	Difference	0.036	0.000	0.036	0.25	0.19	0.44
	% Difference	7.7	-	7.7	22.5	25.7	23.8
9	Simulated	0.000	0.000	0.000	0.15	2.83	2.98
	Predicted	0.000	0.000	0.000	0.04	2.37	2.41
	Difference	0.000	0.000	0.000	-0.11	-0.46	-0.57
	% Difference	-	-	-	-73.3	-16.3	-19.1
10	Simulated	0.121	0.047	0.168	0.74	0.86	1.60
	Predicted	0.144	0.056	0.200	0.68	0.74	1.42
	Difference	0.023	0.009	0.032	-0.06	-0.12	-0.18
	% Difference	19.0	19.1	19.0	-8.1	-14.0	-11.3
11	Simulated	0.000	0.000	0.000	0.40	0.47	0.87
	Predicted	0.000	0.000	0.000	0.38	0.35	0.73
	Difference	0.000	0.000	0.000	-0.02	-0.12	-0.14
	% Difference	-	-	-	-5.0	-25.5	-16.1

Table 7.10, continued overleaf

12	Simulated	0.170	0.000	0.170	0.34	0.28	0.62
	Predicted	0.200	0.000	0.200	0.40	0.19	0.59
	Difference	0.030	0.000	0.030	0.06	-0.09	-0.03
	% Difference	17.6	-	17.6	17.6	-32.1	-4.8
13	Simulated	0.000	0.000	0.000	3.94	3.51	7.45
	Predicted	0.000	0.000	0.000	4.57	4.49	9.06
	Difference	0.000	0.000	0.000	0.63	0.98	1.61
	% Difference	-	-	-	16.0	27.9	21.6
14	Simulated	0.000	0.052	0.052	4.08	1.18	5.26
	Predicted	0.000	0.201	0.201	5.35	0.00	5.35
	Difference	0.000	0.149	0.149	1.27	-1.18	0.09
	% Difference	-	286.5	286.5	31.1	-100.0	1.7
15	Simulated	0.000	0.000	0.000	2.42	6.13	8.55
	Predicted	0.000	0.000	0.000	1.76	7.14	8.90
	Difference	0.000	0.000	0.000	-0.66	1.01	0.35
	% Difference	-	-	-	-27.3	16.5	4.1
16	Simulated	0.000	0.107	0.107	2.36	5.01	7.37
	Predicted	0.016	0.190	0.206	3.19	4.13	7.32
	Difference	0.016	0.083	0.099	0.83	-0.88	-0.05
	% Difference	-	77.6	92.5	35.2	-17.6	-0.7
17	Simulated	0.000	0.000	0.000	1.94	7.73	9.67
	Predicted	0.000	0.000	0.000	0.16	15.58	15.74
	Difference	0.000	0.000	0.000	-1.78	7.85	6.07
	% Difference	-	-	-	-91.8	101.6	62.8
18	Simulated	0.264	0.119	0.383	0.10	6.48	6.58
	Predicted	0.430	0.077	0.507	0.00	7.65	7.65
	Difference	0.166	-0.042	0.124	-0.10	1.17	1.07
	% Difference	62.9	-	32.4	-100.0	18.1	16.3
19	Simulated	0.000	0.041	0.041	2.74	5.35	8.09
	Predicted	0.000	0.000	0.000	5.07	2.95	8.02
	Difference	0.000	-0.041	-0.041	2.33	-2.40	-0.07
	% Difference	-	-100.0	-	85.0	-44.9	-0.9
20	Simulated	0.000	0.637	0.637	2.53	3.73	6.26
	Predicted	0.037	0.444	0.481	5.24	1.69	6.93
	Difference	0.037	-0.193	-0.156	2.71	-2.04	0.67
	% Difference	-	-30.3	-24.5	107.1	-54.7	10.7

Table 7.11: Simulated and predicted results comparison of objective function components for various solutions.

Looking at the thermal discomfort objective components, it is evident that where there is discomfort, generally it is overestimated by the metamodels. The only solutions where this is not the case are some with 9 design variables. However error for all objective components of these solutions is comparatively large, and they are discussed in more detail later. In practical terms, overestimation of comfort means the optimisation is likely to disregard some solutions that achieve optimum comfort. It is therefore possible that solutions achieving optimum

comfort exist with lower energy uses than those given in the Pareto front. It is however interesting to note that no solutions with predicted discomfort of 0, except one with 9 design variables, had any simulated discomfort. This suggests that the trend of overestimating of discomfort is fairly consistent and underestimation is not common among cases of 5 design variables. Overestimation of discomfort, particularly close to the optimum region, can be expected due to the metamodel smoothing inherent in the MLSR process. Wherever a discomfort objective component reaches 0, the nature of MLSR is to smooth this transition into a curve, even if the transition is very sudden and rises steeply. Underestimation of discomfort is more dangerous, as this could lead T-BOT to give solutions with actual discomfort greater than predicted, which in design terms would be rather more of a disaster than achieving comfort with less-than-optimum energy use. Error in the energy use objective components is not consistent in terms of over- and under-estimation.

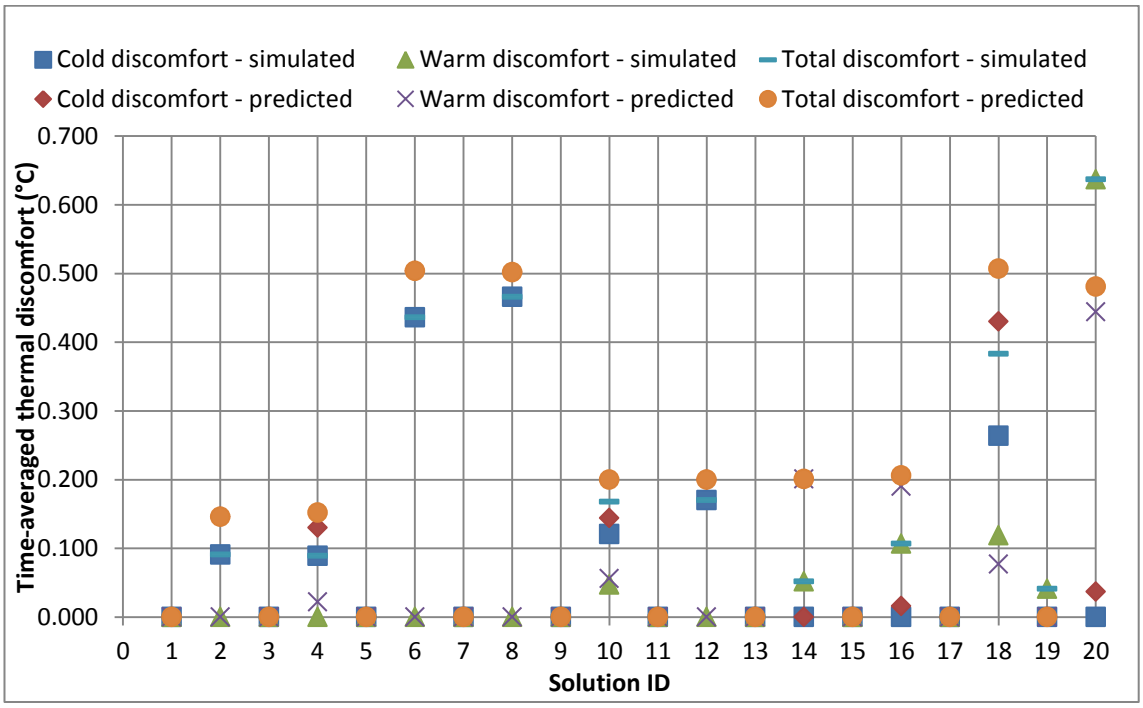

Figure 7.19: Simulated and predicted results comparison of thermal discomfort objective components.

The magnitude of errors seems to broadly correlate with certain variations in conditions. The lowest errors (i.e. the most accurate metamodels) are generally present in solutions with IDs 9-12, which are all from 6 hour period optimisations. Solutions 1-8, which are all 24 hour periods for various average climate conditions, generally have greater errors. Solutions 13-16, which are under extreme summer conditions with 5 design variables, generally have rather greater error than solutions 1-8, and solutions 17-20 with 9 design variables have the greatest errors. This is all fairly intuitive in context. The metamodels for the 6 hour optimisations have the least work to do, as it can reasonably be expected that aggregated or averaged building model

responses over 6 hours are likely to be more straightforward than over 24 hours; there is less variation that the aggregation or averaging procedure must encompass. It is therefore no surprise that metamodel fidelity is superior in these cases.

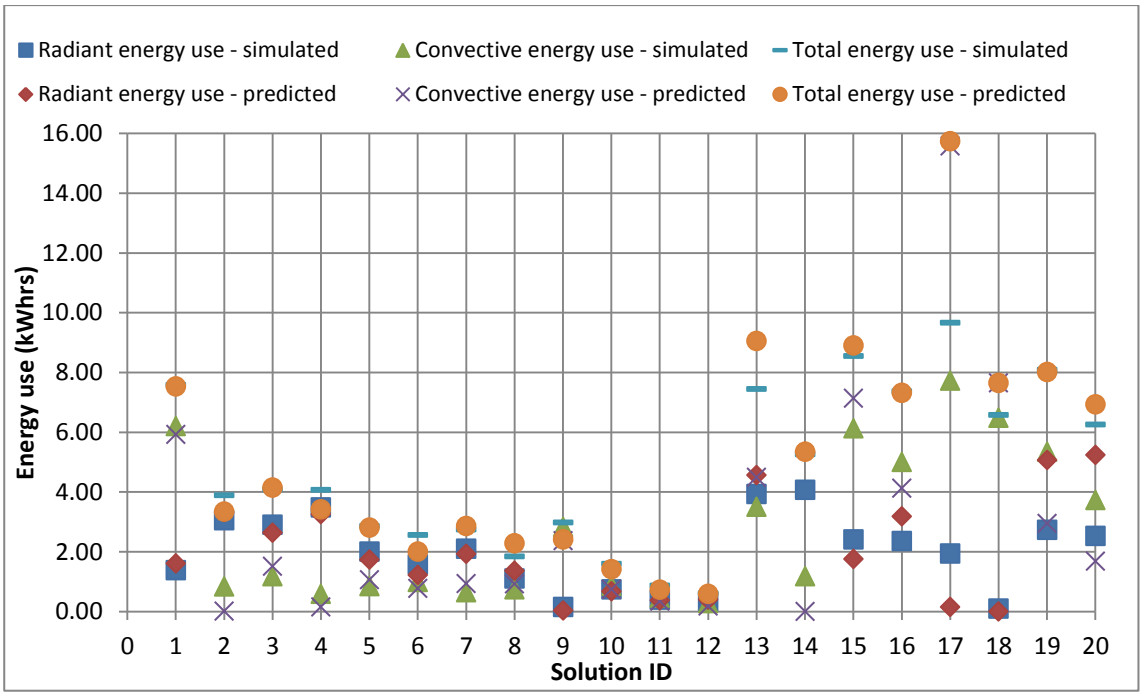

Figure 7.20: Simulated and predicted results comparison of energy use objective components.

Cases under extreme summer conditions with 5 design variables were optimised in terms of all 5 design variables. From analysis of these results it was then decided to run subsequent optimisations (with other climate conditions) in terms of only 3 of the design variables, two were held at fixed values (external wall thermal conductivity and thermal storage metrics) and were not considered variable in the optimisations. This was generally found to stabilise the Pareto fronts significantly, and this is reflected in the lower errors present in the solutions for cases other than extreme summer cases (solutions 13-20). Working with the 9 design variable sample instead of the 5 clearly further reduces metamodel fidelity substantially; the errors for solutions 17-20 are in some cases very large, and gave a clear indication that the window control algorithm and associated design variables had a significant negative effect on metamodel performance. Various metamodel accuracy metrics were extracted during metamodel generation for solutions 13 and 17, and are shown in Table 7.12. The most illustrative results are at the validation stage. It is clear that metamodels of the 9 design variable sample exhibit generally poorer fit to the sample data, though this difference is not as significant as may be expected from the large difference in metamodel fidelity identified by the comparisons above. The discrepancy between the two results sets could indicate that metamodel fidelity varies with the position in the design space; in the optimum region it may be local fidelity that is key rather than averaged metrics such as those shown in Table 7.12.

Sample set	Metamodelled component	Sample sub-set	R-Squared	Relative average absolute error	Maximum absolute error	Root mean square error
5 design variables	Cold thermal discomfort	Build	0.994	0.062	0.287	0.107
		Validation	0.979	0.110	0.591	0.185
		Merged	0.992	0.071	0.320	0.121
	Warm thermal discomfort	Build	0.997	0.039	0.206	0.056
		Validation	0.978	0.114	0.498	0.157
		Merged	0.996	0.049	0.233	0.067
	Radiant energy use	Build	1.000	0.008	0.439	0.084
		Validation	0.998	0.065	1.712	0.606
		Merged	1.000	0.012	0.468	0.118
	Convective energy use	Build	0.989	0.070	1.496	0.381
		Validation	0.885	0.274	3.358	1.299
		Merged	0.984	0.094	1.630	0.466
9 design variables	Cold thermal discomfort	Build	0.992	0.073	0.446	0.123
		Validation	0.976	0.121	0.619	0.206
		Merged	0.989	0.085	0.529	0.143
	Warm thermal discomfort	Build	0.996	0.050	0.157	0.062
		Validation	0.951	0.176	0.576	0.208
		Merged	0.992	0.071	0.252	0.086
	Radiant energy use	Build	1.000	0.015	0.703	0.140
		Validation	0.987	0.081	2.786	0.810
		Merged	0.999	0.022	1.214	0.215
	Convective energy use	Build	0.923	0.223	3.207	1.076
		Validation	0.816	0.326	4.817	1.651
		Merged	0.905	0.245	3.990	1.193

Table 7.12: Metamodel fit metrics for solutions of 5 and 9 design variable cases.

Table 7.13 shows design variable values for the selected solutions. For brevity only the scaled values are shown (all design variables were scaled to vales of 1-11 for the purposes of optimisation as detailed in section 6.4.1), as this provides a clearer assessment of the position of the solution within the range of each design variable. It can be seen that the number of design variables at their limits (close to either the lower limit 1 or the upper limit 11) also broadly correlates with metamodel fidelity. Solutions 1-8 typically have 1 design variable at or close to its limit, in most cases the door discharge factor. Exceptions are solutions 5 and 7, and it can be seen from Table 7.11 that these solutions exhibit relatively good metamodel accuracy. Solutions 10-12 do not have design variable values converged at their limits, and

show notably improved metamodel performance, whereas Solution 9 has two design variables at their limits and this is reflected by significantly poorer metamodel than other solutions for 6 hour periods. Solutions 13-16, with 5 design variables considered in the optimisation, all have at least two design variables converged at their limits. Solutions 17-20 from the 9 design variable sample are varied in this respect.

	Design variable values**										
ID	A	B	C	D	E	F	G	H	I	J	K
1	2.18*	5.44*	9.53	1.15	2.91	-	-	-	-	-	-
2	2.18*	5.44*	8.15	1.01	5.82	-	-	-	-	-	-
3	2.18*	5.44*	8.93	1.00	5.43	-	-	-	-	-	-
4	2.18*	5.44*	8.27	1.00	6.16	-	-	-	-	-	-
5	2.18*	5.44*	2.53	3.62	5.90	-	-	-	-	-	-
6	2.18*	5.44*	1.68	1.00	3.96	-	-	-	-	-	-
7	2.18*	5.44*	2.82	4.25	5.94	-	-	-	-	-	-
8	2.18*	5.44*	1.93	4.26	4.58	-	-	-	-	-	-
9	2.18*	5.44*	9.64	1.02	1.25	-	-	-	-	-	-
10	2.18*	5.44*	8.64	2.25	4.32	-	-	-	-	-	-
11	2.18*	5.44*	2.25	3.47	6.32	-	-	-	-	-	-
12	2.18*	5.44*	2.15	2.72	4.66	-	-	-	-	-	-
13	2.09	11.00	9.89	10.98	2.40	-	-	-	-	-	-
14	3.75	10.99	8.16	11.00	5.35	-	-	-	-	-	-
15	10.42	10.91	8.85	11.00	5.87	-	-	-	-	-	-
16	8.69	10.95	8.73	10.99	6.34	-	-	-	-	-	-
17	-	-	3.09	1.92	3.35	8.62	4.56	8.27	9.61	7.20	1.48
18	-	-	10.92	1.08	1.23	3.28	1.01	1.62	8.38	10.89	3.26
19	-	-	10.81	6.31	5.95	4.12	6.83	4.21	7.61	4.80	5.98
20	-	-	1.13	10.79	1.48	1.09	10.32	10.56	7.17	1.09	3.53

* Held constant during optimisation at default value shown.

** Design variables are assigned letters as follows:
A: External wall thermal conductivity
B: External wall heat storage properties
C: Radiant panel temperature
D: Door discharge factor
E: Corridor temperature
F: Window control opening area ratio 1
G: Window control opening area ratio 2
H: Window control opening area ratio 3
I: Window control temperature setpoint 1
J: Window control temperature setpoint 2
K: Window control temperature setpoint 3

Table 7.13: Design variable values for various solutions.

Generally speaking, a smoother Pareto front typically indicates superior metamodel fidelity. It can be seen throughout the figures in this Chapter and Appendix B that results of optimisations run with 3 design variables (most cases with climate other than extreme

summer) often exhibit smoother Pareto fronts than those run with all of the 5 design variables, or 9 design variables. As previously mentioned, cases with 9 design variables in particular often exhibit large discontinuities in the Pareto front, and it has been shown that these cases have comparatively poor metamodel fidelity.

To summarise this section, predicted results' deviation from simulated results due to metamodel infidelity is significant in many cases. In the case of thermal discomfort objective components, the metamodels tend to overestimate discomfort, and underestimation seems to be uncommon in cases with relatively good metamodel fidelity. Energy use components appear to exhibit less significant error in terms of the typical range of values that the objectives take among the solutions, though these errors are not consistent in sign. Results suggest that cases with smaller time periods and fewer design variables exhibit improved metamodel performance, even if design variables included in the initial sample are held invariant when constructing the metamodels in order to remove them from the design space. Finally, comparison of design variable values and metamodel fidelity suggests that accuracy degrades as design variables converge closer to their limits.

Chapter 8: Discussion

8.1 Chapter Overview

In this chapter the modelling carried out in this thesis is discussed as a whole to evaluate the performance and applicability of T-BOT. First the results from the simulations presented in chapters 5 and 7 are discussed, and conclusions are drawn in terms of model performance. Next, the methodology is critically assessed, highlighting strengths and areas that would be done differently with the benefit of hindsight. Finally the application of T-BOT is considered from a practical perspective, and key opportunities for further work are discussed.

8.2 Discussion of Results

In this study, numerous individual cases have been simulated that fall under four climate cases; extreme summer conditions, average winter conditions, extreme winter conditions and average winter conditions. As mentioned in section 7.1, in the interests of reasonable brevity these results are examined at a high level as opposed to on a case-by-case basis. This section makes reference to results by the case referencing system defined in Appendix B.1 and summarised in Table 7.1. Where the results are presented in Chapter 7 the figure number is stated.

8.2.1 Spatial variation

The summer results generally showed that spatial variation can strongly affect optimum conditions; what gives optimum performance in terms of thermal comfort at one location in the room does not necessarily give optimum performance at another location. The results suggest that this is only significantly true in summer and does not apply in winter; generally comparable solutions for different locations in winter cases were not as different as those in summer cases, particularly solutions close to 0 discomfort. The literature suggests that solar gains are a possible source of this seasonal discrepancy; for example Tzempelikos et al. [2010] demonstrated through a study on shading that solar gains had a significant effect on thermal comfort, and it is intuitive that these effects are likely to be spatially dependent due to the inclusion of per-surface, per-timestep solar incidence in the building modelling. Looking at Figs. 6.15 – 6.20, it is clear that summer climate cases used in the present work generally experience greater direct solar gains for a longer period. Solar shading or blinds were not

included in the model. This could greatly reduce solar gains in the space, which may reduce this seasonal discrepancy.

Metamodel prediction error could potentially be a factor in these conclusions. Thermal discomfort is often overestimated by the metamodels, meaning some solutions that achieve optimum discomfort may be discarded. It is therefore conceivable that more consistent solutions at different locations under summer conditions, have been discarded incorrectly. However the consistency of the conclusions noted above among the myriad cases examined in this work suggests that metamodel error does not play a significant role; it would be expected that such phenomena arising from error would be more inconsistent. Error in energy use objective components is typically of an order that would not significantly affect these conclusions (generally less than 0.5 kWhrs) for optimisations of 3 design variables.

Winter results from the proof of concept study (presented in Chapter 5) exhibited minimal spatial variation in optimum conditions, particularly for solutions close to 0 discomfort, which is broadly consistent with case study results. Mid-season results in the proof of concept study agreed with these conclusions of the winter results, though summer results were found not to be informative so this study did not allow a comparison of summer and winter conditions.

8.2.2 Thermal control

Winter results for the case study consistently showed that practically feasible zero discomfort solutions could be achieved with single set-points for whole 24 hour periods. On the other hand summer results clearly indicated that notable energy savings can be made by controlling the active systems in sub-daily time periods. In the proof of concept study, winter results also exhibited this trend, which is not consistent with the case study results. It is presumed that the sinusoidal climates used for the proof of concept study induced greater daily variation than is typical of real climates, which could explain this. The periodical differences in solutions were found to correlate in the majority of cases with periods of high solar gains on the days in question, which suggests that this phenomenon is also strongly associated with solar gains. Again, the consistency of these conclusions among the many cases studied suggests that they are not an artifact of metamodel prediction error. The proof of concept study results indicated that control periods of less than 6 hours do not give further useful information, though this may also depend on the details of solar gains in the space.

Climate and other boundary conditions vary time step-to-time step, and hence conditions in the room vary time step-to-time step. The longer the control period (with constant setpoints), the greater the variety of conditions under which the control must deliver comfort. It is easiest to rationalise this with discrete time steps. If each time step were optimised

individually, the conditions being optimised are effectively steady-state and hence an optimum solution is easy to obtain; indeed an approximate optimum could probably be obtained analytically under such conditions, as opposed to by simulation. If six time steps are optimised, say for example the period 00:00 – 06:00, this implies that for a thermally optimum solution, the control must deliver thermal comfort under six different conditions with the same setpoints. However in the period 00:00 – 06:00 conditions are not likely to vary by a huge amount; the sun may be just rising around 06:00 depending on the season, but solar gains are not likely to be great until a higher solar altitude is reached. If 24 time steps, i.e. a whole day, is optimised in one go then the control system, with constant setpoints, must deliver comfort under 24 distinct (but not independent) conditions which are likely to vary significantly. For example conditions at midnight will practically always be very different to conditions at midday. In other words the control must normalise conditions; it must warm in cold conditions and cool in warm conditions, all with the same temperatures. This is counter-intuitive by conventional control wisdom; you cool by injecting coolth as compared to the target temperature, and you heat by injecting warmth as compared to the target temperature. In the case study, this seems to be achieved by oppositional use of the two systems; cooling from one and heating from the other, simultaneously.

In a theoretical sense this is clearly an inefficient use of the systems from an energy perspective. However in a practical sense the situation is not implausible. In the building modelled for the case study, air supply in the corridor is presumably implemented largely in deference to ventilation guidelines for UK hospitals (e.g. DoH, 2007), and air is supplied at a constant temperature. Control in individual rooms is then achieved with radiant panels and windows. Particularly if the air supply temperature is not linked with ambient conditions, it is conceivable that the oppositional system use espoused by results from T-BOT could actually take place in reality. Such systems are not uncommon in legacy UK hospital building stock.

This demonstrates a clear advantage of multi-zone optimisation. Whilst the conclusion of supplying cold air into the corridor and then re-heating spaces with a radiant panel is clearly not desirable from an energy perspective, it nonetheless seems to be a valid conclusion that is potentially not far removed from the way some buildings (including the modelled building) operate. This suggests that programs such as T-BOT can help to identify potential for such issues at relatively early design stages. Furthermore the OSMO approach minimises the computational overhead of exploring such issues, as many different scenarios can be optimised from the same sample, including sub-daily control periods which results suggest can mitigate the need for this energy inefficient approach.

8.2.3 Objective functions

The limits of operative temperature calculated from patient PMV limits give a rather narrower comfort range than the limits recommended in CIBSE Guide A (CIBSE, 2006), and also do not depend on the season as there is no adaptive component to the PMV model. Comparison of cases under extreme summer conditions, with operative temperature and PMV comfort criteria (cases *SumHot1_15_1-24_dv5_ot* and *SumHot1_15_1-24_dv5_pmvP*; Figs. 7.9 (a) and (b) respectively), shows that for the thermally optimum solutions the energy use is approximately 20% lower for the *pmvP* case than the *ot* case. However metamodel fidelity was found to be markedly poorer in extreme summer cases likely due to design variables converging at their limits, and error in energy use component metamodels was up to around 1kWh, so this conclusion could be significantly affected by metamodel prediction error. Comparison of similar cases under extreme winter conditions (*WinCol1_15_1-24_dv5_ot* and *WinCol1_15_1-24_dv5_pmvP*; Figs. 7.11 (a) and (b) respectively) shows that for the thermal optimum solutions the energy use of *pmvP* case is more than double that of the *ot* case; an increase in energy use is more intuitive in this case as narrower comfort limits should necessitate tighter control. Given the fact that reducing diurnal temperature variation generally comes at a cost of increased energy use as previously noted, it may be surmised that the most optimal solutions will be obtained with the widest possible thermal comfort range (provided of course that the range is still valid), as these will allow the greatest diurnal variations whilst still maintaining comfort. However the PMV Pareto front has clear discontinuities which were found to be an indication of poor metamodel fidelity, which also casts a certain amount of doubt on these results.

In general, the PMV formulation of the thermal discomfort objective was found to produce more unstable Pareto fronts than the operative temperature formulation. However few cases using PMV were optimised in the present project, and only one was tried with 3 design variables (which gave better metamodel fidelity). The method used to characterise limits of PMV (detailed in section 6.4.2) resulted in them being implemented as limits of operative temperature, the same as other cases. It seems strange then that PMV results would exhibit such poor performance when the actual implementation is the same. The discrepancies are likely to be related to the comfort range of operative temperature, which varies significantly depending on the PMV case as shown in Table 6.6.

This aspect of BTO could benefit from further research in general; studies of different comfort criteria in a BTO context were not found in the literature. The closest found was the work of Bouchlaghem (2000) who compared a variety of thermal discomfort objective function formulations, concluding that mean deviation from a comfort temperature gave the best

performance. This work uses a slight variation of this formulation, replacing the comfort temperature with a comfort range, which is more appropriate for adaptive comfort criteria. A similar objective formulation was also used by Al-Homoud (1994), with the notable exception that the quantity was not time-averaged. The majority of BTO studies use PMV or PPD, for example Eisenhower et al. (2012), Wright et al. (2002), Xu et al. (2008), Pantelic (2012) and Chantrelle (2011) all used one of these two comfort metrics. It is worth noting that the OSMO approach could significantly reduce the computational overhead of such a study, as any comfort criteria could be optimised from a single sample (assuming requisite data was extracted from the sample) making T-BOT well placed to facilitate such an investigation.

In terms of the energy objective, a more comprehensive model that includes other zones connected to the corridor would allow a more deterministic metric of ventilation energy use, as the corridor could then be modelled in full and hence the energy input required to maintain it at the target temperature could be directly extracted from the model. This may reveal different tendencies in convective energy use than those suggested by the results with the present methodology (i.e. energy flow from corridor to room due to airflow).

8.2.4 Design variables

The optimisations for the extreme summer climate cases were performed including the external wall properties design variables where appropriate, these were not included in optimisations for other climate cases. This was due to the fact that the predictions given by the metamodels were found to be markedly less reliable when these two design variables were at their limits, which was often the case in optimal solution. Also, while it was hoped that by including these design variables in the optimisations valuable information pertaining to the use of variable-properties construction materials such as phase change materials (PCMs) may be gained, this was not the case. The lower degree of metamodel fidelity in the extreme summer cases and the poor consistency of the results casts a certain amount of doubt on the validity of the results. Far greater consistency was achieved in other climate cases where the external wall property design variables were omitted, and only 3 design variables were optimised. Eisenhower et al. (2012) also found that good results could be obtained optimising using metamodels over a subset of design variables included in the initial sample, though the principal was not linked to metamodel fidelity in their paper.

The use of the four-stage window control algorithm in the model was intended to allow the optimisation to converge to solutions where the window would be open wider during warmer periods of the day. Whilst this behaviour was evident in some solutions, for example the *SumAve1_13_1-24_dv9_ot* case, for the most part solutions converged entirely within one or

two of four possible control regions, and hence some of the 6 design variables used to represent this algorithm in terms of the optimisation became moot. This effectively adds superfluous design variables to the optimisation problems, which is highly undesirable as it increases the complexity of the metamodels needed to estimate the model responses. This is reflected in metamodel fidelity, which was found to be comparatively very poor for the 9 design variable cases that used the window control algorithm. In one case metamodel prediction error was found to be approximately half of the predicted value of the radiant energy use component. Clearly it is therefore difficult to draw conclusions from these results as there is significant doubt over their validity. Nonetheless, the general trend in directly comparable cases using the 5 design variable sample and the 9 design variable sample (i.e. those for the average summer climate cases; those for the extreme summer climate cases are not directly comparable as the *dv5* case includes external wall properties as design variables whereas the *dv9* case does not) does suggest that allowing the window to be opened does reduce energy use, as would be expected in summer due to the passive cooling that is typically induced. It may be concluded then that for the purposes of BTO, results suggest it is more useful to optimise the day in finer time periods to allow for changes in the window opening area than to include a more complex window control system and attempt to encapsulate the behaviour within time periods. This conclusion ties in nicely with that fact that windows are more likely to be open in summer when finer optimisation periods have been shown to provide useful information, though the principle may apply to any case with significant solar gains.

The notion of including window opening in a procedure such as T-BOT is worth further investigation, as such behavioural aspects of building performance are seldom considered at early design stages; however results of this work give a clear warning against complex characterisations. These issues also highlight the need to carefully consider the selection of design variables. Complex characterisation of intricate control mechanisms are not recommended due to the performance of the windows control algorithm in the present work. Furthermore, care must be taken to use design variables which significantly affect salient elements of the room performance, and that these design variables behave in a manner which is constructive and appropriate for optimisation. For example, the design variables pertaining to the building fabric in the present work did not prove particularly useful in the results as they were not constrained for behaviour corresponding with active building materials such as PCMs. Indeed in many cases these design variables were actually detrimental, as they often converged to their limits at optimum solutions where metamodel fidelity is not as reliable as demonstrated in section 7.7.

It could be useful to consider implementing a design variable sensitivity analysis step in the process as in the method of Eisenhower et al. (2012); this could help in a number of ways. Firstly and most obviously, it could ensure that the smallest possible number of design variables are optimised by removing design variables which are not likely to provide useful design information. Secondly if appropriate output and interrogation facilities were provided, it could allow evaluation of the range of possible values for design variables, which could inform limits imposed for the purposes of optimisation. Of course, including design variables in the initial sample that are then identified as superfluous is not desirable as it increases the dimensionality of the sample space and therefore decreases the coverage of any sample of a given size, and due to the computational requirement of models including CFD minimising sample size is clearly of paramount importance in terms of simulation time. For the purposes of T-BOT it could be worth investigating pre-sampling sensitivity studies.

8.2.5 Case study applicability to modelled building

Whilst the case study was primarily a vehicle for an in-depth exploration of the performance of T-BOT, it is worthwhile to summarise the applicability of the case study results to the real modelled building. This is largely a matter of how well the model represented the building, and how well the design variables reflected the capabilities of the building. Note that the validity of the results in absolute terms (i.e. metamodel fidelity) is considered as a separate issue, discussed principally in section 8.3.2.

There are number of aspects in which the case study did not reflect the modelled building well. Firstly, the use of constant heat gains is likely to be unrealistic. In practice occupancy, lighting use and equipment use are all likely to vary even through the course of a day. The model was set up with constant heat gains as little information was available on typical usage of these rooms, and hence an assumed profile with greater variation could not be developed with confidence. The impact of heat gains on results was not investigated, which could be an informative subject. This is discussed along with other further work in section 8.5.

Secondly the radiant ceiling was modelled as being able to provide cooling as well as heating, which is not realistic. In practice, supplying water at a temperature commensurate with cooling to such systems that are not designed for this, could cause issues such as condensation risk. As mentioned in section 6.2.2, this was done in response to outcomes of the proof of concept study, in order to provide greater freedom for the optimisation to find thermally optimum solutions, particularly under summer conditions.

Furthermore the lack of solar shading (e.g. blinds) in the model is a potentially significant source of disparity between the model and the real building, as discussed in section 8.2.1.

If T-BOT is used in practice in the future, clearly it is critical that the building model represents the real building as closely as possible. Hence it is worth noting at this point that the ability to characterise these issues is only related to the T-BOT methodology insofar as the limitations of the building simulation program. Currently T-BOT is configured to work with ESP-r, so the building modelling functionality of T-BOT is only limited by that of ESP-r. Considerations of using T-BOT in practice are discussed in detail in section 8.4.

In summary, applying practical conclusions from the results of the case study to the real building is heavily caveated. The model diverges from the real building in a number of important aspects. However, the exercise has highlighted the importance of accurate building modelling in BTO, which is not commonly considered in the literature.

As a final note, it is worth mentioning here that whilst the program was developed with functionality aimed at optimising hospital environments, a modular program architecture was established as shown in Fig. 4.18. Hence the process can readily be applied to other contexts by simply adding the appropriate design variable writer scripts. However, further investigation may be necessary to verify that program performance is consistent in other contexts.

8.3 Critical Evaluation

The methodology of T-BOT detailed herein seeks to provide an efficient framework to use highly computationally intensive building simulation for the purposes of detailed optimisation at a fine resolution in both space and time. In this respect, T-BOT is considered broadly successful. Any combination of location, optimisation time period and comfort criteria can be optimised with greatly reduced computational overhead compared to a direct search method (i.e. one that does not use metamodels) due to the OSMO approach; the principal that by saving primitive information from an initial sample of the design space, the sample becomes independent of such contextual variables. Values of these variables are then specified for optimisations, not sampling, which greatly expands the potential usefulness of a sample.

The following highlights key aspects of the methodology and discusses associated advantages and challenges.

8.3.1 Inclusion of CFD

In order to optimise at a local spatial resolution, values of state variables such as temperature and air velocity are required at different points in a room throughout the period of the simulation. Whilst the local radiant environment can be evaluated relatively simply by

calculating view factors for all surfaces surrounding a point, the local convective environment is far more difficult to evaluate. CFD is the current *de facto* method of accomplishing local air flow simulation. T-BOT uses the dynamic domain coupling functionality in the building simulation program ESP-r to link the various simulation domains, such that the outcomes of the building simulation (surface temperatures, radiant flux, etc.) become the boundary conditions for the CFD at each time step. The building and airflow domains therefore evolve in tandem. Whilst this is an efficient means for harmonious operation of the simulation, it does not nullify the principal disadvantage of using CFD, which is comparatively very high computational requirement.

Quantifying this in general terms, a week-long simulation without CFD would typically take a few seconds. When CFD is included, simulation time increases many-fold to minutes or hours. It has been demonstrated in section 4.3.2 that coarse-grid CFD gives an acceptable approximation of general flow patterns and temperatures obtained with more detailed simulation; use of coarse CFD was found to greatly reduce simulation time so this is highly recommended for BTO applications.

Practically speaking, the simulation time required by CFD constrains the simulation period, however this is linearly scalable to the resolution of a single simulation with parallelisation. For a 100 point sample the simulation process would lose no efficiency in terms of run-time if up to 100 separate processors were used. Given the increasing availability of vast HPC resources, in this respect the T-BOT methodology was considered to be well matched to future trends. Nonetheless, it is not considered likely to be practical to run full-year CFD simulations in a reasonable time on even the most advanced HPC systems in the foreseeable future. This necessitates that samples be run with representative weather data, and therefore it may be difficult to fully encapsulate seasonal variation and weather extremes in the design process. The approach taken in the present work was considered to be the most pragmatic way to work within this constraint; analyse the available weather data and extract the most extreme days and the most average days for summer and winter, and run samples under these conditions. Given that buildings in the UK generally tend to use their active thermal control systems (ie. heating and cooling) mainly during these seasons, hence consuming the most energy, it is considered appropriate to use T-BOT to inform design and operation primarily under these conditions.

However, it may be possible to further reduce simulation computational requirement by reducing the amount of CFD simulations required. Not all applications of T-BOT may warrant inclusion of CFD at all, for example cases where it is principally the radiant environment which is likely to induce spatial variation of thermal conditions. Also, it may be worth exploring the

possibility of only running CFD simulations at intervals of time steps, for example every 3rd or 6th time step. This functionality is currently outwith the capabilities of ESP-r, but the open-source model means it could viably be developed to allow such functionality.

8.3.2 Optimisation and metamodeling approaches

As noted by Eisenhower et al. [2012], the optimisation methods required to tackle large scale and many-faceted problems such as building optimisation typically require a large number of objective function evaluations. In direct search optimisation, each objective function evaluation is a building simulation. Hence as the computational requirement of each simulation increases, direct search optimisations become less feasible. Even a small increase in the time taken to run a single simulation can have enormous aggregate effects on the run-time of the process as a whole, due to the high number of simulations required. To include CFD in the building simulation is to push this caveat to extremes, as the computational requirement of each simulation is raised many-fold. It was therefore necessary to apply meta-modelling to reduce the number of simulations required, in order to achieve a feasible method.

In addition to enabling a feasible method, the metamodeling creates a disconnect between the simulation and optimisation functionality. In T-BOT this is exploited to allow optimisations to be performed over any time period, at any location, for any comfort criteria, from a single initial sample; this was termed a “one sample many optimisations” or OSMO approach. Although many researchers have successfully applied metamodeling to BTO e.g. Eisenhower et al. [2012], Xu et al. [2008], Gengembre et al. [2012] and Magnier and Haghighat [2010], to the author’s knowledge the OSMO approach is a unique feature of T-BOT. Also any subset of design variables in the initial sample can be optimised, though this is not a unique feature of T-BOT as it is also demonstrated by Eisenhower et al. [2012]. The computational requirement of the sample is not significantly increased to accommodate this functionality; only one CFD solution per time step is required for each sample point (though the coupling mechanisms of ESP-r result in multiple simulations being run in order to arrive at a time steps final solution). The OSMO approach is accomplished by saving a larger amount of more primitive information from the sample, which may then be compiled in different ways into data to train the metamodels. These economies add extra value to the T-BOT methodology when compared to a direct search method, above and beyond the implicit benefit of fewer objective function evaluations due to the metamodelling.

Consider a sample of a model representing an operating theatre run for 1 day. The sampling and metamodeling process has already reduced the time required compared with a direct

search method, by a significant margin. Instead of this sample data being used once to optimise the room and then discarded, many different options can be explored by the user. Results could be quickly obtained for a variety of surgeon positions and comfort requirements, perhaps influenced by different levels of metabolic rate. Optimisations could be done for the patient for a range of clothing indices, to see if it is more energy efficient to reduce heating and simply place a blanket over the patient. Different locations for the patient could be examined, to see if energy savings could be achieved by moving the operating table to the other side of the room. The whole day could be optimised, giving robust set points that should maintain a reasonable level of comfort during periods of similar climate. Each hour could be optimised individually to see if significant energy savings could be achieved with finer control periods. With this wealth of information, well informed design choices can be made.

However, the method is not without caveats. Clearly the accuracy of the metamodel predictions is paramount for the results of T-BOT to have any value. The metamodeling method currently employed by T-BOT (moving least squares regression, MLSR) is well suited to non-linear or numerically noisy data [Toropov et al., 2005], but has been found to exhibit a number of disadvantages in this context. Firstly, unlike some other methods (e.g. Kriging) the metamodel responses do not necessarily pass through every sample point [Gilkeson et al., 2013]. Given that every sample point is "correct" within the context of the study (it is not suggested that building simulation is automatically true-to-life) this can be a severe disadvantage. The problem is particularly acute in cases where at least one objective function component has 0 value over a significant portion of the design space. This can introduce sharp changes of gradient into the response surface, which MLSR does not model well; the method is continuous and tends to smooth such discontinuities. Hence where there should be a sharp fall to 0 in the response surface, the MLSR approximation will generally exhibit a smooth curve instead. A safeguard to ensure non-negative objective component values was implemented in the methodology, which went some way to mitigating this, however over-smoothing of metamodel responses close to 0 values was still found to be a consistent issue. This reduces metamodel fidelity in the region; the metamodel gives non-zero predictions for solutions which should have a value of zero as demonstrated in section 7.7. This is a particularly undesirable tendency as optimum solutions are likely to lie in these regions.

For example, consider the thermal discomfort criterion used herein. This was calculated by summing two separate components, warm discomfort and cold discomfort, which were metamodelled separately. An example of typical response surfaces for these is shown in Fig. 8.1(a). It is evident from this figure that the optimum solutions in terms of thermal comfort lie along the valley formed where the two components' response surfaces cross one another; this

region is close to where the cold discomfort surface moves away from 0. If the curve in this region were to become more of a sharp angle, it could significantly change the position of the optimal valley; similarly so for energy use With optima being sensitive to relatively minor features of the metamodelled responses, it could be beneficial to investigate other metamodelling methods in future research; this is discussed further in section 8.5.

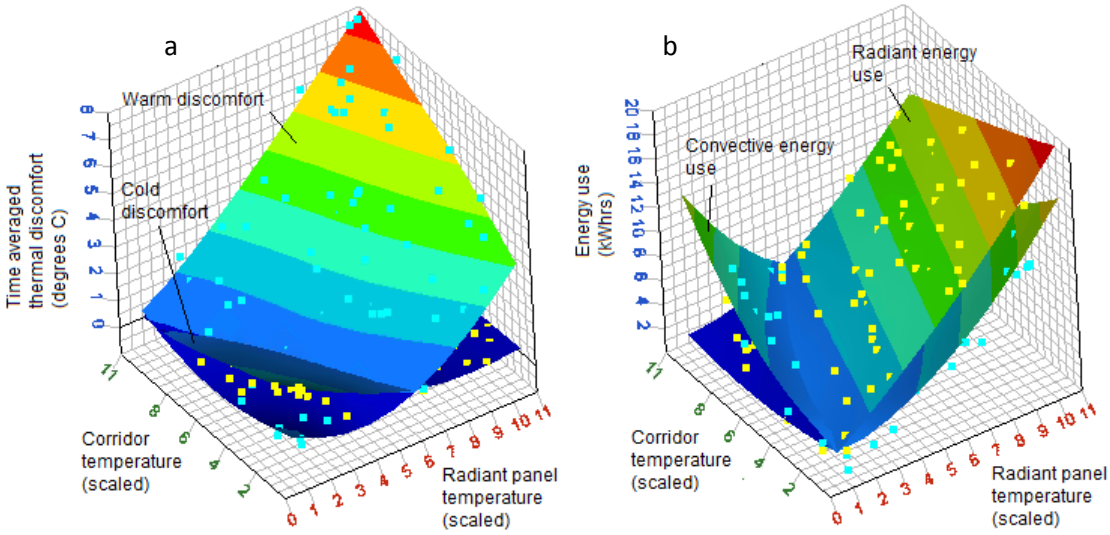

Figure 8.1: Examples of metamodel responses, showing thermal discomfort objective components (a) and energy use objective components (b). Points are sampled data.

8.3.3 Design space

The design space is finite, being bounded by the limiting values of each design variable. It is assumed that only solutions inside the design space are feasible solutions, and therefore sample points are not placed outside the boundaries. At the boundaries of the design space metamodel fidelity has been shown to degrade; this is presumably due to there being fewer nearby sample points as there are no sample points on the unfeasible side of the boundary. It is therefore desirable to keep solutions away from the edges of the design space wherever possible, which has implications on selection of design variables and their limits. For example design variables and their limits could be selected such that an optimum solution is highly unlikely to lie on any boundary due to extreme limiting values. However, this then increases the size of the design space and hence increases the area must be sampled. Furthermore, design variables that converge to their limits are potentially not as informative as they could be. As a consequence of converging to its limits, the design variable often becomes invariant over the solutions comprising the Pareto front; it effectively becomes constant and therefore adds unnecessary dimensionality to the design space. The more design variables that are included, the more likely it is that such issues will arise and adversely affect results.

Such issues are often collectively termed the “curse of dimensionality”; the fact that each design variable, and hence each dimension of the design space, exponentially increases the complexity of the design space. A key issue with this is that this increases the required size of the sample proportionately. As the majority of the runtime for the process is taken by the initial sample, reducing the size of the sample is clearly beneficial. Hence the method presented in this thesis is better suited to cases of few design variables, which aligns well with the intended remit of detailed optimisation of single spaces. Of course in theory the method could be applied to larger areas, even whole buildings, though the task of characterising an entire building in terms of CFD would be daunting at best. The literature supports this conclusion, as other BTO studies are typically performed with many more design variables; Eisenhower et al. [2012] used up to 1009, Wright et al. (2002) used up to 200, Al-Homoud (1997) used up to 16. Results from the present study suggest that 3 design variables is appropriate for a sample size of 150, though of course this depends on the design variables and their ranges as discussed, and further developments to the T-BOT methodology may improve this.

8.4 Use of T-BOT in Practice

This section briefly examines the potential uses of T-BOT in practice. First though it is necessary to point out that further work is needed before T-BOT should be used beyond research. Many of the issues discussed in previous sections (e.g. metamodel fidelity, sensitivity of solutions) as well as practical issues of usability remain to be addressed. These are discussed in section 8.5.

The results indicate that T-BOT cannot be regarded as a black box system; it is not designed to spit out a single solution and tell the user “this is the best design”. Rather the benefits of the program are as a support tool for informed decision making, providing a flexible and time-efficient information gathering framework. Also, adopting an optimisation framework provides convenient nomenclature and already well-defined processes for examining the combined influence of multiple parameters.

The underlying rationale for this is the premise that design should always be a human-led process. With currently available technology, no artificial intelligence can fully replace the quality of human judgement based on practical experience. This may manifest in fairly trivial ways, such as discarding a solution which advocates a radiator right next to a door, which may restrict the opening of said door. On the other hand human judgement might be the difference that averts disaster; for example discarding a solution which advocates a high-temperature radiator next to a flammable furnishing. The automatic systems cannot hope to

be fully aware of any and all practical implications of its design choices. However, whilst they may not be able to replicate human judgement, automated systems are exceptionally good at other parts of the design process. Generating and simulating building models for example, can be accomplished automatically in a fraction of the time that a human would need. It is therefore proposed to use BTO as a tool, rather than a replacement, for human-led building design.

The T-BOT system was designed as an approach that could ultimately fulfil this role in a practical setting. The vast majority of the time required to use T-BOT is in simulating the initial sample of solutions; optimisations then performed using the sample data are much quicker. This splits the workload associated with the use of T-BOT into two distinct phases. First, the initial "seed" building model must be constructed with all required design variables represented appropriately therein. T-BOT must also be configured to work with this model, which currently involves coding bespoke design variable writer scripts (as shown in Fig. 4.16) among other things, though it is anticipated that future iterations of T-BOT could significantly streamline the configuration functionality. The initial sample can then be run, which may take anywhere from hours to weeks depending on the complexity of the model and the computational resources available. It is worth noting that it is not considered likely that an undertaking like this that includes CFD will ever be feasible on a desktop computer; it is assumed that HPC resources will be available to run such initial samples.

Given the level of time and computational effort that is put into the T-BOT process, it is worthwhile to briefly consider the type of application. Clearly this kind of process is not appropriate for every room in every building ever designed; for example it is probably not necessary to optimise every single bathroom in a high rise office block. Firstly, there is decades of world-wide precedent to inform the design of this very common type of room. Secondly many of these rooms are likely to be of similar size and shape with similar furnishings, and whilst the position in the building may influence the exact details of solar incidence for example, such variation is not likely to significantly affect even locally evaluated metrics. Finally whilst it is appreciated that achieving efficiency and comfort in building design is an intrinsically distributed problem, the fact remains that such rooms are generally less important areas. Bathrooms in commercial buildings are not typically sources of high energy use, and are usually only occupied briefly and intermittently. Therefore, it seems reasonable to expect applications of T-BOT to be focussed on two broad types – archetypes and critical spaces.

Application of T-BOT to room archetypes could be useful in both design and research contexts. For example a design context might see a user exploring the optimum control period in a

typical office room to ensure comfort may be robustly maintained in a particular building. A research context might see a user exploring the benefits of a wallboard impregnated with active materials such as PCMs. However, application in this context must be done with care. For example consider the design context suggested above; optimisation of an office room archetype. It would not be prudent to consider solar shading as design variable(s) in this case, as solar incidence angles will vary depending on the exact position of the room in the building as already mentioned. Whilst this is not likely to have a large impact on general metrics such as thermal comfort and energy use, it is clearly critical to the operation of solar shading. In this instance it may be better to use separate models for each storey of the building in the case of horizontal shading, or at regular intervals along the façade for vertical shading.

Application of T-BOT to critical spaces could be particularly useful in design and policy contexts. For example T-BOT could be applied to the design of an operating theatre, ensuring the HVAC systems available are able to provide robust maintenance of thermal comfort for surgeons, support staff and patients. Similarly, T-BOT could be applied to an operating theatre archetype to try to generalise this design information into design guidance that could then inform standards and policy, though of course this would be subject to the same caveats as any archetype study. In the present work T-BOT was developed and applied with hospital bedrooms in mind; these may also be considered critical spaces as they typically accommodate the most vulnerable among the population.

8.4.1 Further potential outputs from T-BOT

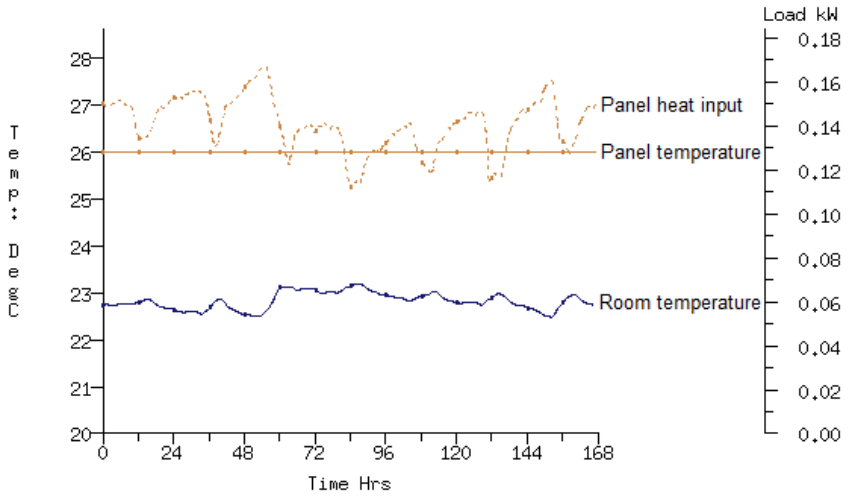

Figure 8.2: Example of profiles extracted from a 1 week simulation of a thermally optimum winter solution.

This section briefly presents examples of further potential outputs from T-BOT that may be useful in practice.

Firstly, it is worth highlighting one aspect of the T-BOT methodology that has not been well explored in the present thesis. Each point on the Pareto front is in theory a design solution, so it is likely to be useful in practice to examine these design solutions outside the remit of the optimisation. For example, temperature and energy use profiles and CFD visualisations could be used to assess model performance in far greater detail than is encapsulated by the objective function and design variable values.

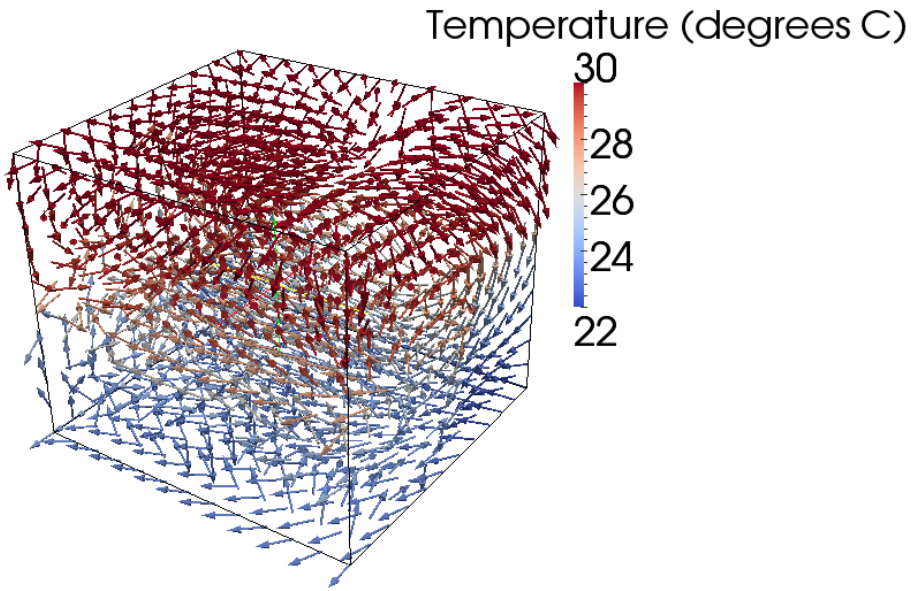

Figure 8.3: Example of a CFD visualisation extracted from simulation of a solution with high panel temperature.

The design variable values of the solution effectively give a "mapping" which enables the user to configure the building model, automatically or manually, to give the simulated performance characterised by the objective function values (subject to the accuracy of the metamodels). Such subsequent simulation is not limited by the caveats of the T-BOT methodology in terms of run-time, and hence is not constrained to one-day periods for example. Fig. 8.2 shows an example of profiles extracted from a 1 week simulation of the case study model for a thermally optimum winter solution; it can be seen that the room temperature does not vary from the comfort range of 22-24 °C, and the panel is able to maintain its target temperature for the entire week. Fig. 8.3 shows an example of a CFD visualisation from a winter solution with a high panel temperature; it can be seen that there is significant vertical variation of air temperature. These conclusions may not be apparent from the objective function and design variable values given by T-BOT, and hence demonstrate the value of this potential.

As well as investigating model performance, design variable sensitivity and metamodel fit to the sample data can also be examined in greater detail using outputs from HyperStudy. Fig. 8.4 shows an example of the sensitivity of energy use components to the door discharge factor

design variable. In this case it can be seen that whilst variation of other parameters results in significant spread of the data, generally speaking lower values of the design variable results in lower energy use. Fig. 8.5 shows an example of predicted vs. observed values for a metamodel of the warm discomfort objective component.

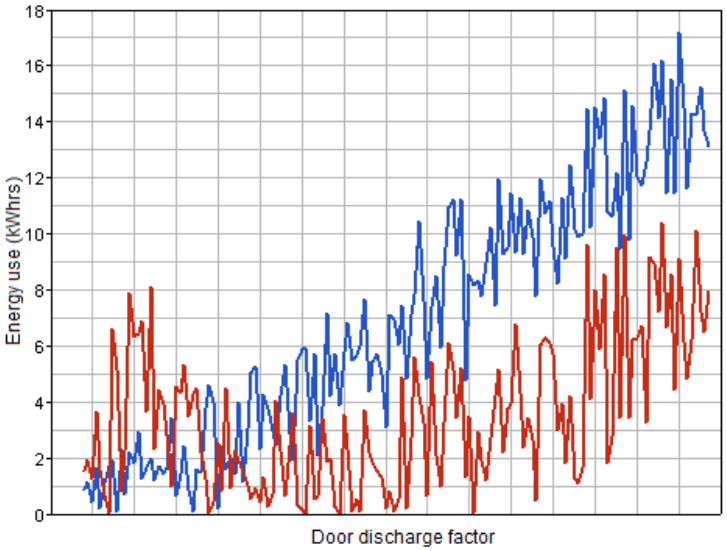

Figure 8.4: Example of sensitivity of radiant (blue) and convective (red) components of energy use to the door discharge factor design variable.

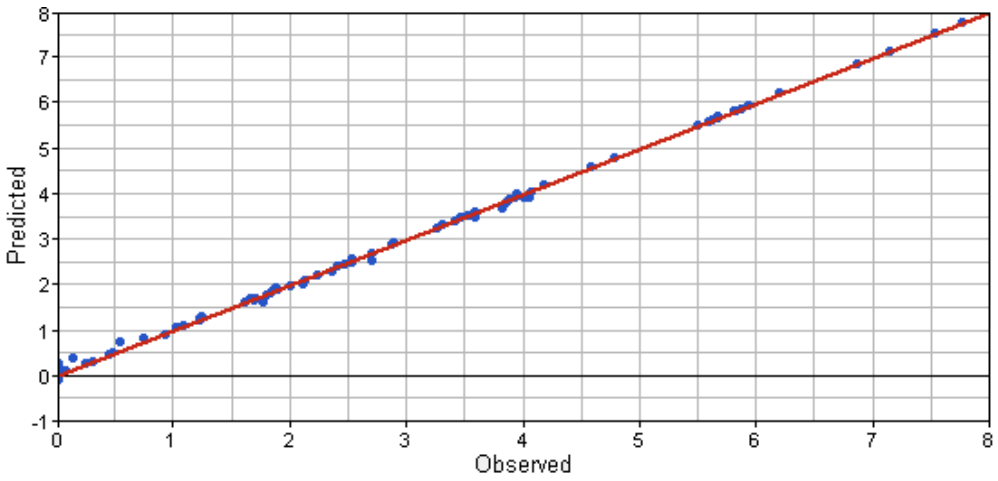

Figure 8.5: Example of predicted vs observed values for a warm discomfort metamodel.

8.5 Further Work

This section draws on previous discussions and summarises key opportunities for further work to develop T-BOT towards becoming a usable tool in practice.

First, results indicate that thermal comfort criteria are a primary concern and strongly affect the results given by T-BOT. Many human thermal comfort models exist in the literature (e.g. Zolfaghari and Maerefat [2010a], [Zhang et al., 2010a, 2010b and 2010c] and Al-Othmani et al.

[2008]), both empirical and theoretical, but few have been effectively integrated into building simulation software, especially more advanced modern examples. Given the high level of modelling detail undertaken in the T-BOT methodology, it could be worth exploring whether the process would benefit from a detailed, comprehensive thermal comfort model that takes into account subjective, objective and adaptive influences on thermal comfort. This would allow far greater confidence to set definitive comfort criteria for T-BOT that would ensure practically accurate and fully optimal results. Pantelic et al. [2012] implemented a fairly advanced thermal mannequin-based thermal comfort model for their study on HVAC optimisation, so there is precedent for such developments.

Whilst the program was designed with hospital environments in mind, in terms of design variables and characterisation of thermal comfort, it was found that a modular program architecture was appropriate as shown in Fig. 4.18. This means that the program can be readily adapted to other contexts; other comfort metrics can be calculated due to the OSMO approach, though additional data may need to be extracted from the samples for this, and additional design variable writer scripts can be easily added into the program.

As noted previously, self-validation of the optimality of the solutions in absolute terms is likely needed in the T-BOT methodology before it can be used with confidence in practice. Given the comparative length of time of running initial samples and performing optimisations within the framework of T-BOT, it would probably not increase total run-time by a significant proportion to simulate one or more solutions with the actual building model after the optimal Pareto front has been identified, up until agreement is found between the predicted responses and the actual responses to within a certain tolerance. This may be only the thermally optimal solution (ie. at 0 discomfort), or it may be a few of the solutions along the length of the Pareto front; further work would be needed to ascertain the most effective way of doing this. This would allow evaluation of the true optimality of the optimal solutions identified by T-BOT, but also the fidelity of the metamodels in the region of these optima. Moreover, if the predicted and simulated performance do not match to within the tolerance the newly run simulation results may be added to the sample space and further refine the metamodel in the region of the optimum. This would combine the economies of the OSMO approach with the targeted successive model refinement process present in other BTO methodologies in the literature; for example that of Gengembre et al. [2012].

It is worth noting that this self-validation could not be achieved within the framework of T-BOT as it currently stands; a greater degree of automation would be necessary. In particular, the transfer of data between the simulation and optimisation phases would need to be fully

automated and made bi-directional. However further development of usability and automation is required anyway for T-BOT to be generally usable in a professional context.

As a further note on validation, it is worth considering the CFD in the building simulation. CFD is a highly complex modelling domain, and is generally not able to be solved analytically. Hence it is solved iteratively to within a tolerance of the residuals, essentially a measure of the stability of the variables between iterations. This introduces a certain element of uncertainty within the solutions that is not present in DTM. As previously mentioned, the conflation procedures and general formulation of the CFD in ESP-r goes some way to mitigating this problem, as it is specifically formulated to tailor the CFD to the specific building simulation that is undertaken at each time-step. In the context of T-BOT, the vast majority of CFD solutions were found to converge to sensible solutions that were consistent with the other modelling domains; only around 3 CFD solutions among the 39,600 used in the case studies were found to have converged to nonsensical solutions, and these were identified and removed from the samples during the manual metamodelling checking process. However, some further work to identify more robust ways of automating the CFD and checking the solutions may help to improve the consistency of the solutions.

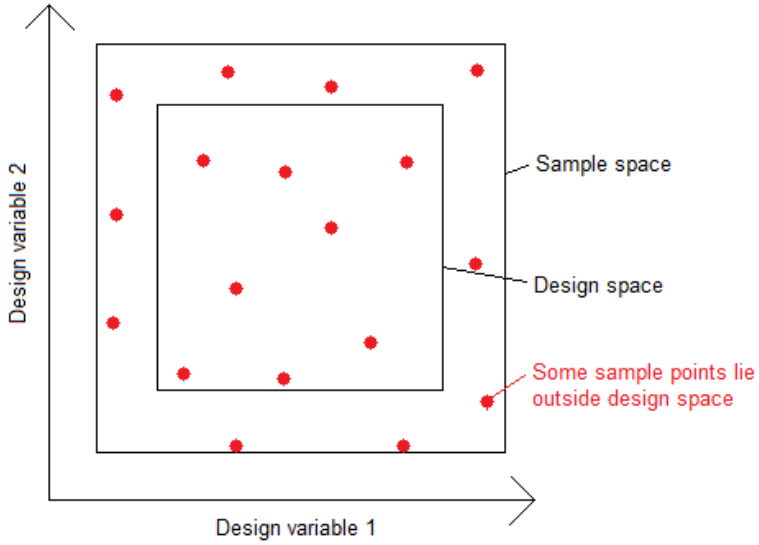

Figure 8.6: Diagram of expanding the sample space beyond the limits of the design space, for an example of 2 design variables.

Significant opportunity also exists to reduce the time required to run the initial sample if the CFD model can be improved or replaced to reduce computational requirement. For example, the CFD coupling procedures in ESP-r could be reviewed and developed to try to reduce the number of simulations necessary to arrive at the final results for the time step. Alternatively the solution procedure of the CFD could be entirely replaced by a less computationally

intensive one, for example an alternative method known as fast fluid dynamics (FFD) shows great potential in the field of building simulation [Zuo & Chen, 2010].

The final solution self-validation procedure described above links in with development of the sampling DOE, as it adds a progressive refinement element. However it may be worth examining the potential to also improve the initial sampling DOE. A number of possibilities are available for this.

Firstly, in order to address the issue of metamodel fidelity deteriorating close to design variable limits, it is conceivable that making the design variable limits of the sample wider than the limits of the design space, as demonstrated in Fig. 8.6, may be beneficial. This would give precedent for responses outside of the boundaries of the design space, which would allow the metamodels to far more accurately approximate those responses in these regions. However, this would also implicitly reduce sample density within the design space, which may decrease the fidelity of the metamodels in a general sense. Further work would be necessary to evaluate this trade-off and see if it would be valuable to implement.

Secondly, it may be beneficial to integrate a certain degree of optimisation functionality into the simulation phase. This may allow a more intelligent variable DOE, tailored to the individual design space; sample distribution could be skewed to cover in more detail areas within the design space where the optima are likely to lie. This could take the form of a coarser initial sample with a uniform DOE to assess general trends of the sample space, then a more focussed sample with a DOE that is skewed to be denser in regions where optimum solutions are likely to lie, evaluated by an optimisation algorithm. However, it is likely that this would be difficult to achieve whilst maintaining the OSMO economies. The OSMO approach allows multiple optimisation problems to be performed using the same sample; unless the optima for every single one of these separate optimisations lay in approximately the same region of the design space, the value of this approach would be limited.

As a final point, constraints within the optimisation problems could go some way to improving the usefulness of the outputs from T-BOT. For example, a constraint could be specified to set a maximum value of the thermal discomfort objective. This could force the Pareto front to only cover a region within which deviation from the comfort criteria may be acceptable, for example. Also, constraints could be set to prevent radiant panel or corridor temperatures from rising above the comfort temperature range during the cooling season or below it during the heating season; this may reduce any tendencies to advocate energy-inefficient strategies such as the peak-reducing behaviour manifested in the case study results.

Chapter 9: Conclusions

9.1 Chapter Overview

This Chapter summarises key findings of the work and relates these back to the project aims and objectives, explicitly defines the contributions to knowledge, and summarises key opportunities for further work.

9.2 Key Findings

This section summarises the key findings of the work:

1. The developed method enables simultaneous optimisation of the radiant and convective thermal environment, at a local spatial resolution, in a strongly coupled manner. Results suggested that optimising at different locations can result in different optimum conditions. For example the summer cases that were examined in detail exhibited a difference in energy use of over 3 kWh between thermally optimum conditions, though winter results did not show the same trend.
2. The method also enables optimisation of variable time periods. Results suggested that optimising at a finer temporal resolution (consistent with finer control periods in the case of constant design variables pertaining to system operation e.g. radiator temperature) can have benefits to energy use. For example the summer cases examined in detail showed a 13% saving in energy use of thermally optimum solutions by optimising four 6 hour periods rather than one 24 hour period. Winter results did not exhibit this saving. However, results also suggested that there are limits to the benefits that can be gained from finer temporal resolution (e.g. no further benefit from 3 hour periods over 6 hour periods).
3. Any thermal comfort metrics that can be calculated from the primitive data stored from the sample can be used in the thermal discomfort objective function. Results suggested that using different thermal comfort criteria can have a large impact on results; results that characterised comfort limits in terms of PMV were found to be markedly less reliable due to poorer metamodel fidelity and also gave different optimum conditions (e.g. difference of over 4 kWh for thermally optimum winter results).
4. The methodology is efficient in that any combination of location (defined pre-sampling), time period (within the simulation period) and thermal comfort criteria

(subject to the data stored from the sample) can be optimised from a single sample. This represents substantial time-saving economies over a direct search optimisation approach, as well as other metamodel based BTO procedures in the literature (see section 9.3).

5. The addition of CFD into the building simulation places the majority of the computational burden on the simulation of the sample; this aligns well with point 4 above, as many quick optimisations can be performed from a single slow sample (in relative terms).
6. Computational parallelisation is particularly useful for the sampling, as each sample is independent and therefore the run-time of the sample scales linearly with parallelisation down to the level of one core per sample. Benefits of further parallelisation were not investigated.
7. The definition of the sample/design space is critical to the effective operation of the program. This is defined by the design variables and their limits. Care should be taken to minimise the amount of design variables and the range of values wherever possible. The sample size is critical to the accuracy of the metamodels. Results suggested that a sample of 150 simulations (100 in the build set, 50 in the validate set) is generally sufficient to obtain representative results, for a sample/design space of 5 dimensions (i.e. 5 design variables), provided solutions do not converge to the edge of the design space in terms of more than 1 design variable (see point 8 below).
8. The method suffers from degraded metamodel fidelity near boundaries, and particularly corners, of the design space. Care should be taken to set design variable ranges to try to minimise the possibility of optimum solutions converging near design space boundaries, and particularly corners, wherever possible. Further development may address this issue (see section 9.4).
9. It is critical that the model used in the sampling represents the space to be optimised as closely as possible. In the present work a number of modelling disparities were identified, which limited the development of results from T-BOT into practical design and operation conclusions.
10. Seasonal variations were found to have significant effects on results in multiple aspects. Large simulation periods may make sample run-times prohibitive, so it is recommended to establish smaller representative periods and run samples for each of these individually. Methods for obtaining representative extreme and average periods of 1 day have been suggested.
11. T-BOT should not be regarded as a black-box model, rather an information gathering framework and decision support tool. The characterisation of results as a Pareto front

in terms of thermal discomfort and energy use, allows the user to assess the optimum trade-off between these criteria and select practically appropriate solutions based on their own judgement. The optimality of solutions on the Pareto front in an absolute sense is not implicitly assured with the current methodology. For example results indicated that thermal discomfort was often overestimated and hence lower energy solutions may be incorrectly discarded from the Pareto front, particularly solutions with discomfort close to 0. Further development may address this issue (see section 9.4).

9.2.1 Satisfaction of project aims and objectives

It is suggested that based on the criteria set out in section 1.8, the project has been broadly successful. A BTO methodology that provides the functionality of "quantitative evaluation of an optimum trade-off between thermal comfort and energy use, evaluated at a local level within individual spaces" has been developed, implemented and tested based on "the state-of-the-art art in the fields of thermal comfort, building simulation and optimisation and key research gaps in the field of BTO".

Objectives 1 and 2 have been fulfilled by the comprehensive literature review presented in Chapter 2. Objective 3 has been fulfilled by the development detailed in Chapters 4 and 6. Objective 4 has been fulfilled by the proof-of-concept study detailed in Chapter 5, and objective 5 has been fulfilled by the case study detailed in Chapter 7.

9.3 Contributions to State-of-the-Art

The T-BOT methodology that has been developed is, to the author's knowledge, unique in a number of aspects:

1. The use of coupled DTM and CFD in a BTO context is not well studied in the literature. To the author's knowledge, adaptive coupling of comprehensive, independent DTM and CFD domains has not been applied in a BTO context.
2. The disconnect between the simulation and optimisation phases, induced by metamodeling, enables additional economies over a direct search optimisation approach. These economies arise from saving a larger amount of primitive information from the sample, which allows the data to be compiled in a variety of contexts to train metamodels for optimisation. Essentially, many different contexts

can be optimised from a single sample. This has been termed an OSMO approach, and allows contextual variation in optimisations in three aspects:

a. In the current methodology, air temperature, air velocity and MRT are saved for a variety of locations defined before the simulation stage. This allows operative temperature to be calculated at any of these locations and used to train metamodels, from a single sample.
b. Information is saved at every time step of the building simulation, which means objective function metrics can be calculated from aggregation of any combination of time steps. This means any period within the simulation period may be optimised from a single sample.
c. Additionally, the primitive information saved from the sample allows other comfort metrics to be calculated as well as operative temperature, though this may require the methodology to be developed to extract additional information from the sample simulations. This could allow multiple comfort metrics to be explored from a single sample.

Furthermore, results from testing have exposed some aspects of BTO that are not well studied in the literature. These factors can be explored more efficiently with T-BOT than other BTO programs found in the literature:

3. Results indicate that spatial variation of thermal conditions may be a significant factor in optimum conditions; what provides optimum conditions at one point in the room may not necessarily provide optimum conditions at another. This is not commonly considered in BTO studies found in the literature.
4. Variation in temporal resolution has been shown to significantly and consistently affect results, though case study results suggested that this can be strongly dependent on seasonal variations. In practical terms this relates largely to control periods, which to the author's knowledge has not been studied in a BTO context.
5. Results also indicate that the choice of thermal comfort criteria may have a significant effect on results, which exposes a potential weakness of other BTO programs in the literature that are limited to single comfort metrics, commonly PMV and/or PPD.

9.4 Further Work

This section summarises key opportunities for further work suggested by the testing of the program. For a more in-depth discussion of these points the reader is referred to section 8.5.

1. Metamodel fidelity is key to the accuracy of results; a number of potential developments could assist in ensuring confidence in this critical aspect:
    a. A post-optimisation sequential simulation and metamodel reformulation procedure could be implemented to ensure absolute accuracy of optimum solutions. For example, the user runs an optimisation and is presented with a Pareto front of solutions. From these solutions they select a few that are of most interest to them. These solutions are then simulated with the actual building model (clearly this is subject to the caveats of simulation time), and the simulated results are compared to the results predicted by the metamodels. If they do not agree to within a certain tolerance, then the metamodels are reformulated including the new simulations as additional sample points, and the optimisation continues for a few generations. This process is repeated until the user's selected optimum solutions agree with the simulation results. A similar procedure that includes progressive simulation in the metamodeling procedure could also be applied; there is precedent in the literature for such procedures, such as the work of Gengembre et al. [2012], however this may compromise the OSMO economies of the current T-BOT methodology.
    b. Metamodel fidelity degrading at design space boundaries could be addressed by making the sample space larger than the design space, such that sample points lie outside the design space and hence provide greater precedent for predictions at the boundaries of the design space. This could be accomplished practically by extending the design variable limits beyond those specified for the design space, though clearly this may cause problems where the limit of a design variable is close to 0 and cannot take negative values.
    c. It may be worth investigating a broad range of metamodeling methods in greater detail, as only basic implementations of neural networks and support vector regression were considered as alternatives to MLSR in the present work.
2. Given the strong dependence of simulation run-time on the CFD, it seems prudent to investigate less computationally intensive alternatives, for example the fast fluid dynamics (FFD) method proposed by Zuo and Chen [2010]. Alternatively, the possibility of reducing the number of CFD simulations in each simulation could be considered, for example only running the CFD every other time step.

3. Implementation of more modern thermal mannequin comfort models could also be worth investigating, though this may affect the OSMO economies depending on what information is required as inputs for these models.
4. Design variable sensitivity was not considered in great detail in the present work, but given the dependence of the solutions on design variables as highlighted in section 9.3.1, it may be useful to add an additional step into the process to examine design variable sensitivity before optimisations, or even the sampling, is performed. Precedent exists in the literature for such an addition, for example the work of Eisenhower et al. [2012].
5. Also, whilst the program can be fairly easily adapted to contexts outwith hospital environments as mentioned in section 9.2, it may be necessary to investigate program performance in other contexts before application.
6. Finally, it may be worth investigating whether further constraints could improve the procedure. For example, constraints could be placed on objective function values, or further functions of design variable values.

Appendix A: Matlab Code for Calculation of PMV

```
function [PMV]=calc_PMV(clo,met,work,airT,MRT,airV,hum)


% calculate vapour pressure
PS=hum*10*exp(16.6536-4030.183/(airT+235));
% convert units
ICL=clo*0.155;
M=met*58.15;
W=work*58.15;
% internal heat production
MW=M-W;
% clothing area factor
if ICL<=0.078
   FCL=1+1.29*ICL;
else
   FCL=1.05+0.645*ICL;
end
% heat transfer coefficient?
HCF=12.1*sqrt(airV);
% degrees C to kelvins
TAA=airT+273;
TRA=MRT+273;
```

```
% --- CALCULATE SURFACE TEMP OF CLOTHING BY ITERATION ---

% initial guess

TCLAA=TAA+(35.5-airT)/(3.5*(6.45*ICL+0.1));

TCLA=TCLAA-273;

% calculation terms

EPS=0.00015;

N=0;

while N<=151

   HCN=2.38*abs(TCLA-airT)^0.25;

   if HCF>HCN

      HC=HCF;

   else

      HC=HCN;

   end

   TCLB=35.7-0.028*MW-ICL*(3.96e-8*FCL*(TCLAA^4-TRA^4)+FCL*HC*(TCLA-airT));

   N=N+1;

   if abs(TCLA-TCLB)<=EPS

      N=152;

   end

   TCLA=(TCLB+TCLA)/2;

   TCLAA=TCLA+273;

   if N==151

      error('iteration did not converge below tolerance');

   end
```

```
end

% surface temp of clothes

TCL=TCLB;

TCLA=TCL+273;



% --- HEAT LOSS COMPONENTS ---

% skin

HL1=0.00305*(5733-6.99*MW-PS);

% sweating

HL2=0.42*(MW-58.15);

% latent respiration

HL3=1.7e-5*M*(5867-PS);

% dry respiration

HL4=0.0014*M*(34-airT);

% radiation

HL5=3.96e-8*FCL*(TCLA^4-TRA^4);

% convection

HL6=FCL*HC*(TCL-airT);



% --- CALCULATE PMV ---

% thermal sensation trans coefficient

TS=0.303*exp(-0.036*M)+0.028;

PMV=TS*(MW-HL1-HL2-HL3-HL4-HL5-HL6);

End
```

Appendix B: Full Results of Case Study

B.1 Case Indexing System

In order to concisely describe the myriad cases examined in this work, a referencing system is described here and used hereafter in the appendix.

The case references consist of 5 distinct elements, and are of the following form: [climate case]_[evaluation location]_[optimisation period]_[number of design variables]_[thermal discomfort objective formulation]. Possible values and their meanings for each of these elements are given below:

- Climate case (see section 6.3.6 for more information)
    - "SumHot1" – Extreme summer conditions
    - "SumAve1" – Average summer conditions, case 1
    - "SumAve2" – Average summer conditions, case 2
    - "WinCol1" – Extreme winter conditions
    - "WinAve1" – Average winter conditions, case 1
    - "WinAve2" – Average winter conditions, case 2
- Evaluation location
    - A numeric reference, which refers to a location within the grid of MRT sensors specified at run-time of each sample. This grid and its referencing are shown in Fig. 7.1
- Optimisation period
    - Two numbers, separated by a hyphen; in some cases there may be multiple instances of this separated by commas. This describes the hourly periods over which the optimisation is performed. See examples below for clarification.
- Number of design variables
    - "dvX" where "X" is an integer. This describes the number of design variables in the sample used in that particular case; this is essentially included to provide a convenient means of distinguishing between different samples run for the same climate case.
- Thermal discomfort objective formulation
    - ot – operative temperature formulation, ie. default values of operative temperature limits for comfort region should be assumed; these are given in section 4.6.1.

- pmvX – pmv formulation, ie. operative temperature comfort limits are calculated from PMV comfort limits. "X" is one of three options; "P"atient, "V"isitor or "S"taff. This refers to the case of PMV input variables; parameters of these cases can be found, along with a description of the methodology used to calculate operative temperature limits from the PMV limits, in section 6.4.2.

A few examples are provided here to clarify the above information:

- *Example case 1: SumHot1_1_1-24_dv2_ot*
  - Extreme summer conditions,
  - Evaluation location 1,
  - A single optimisation over the full 24-hour period,
  - A sample of 2 design variables,
  - Default operative temperature limits.
- *Example case 2: WinAve2_15_1-6,7-12,13-18,19-24_dv9_pmvP*
  - Average winter conditions, day 2 (see section 6.3.6 for definition of this),
  - Evaluation location 15,
  - 4 separate optimisations, each for a 6-hour period of the day,
  - A sample of 9 design variables,
  - Operative temperature limits calculated from PMV limits, for the case of patient comfort.

B.2 Extreme Summer Conditions

This case was simulated for the warmest summer conditions, as detailed in section 6.3.6. Two samples were simulated under these conditions; one with 5 design variables and one with 9 design variables. The first of these samples, with 5 design variables, operated under the assumption that the window remained closed at all times. This sample was of the default size; 100 build samples and 50 validation samples. The design variables included were:

- External wall thermal conductivity (0.1-1.5 W/mK)
- External wall thermal storage variables (500-2000 kg/m3 & J/kgK)
- Radiant panel temperature (5-30 °C)
- Door discharge factor (0.1-0.9)
- Corridor air temperature (18-26 °C)

The second sample with 9 design variables investigated the effect of allowing the window to be open. Due the greater dimensionality of the problem, a larger sample was needed, so the number of samples in this set was doubled; 200 build samples and 100 validation samples. The design variables included were:

- Window control opening area ratios (3 variables) (0.1-11)
- Window control temperature set-points (3 variables) (18-28 °C)
- Radiant panel temperature (5-30 °C)
- Door discharge factor (0.1-0.9)
- Corridor air temperature (18-26 °C)

B.2.1 Patient comfort optimisation

Case: SumHot1_15_1-24_dv5_ot

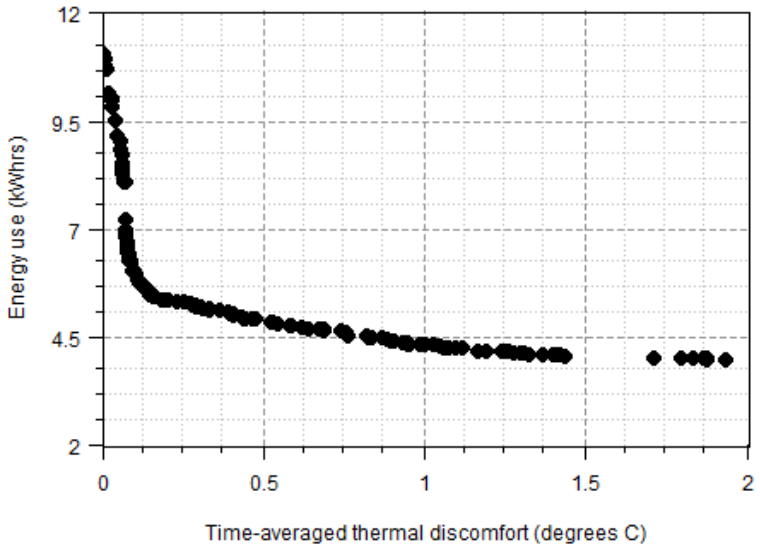

Figure B.1: Pareto front for the *SumHot1_15_1-24_dv5_ot* case shown as a trade-off between the two objectives.

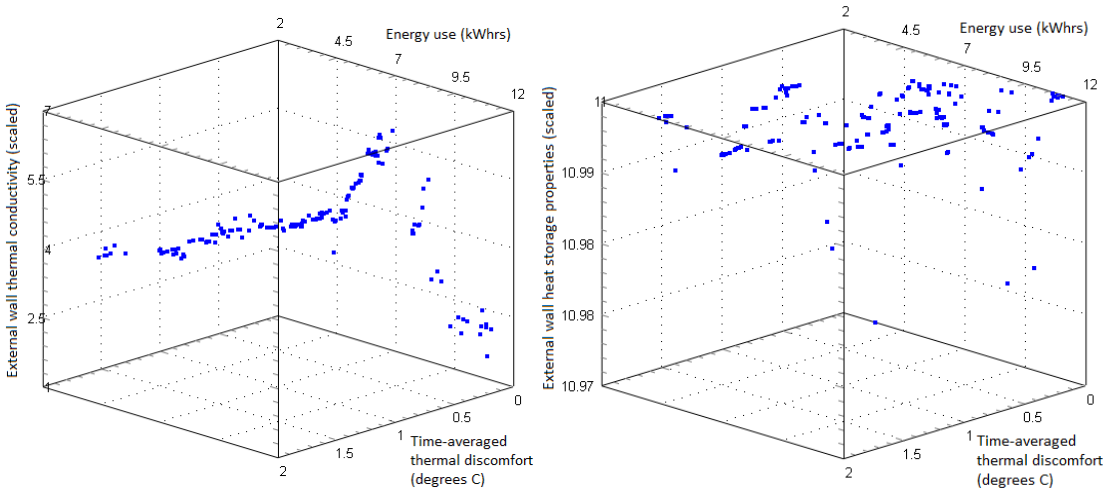

Figure B.2, continued overleaf

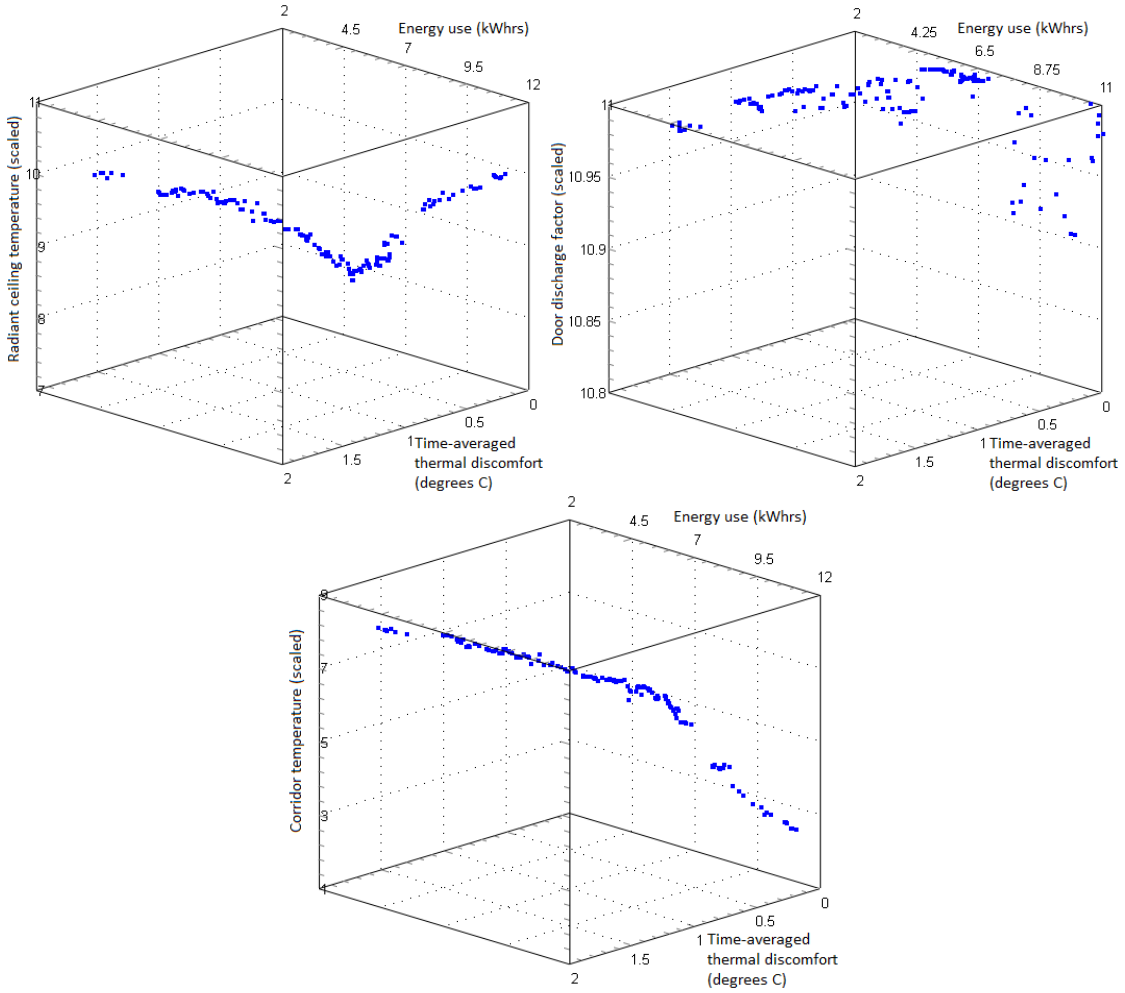

Figure B.2: Pareto front (as shown in Fig. B.1) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

Case: SumHot1_15_1-24_dv5_pmvP

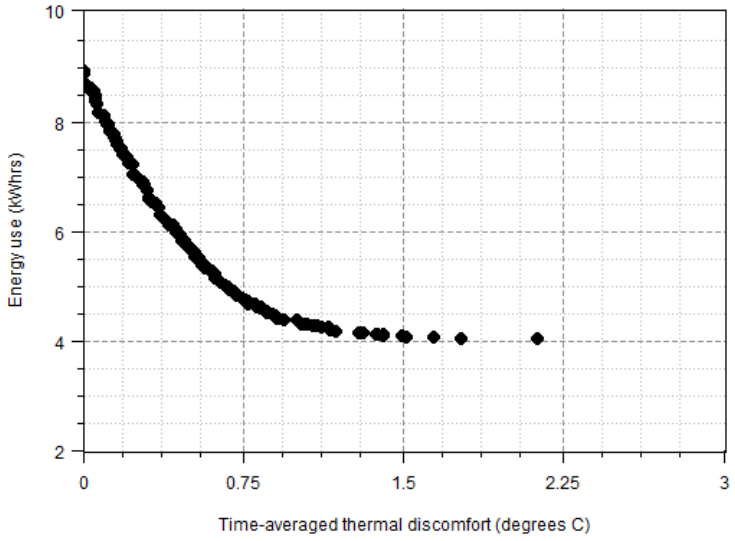

Figure B.3: Pareto front for the *SumHot1_15_1-24_dv5_pmv* case shown as a trade-off between the two objectives.

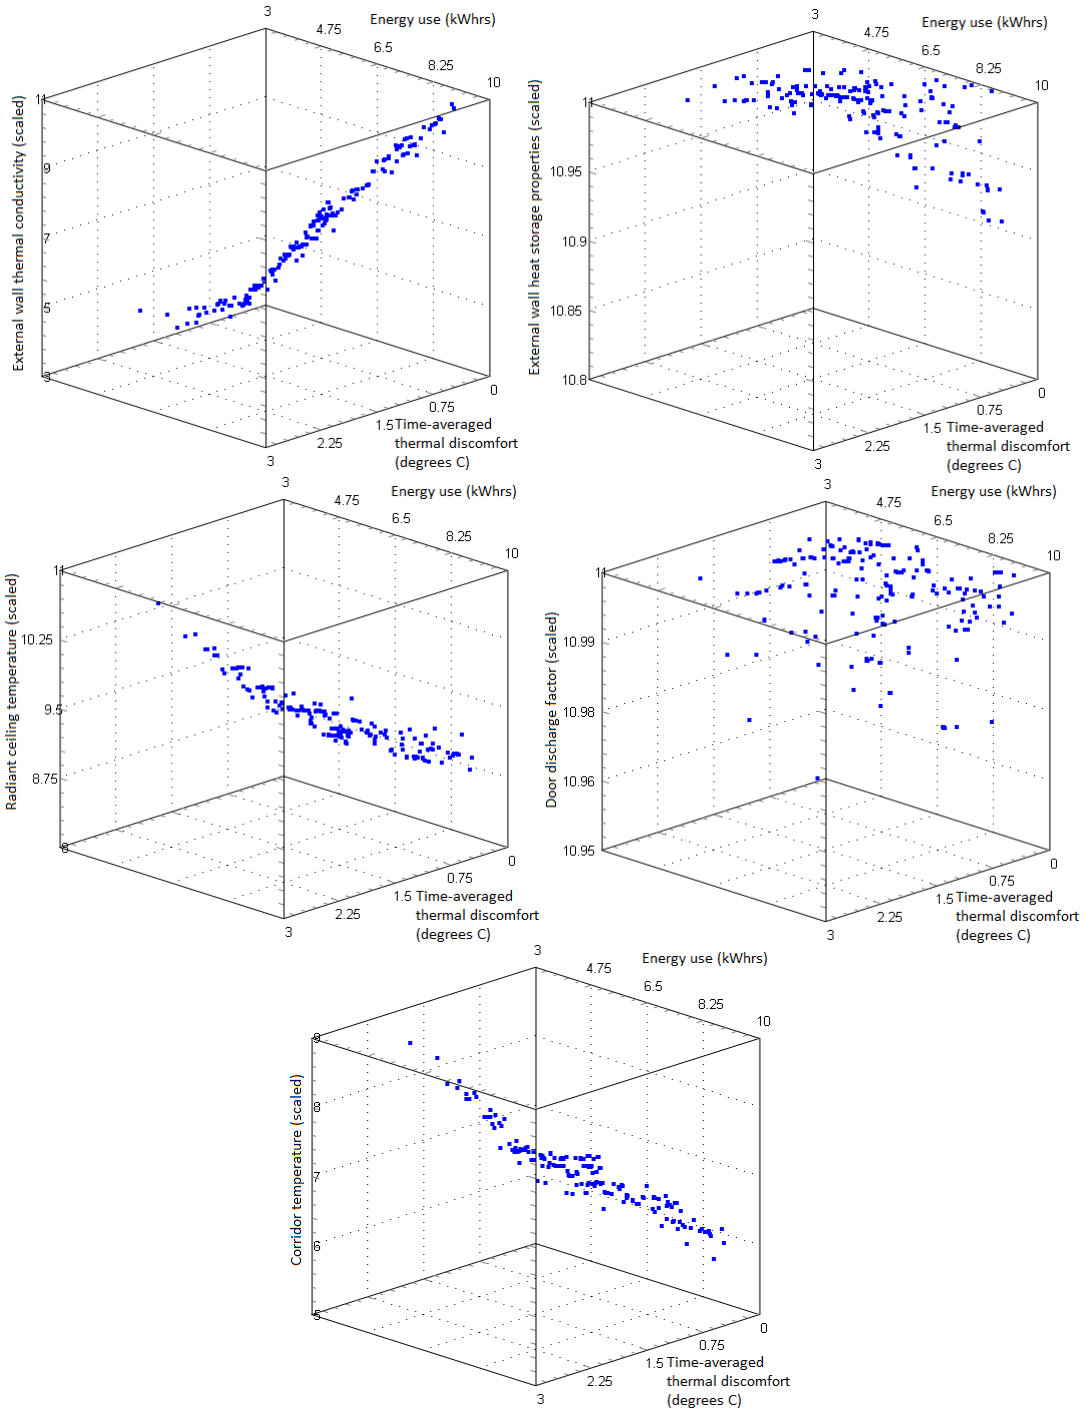

Figure B.4: Pareto front (as shown in Fig. B.3) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

Case: SumHot1_15_1-24_dv9_ot

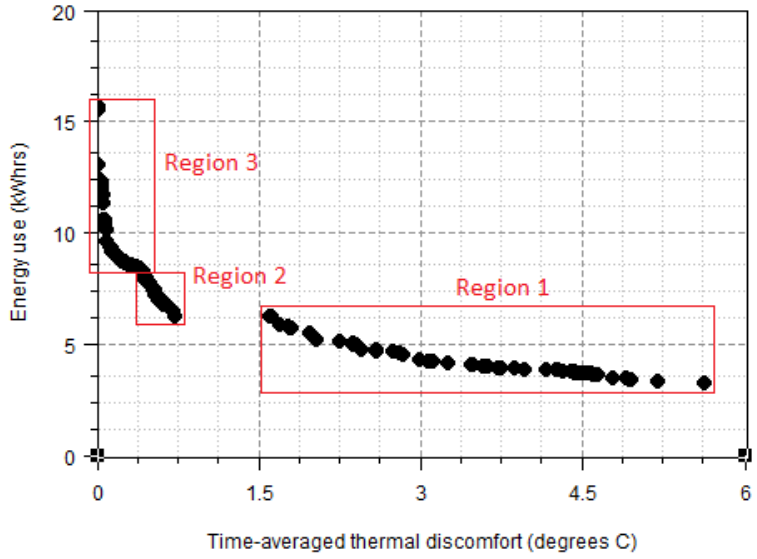

Figure B.5: Pareto front for the *SumHot1_15_1-24_dv9_ot* case shown as a trade-off between the two objectives, with distinct regions marked.

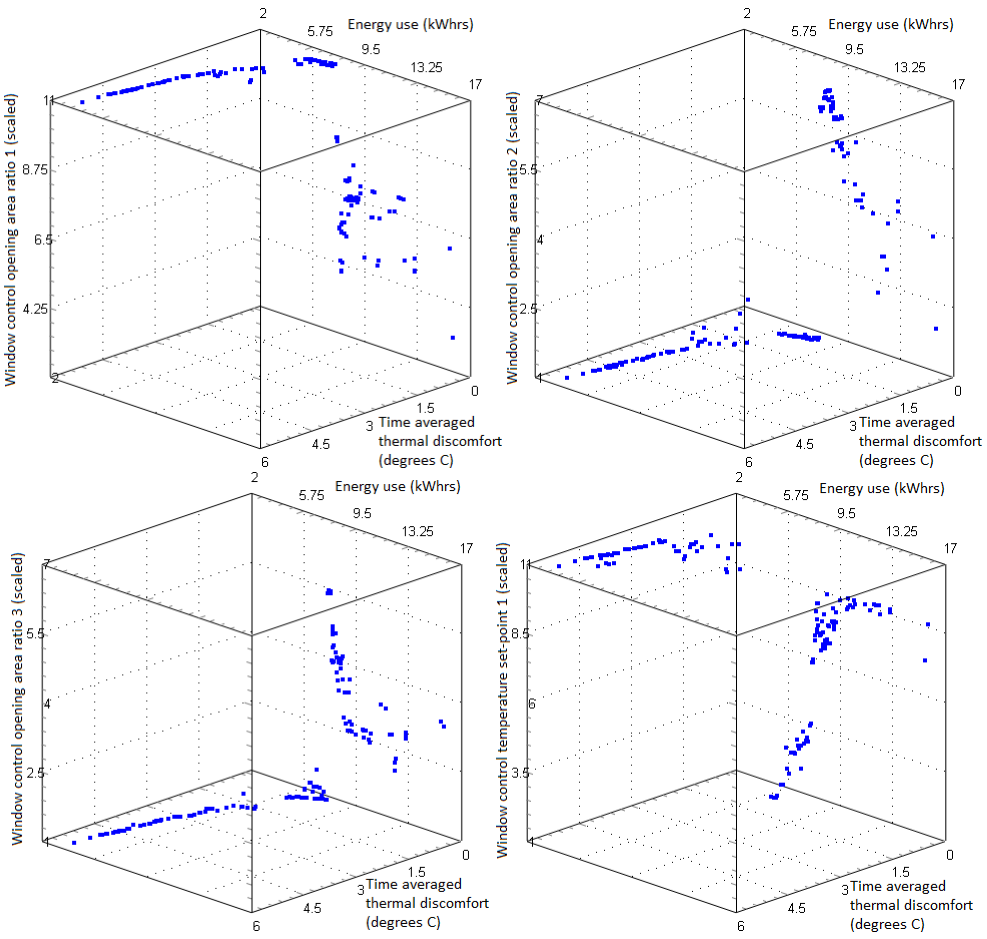

Figure B.6, continued overleaf

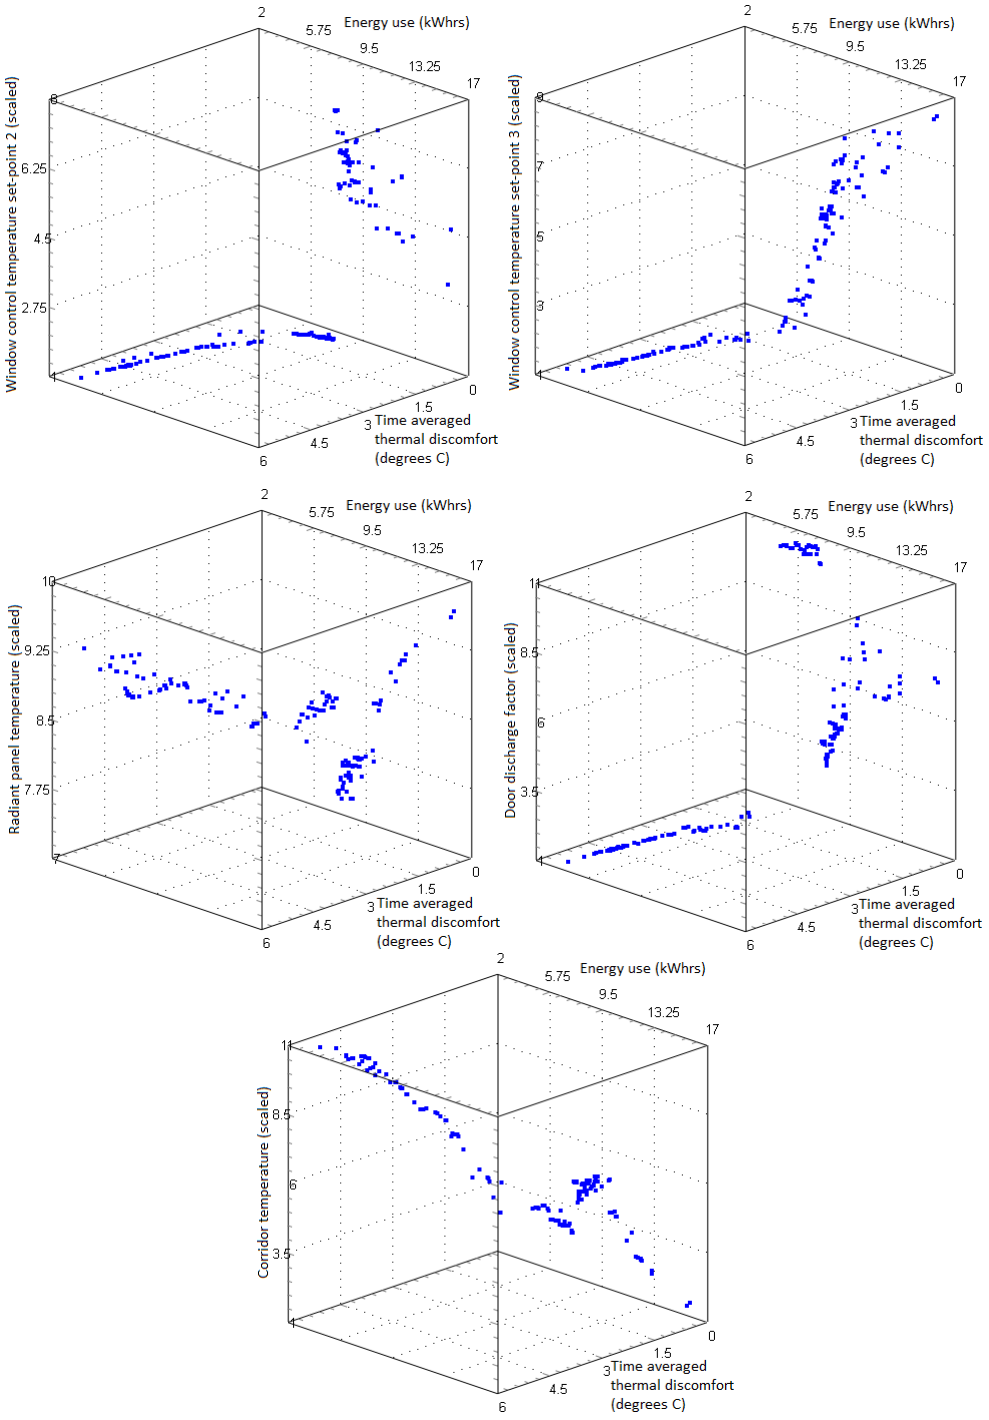

Figure B.6: Pareto front (as shown in Fig. B.5) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

Case: SumHot1_15_1-24_dv9_pmvP

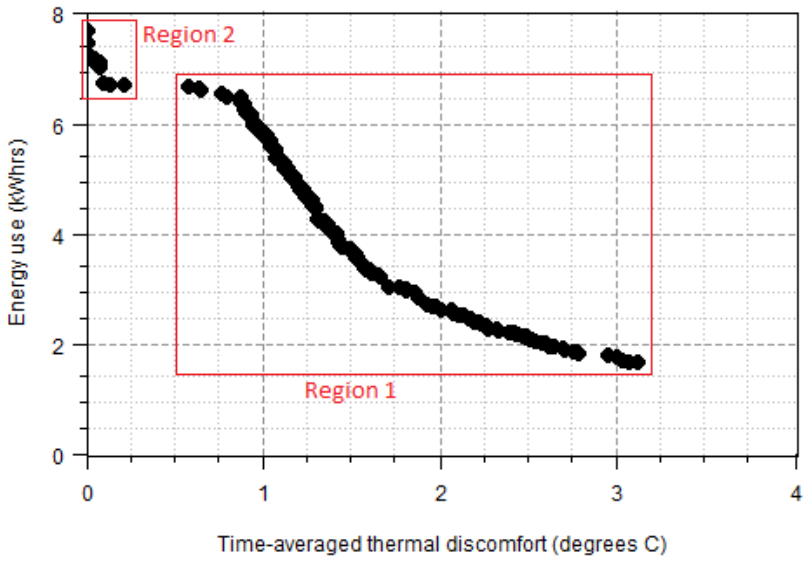

Figure B.7: Pareto front for the *SumHot1_15_1-24_dv9_pmv* case shown as a trade-off between the two objectives, with distinct regions marked.

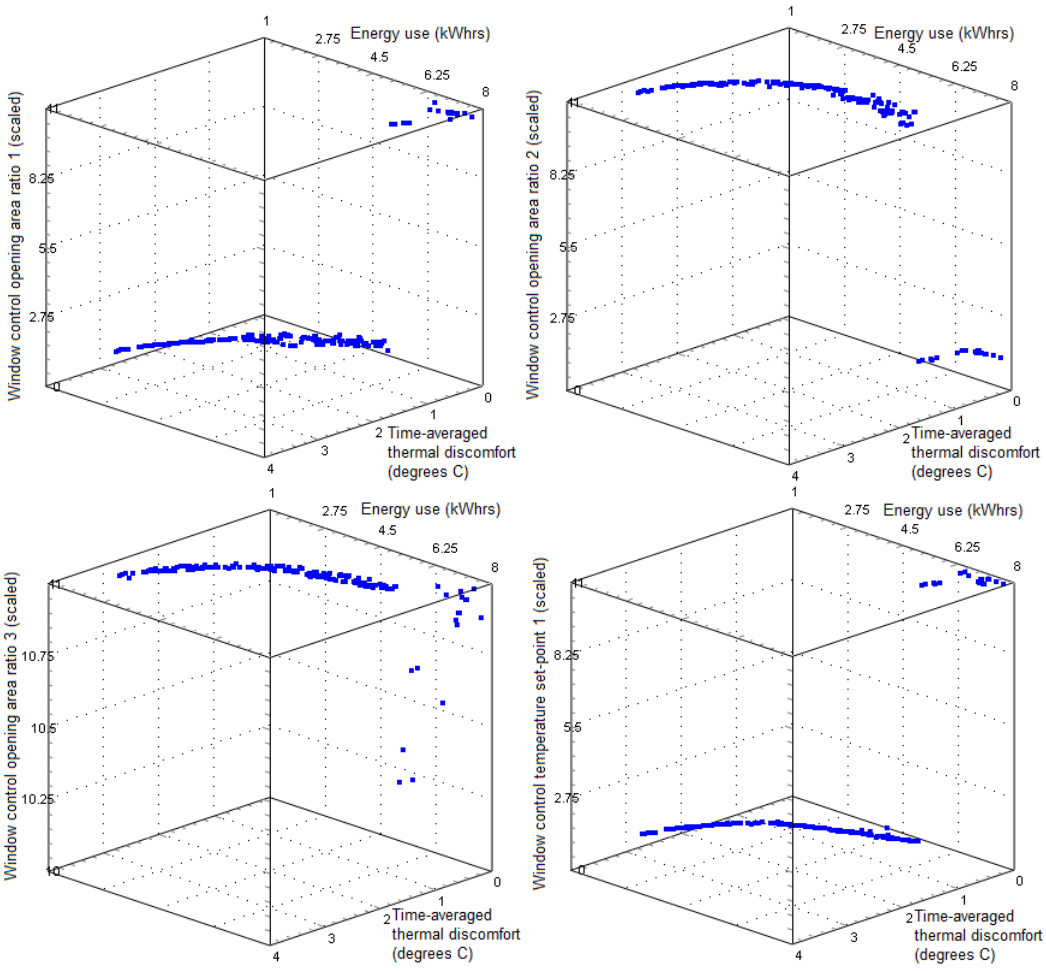

Figure B.8, continued overleaf

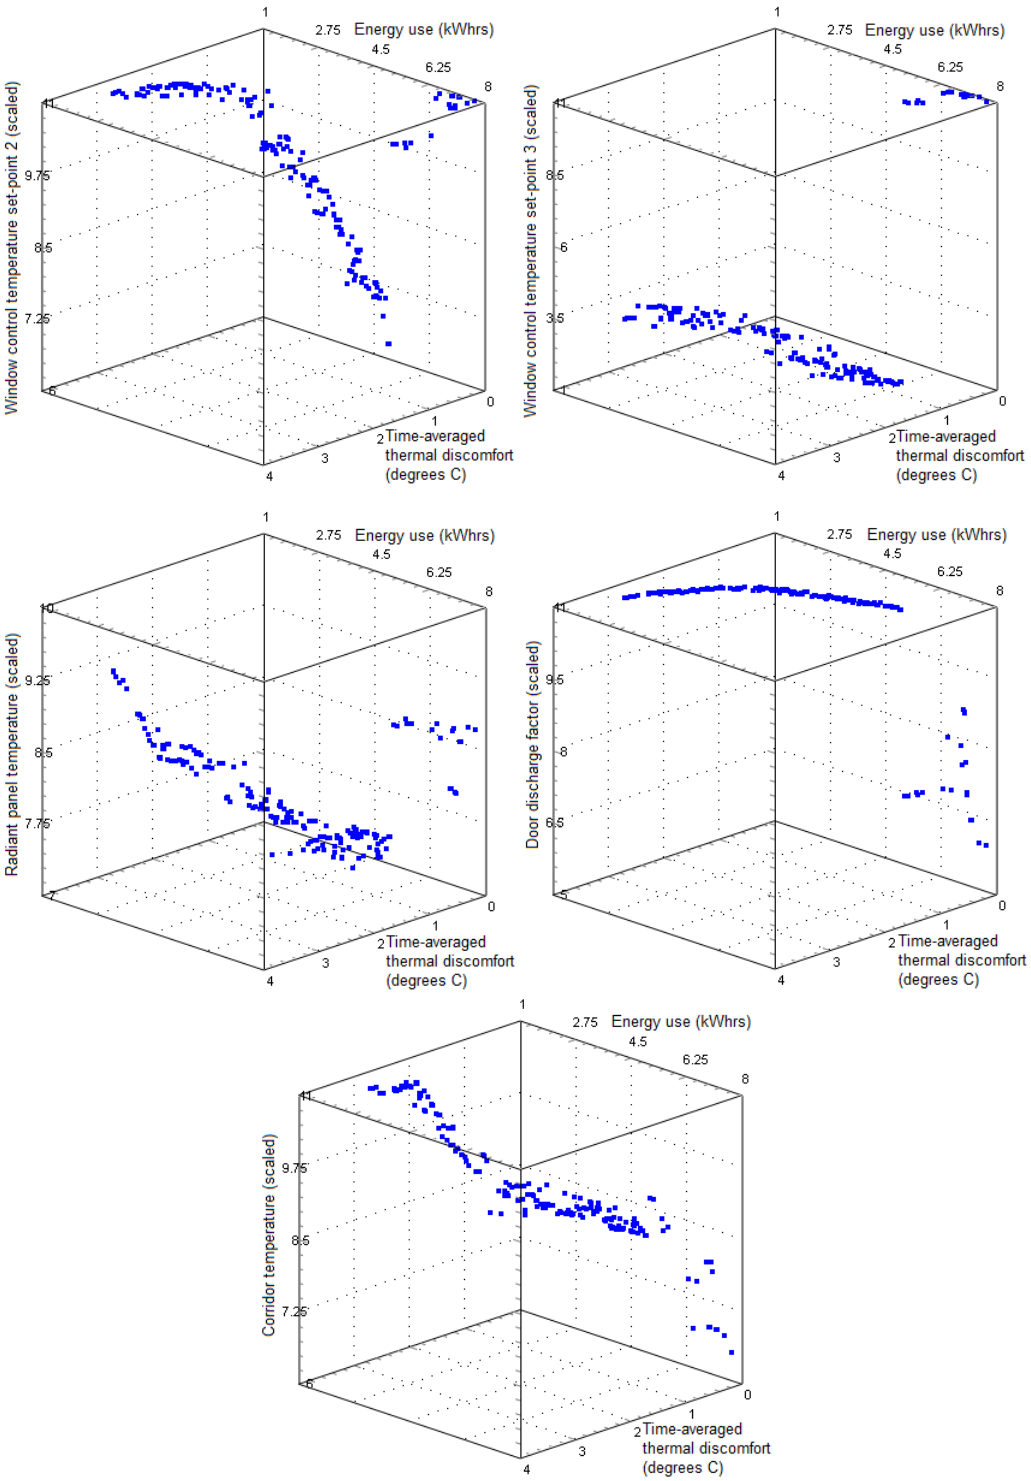

Figure 7.9: Pareto front (as shown in Fig. 7.8) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see above.

Case: SumHot1_15_1-6,7-12,13-18,19-24_dv5_ot

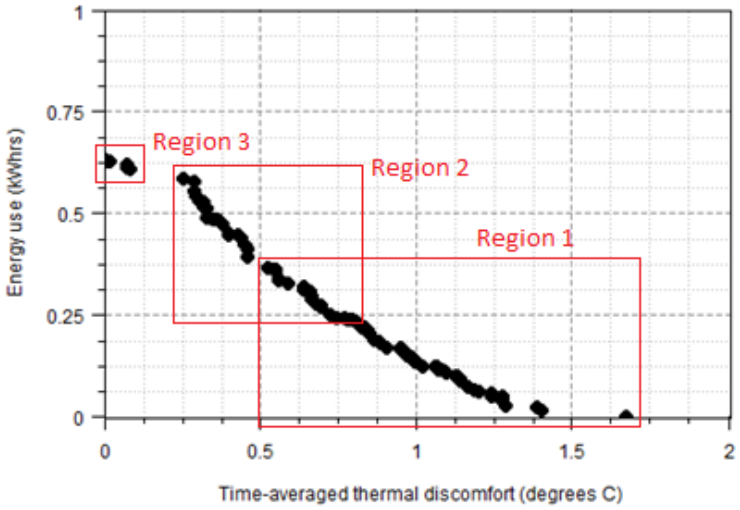

Figure B.9: Pareto front for the *SumHot1_15_1-6_dv5_ot* case shown as a trade-off between the two objectives, with regions marked.

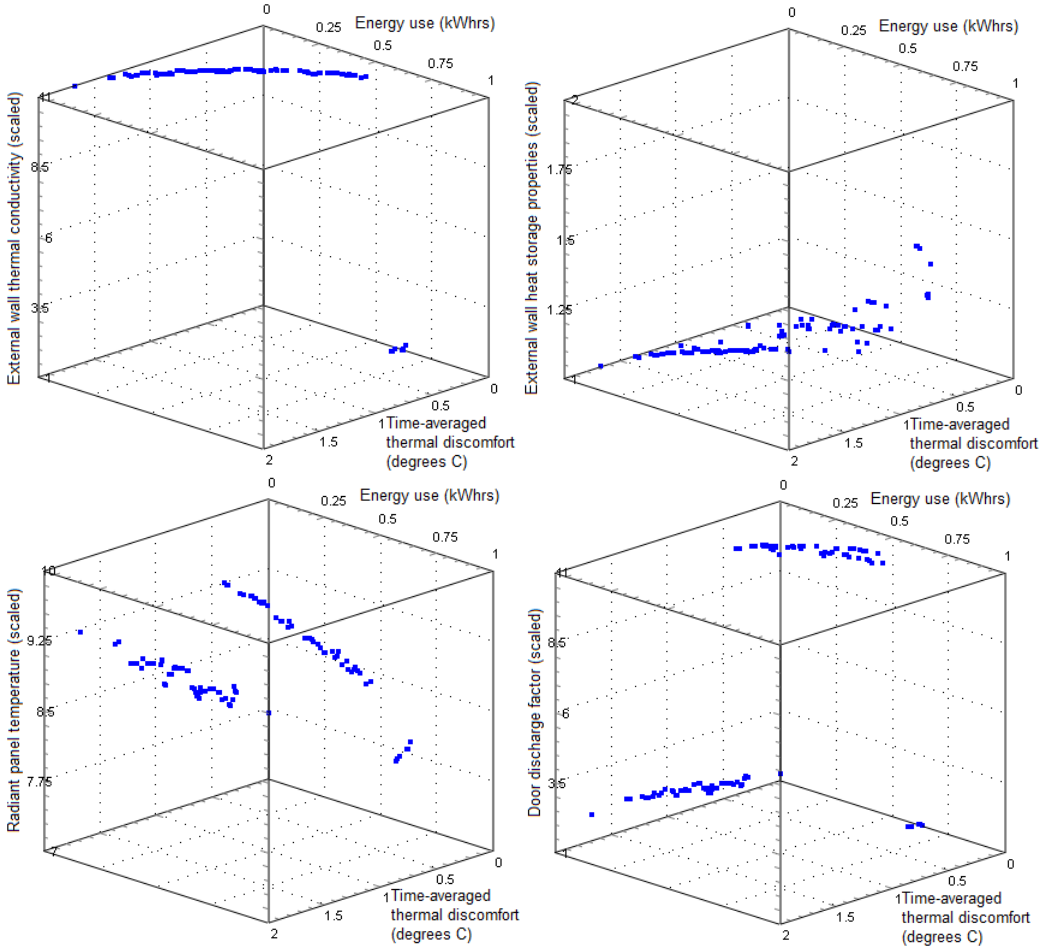

Figure B.10, continued overleaf

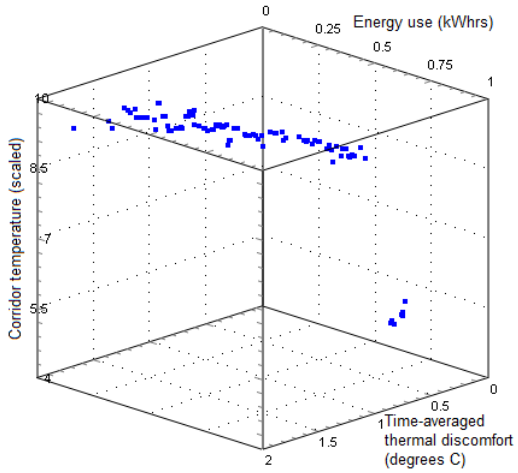

Figure B.10: Pareto front (as shown in Fig. B.9) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

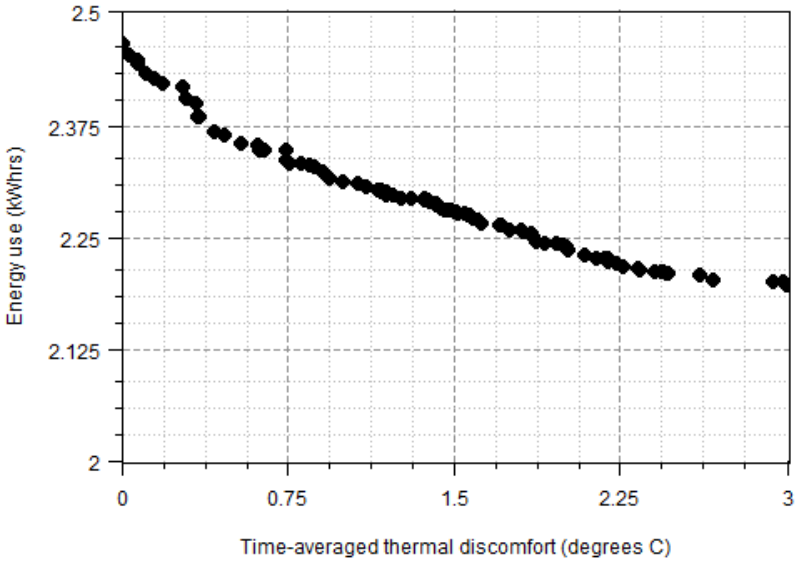

Figure B.11: Pareto front for the *SumHot1_15_7-12_dv5_ot* case shown as a trade-off between the two objectives.

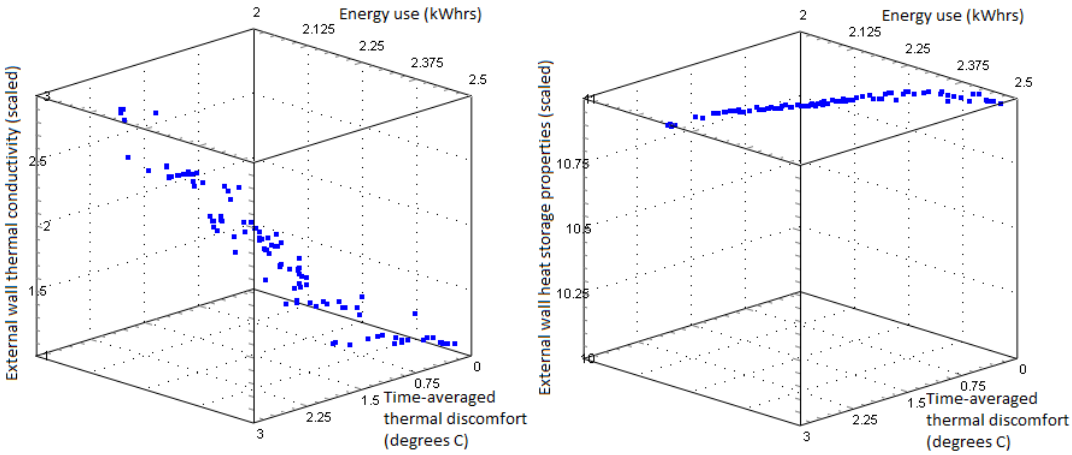

Figure B.12, continued overleaf

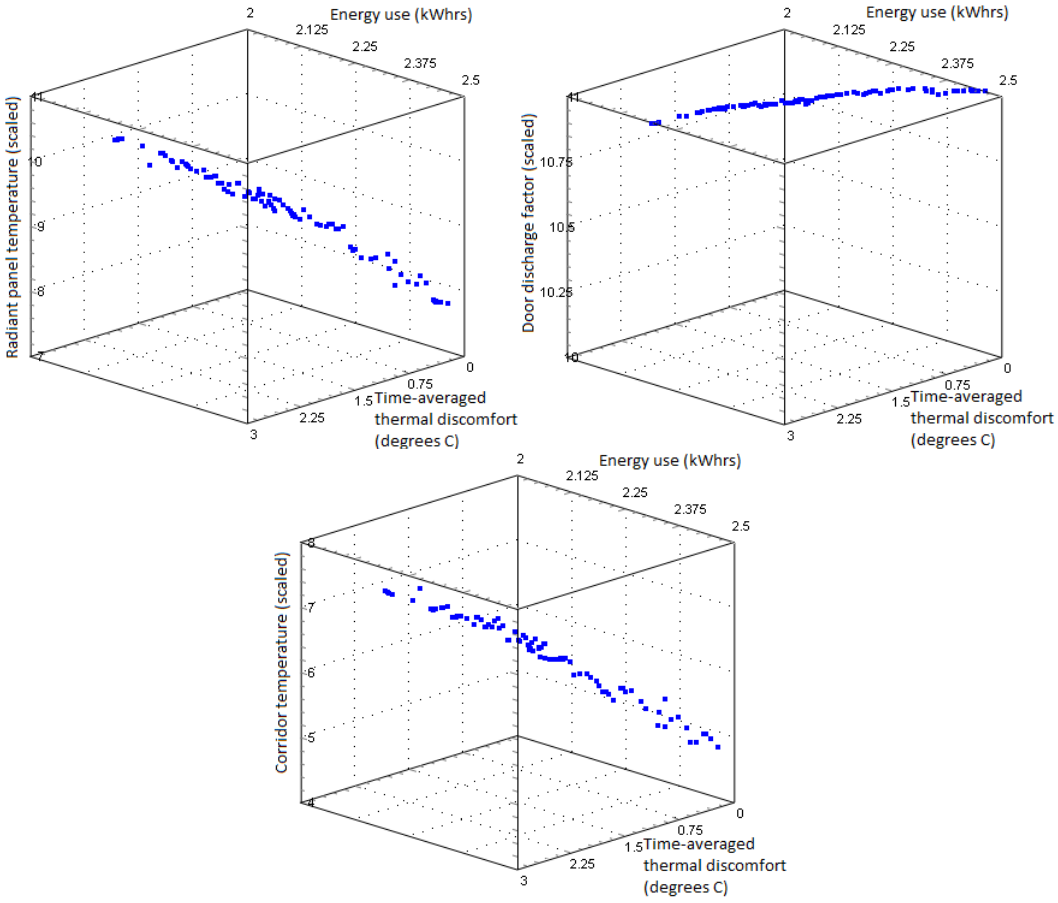

Figure B.12: Pareto front (as shown in Fig. B.11) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

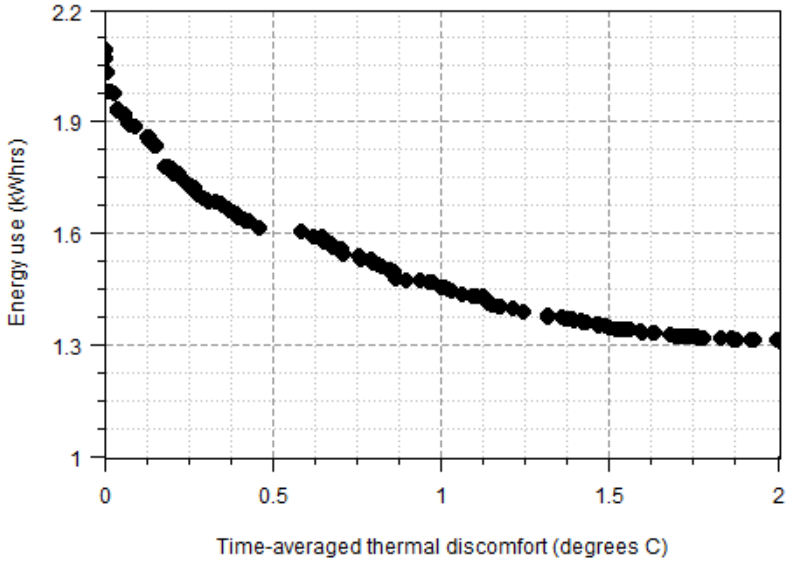

Figure B.13: Pareto front for the *SumHot1_15_13-18_dv5_ot* case shown as a trade-off between the two objectives.

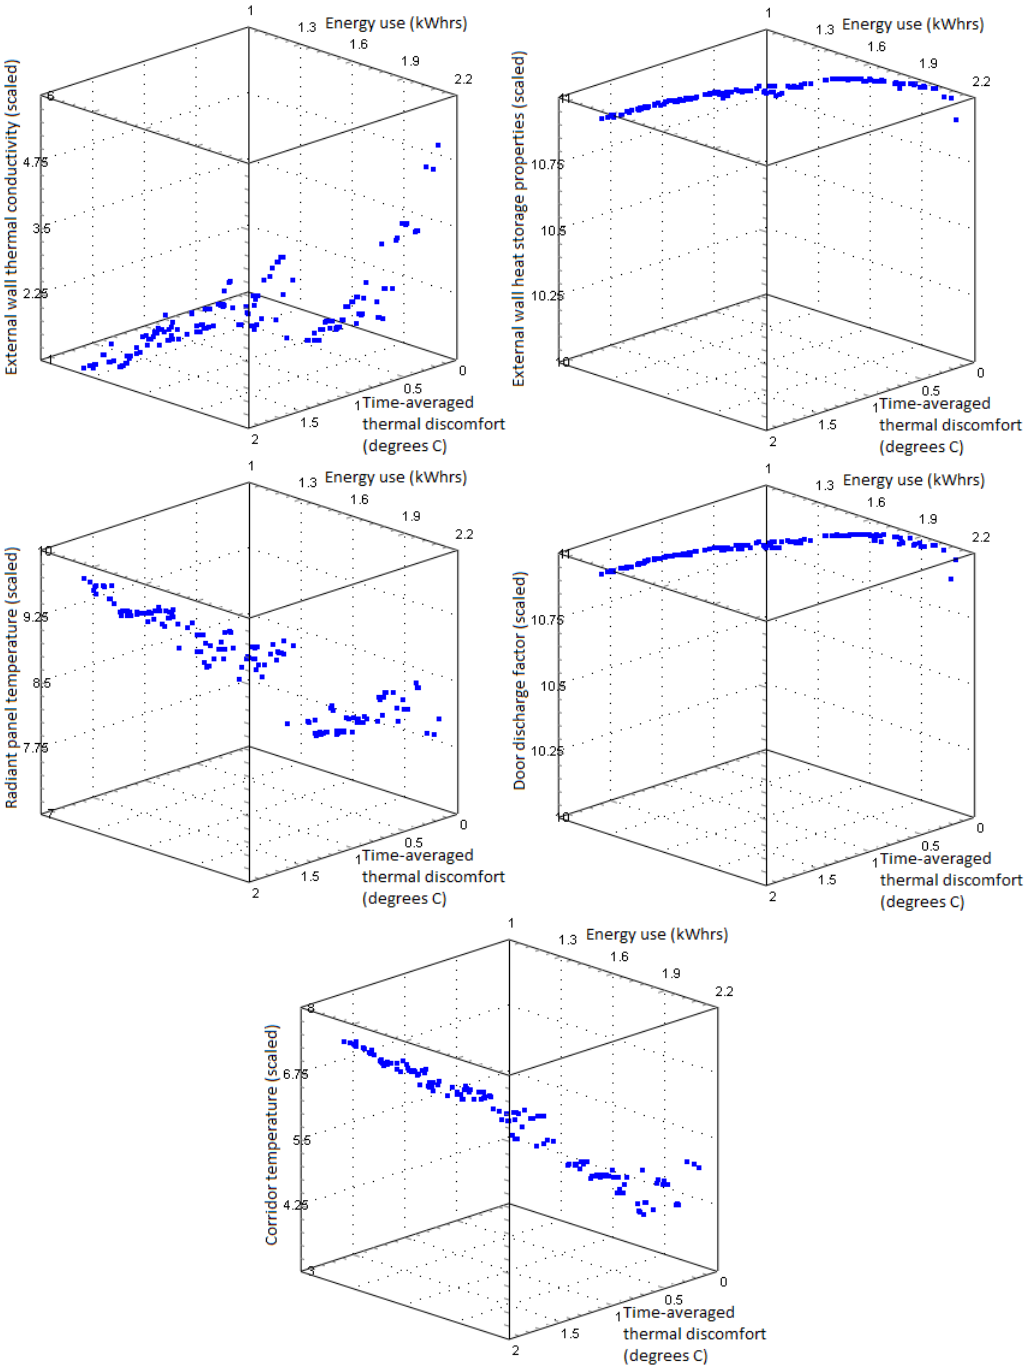

Figure B.14: Pareto front (as shown in Fig. B.13) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

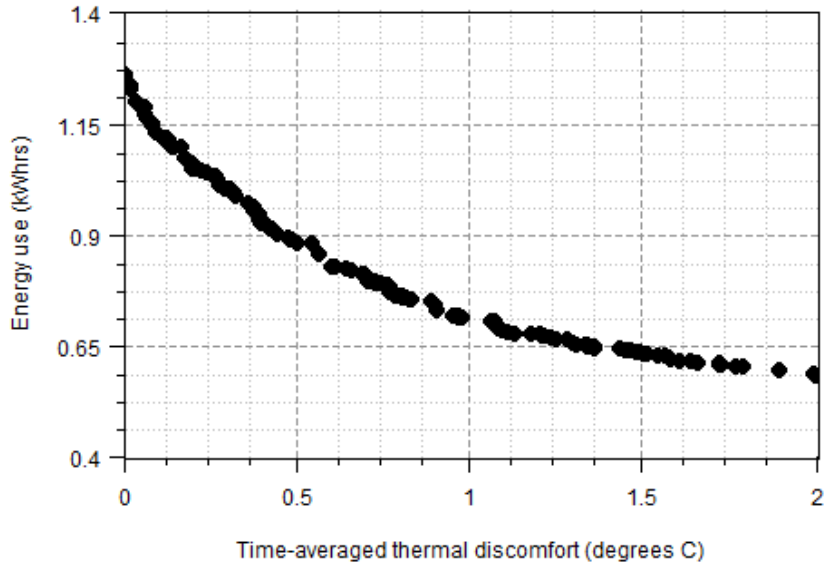

Figure B.15: Pareto front for the *SumHot1_15_19-24_dv5_ot* case shown as a trade-off between the two objectives.

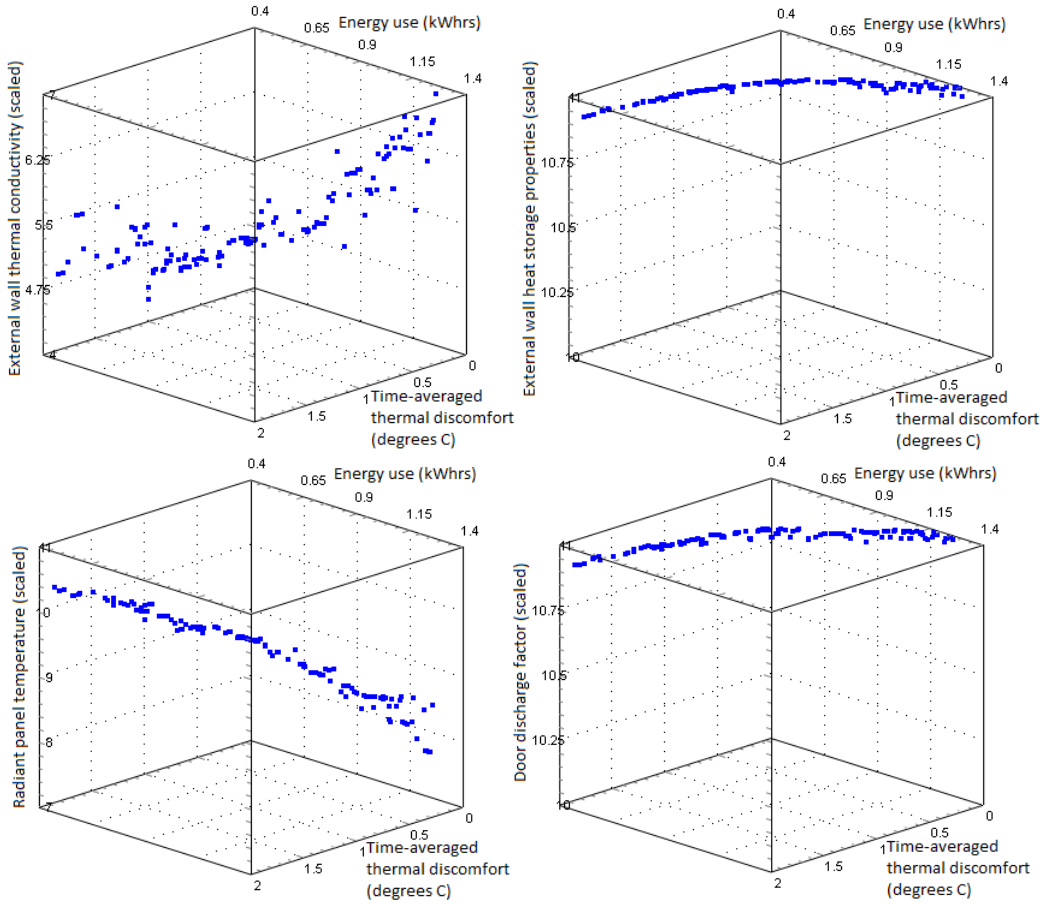

Figure B.16, continued overleaf

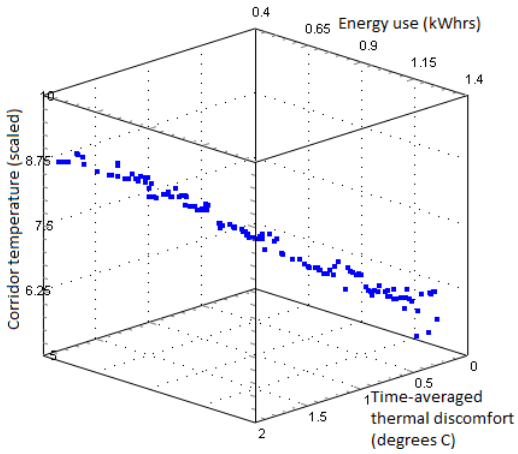

Figure B.16: Pareto front (as shown in Fig. B.15) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

Case: SumHot1_15_1-6,7-12,13-18,19-24_dv9_ot

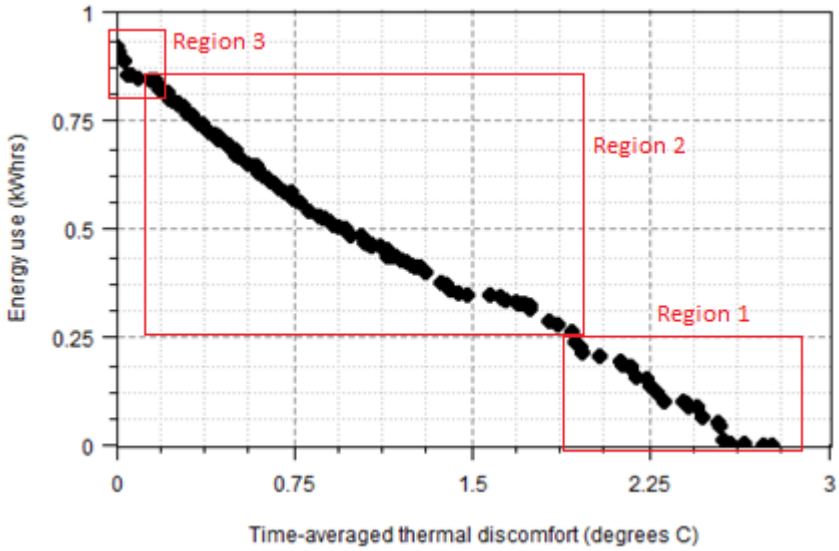

Figure B.17: Pareto front for the *SumHot1_15_1-6_dv9_ot* case shown as a trade-off between the two objectives, with distinct regions marked.

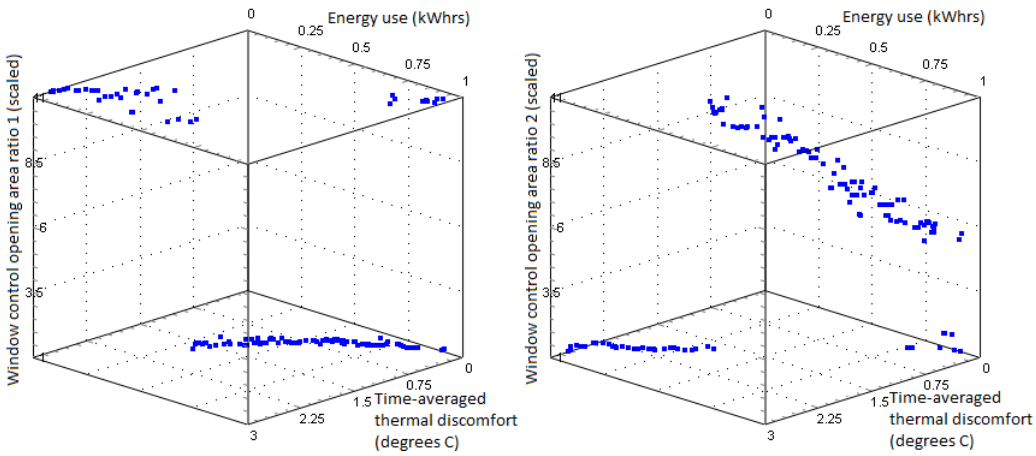

Figure B.18, continued overleaf

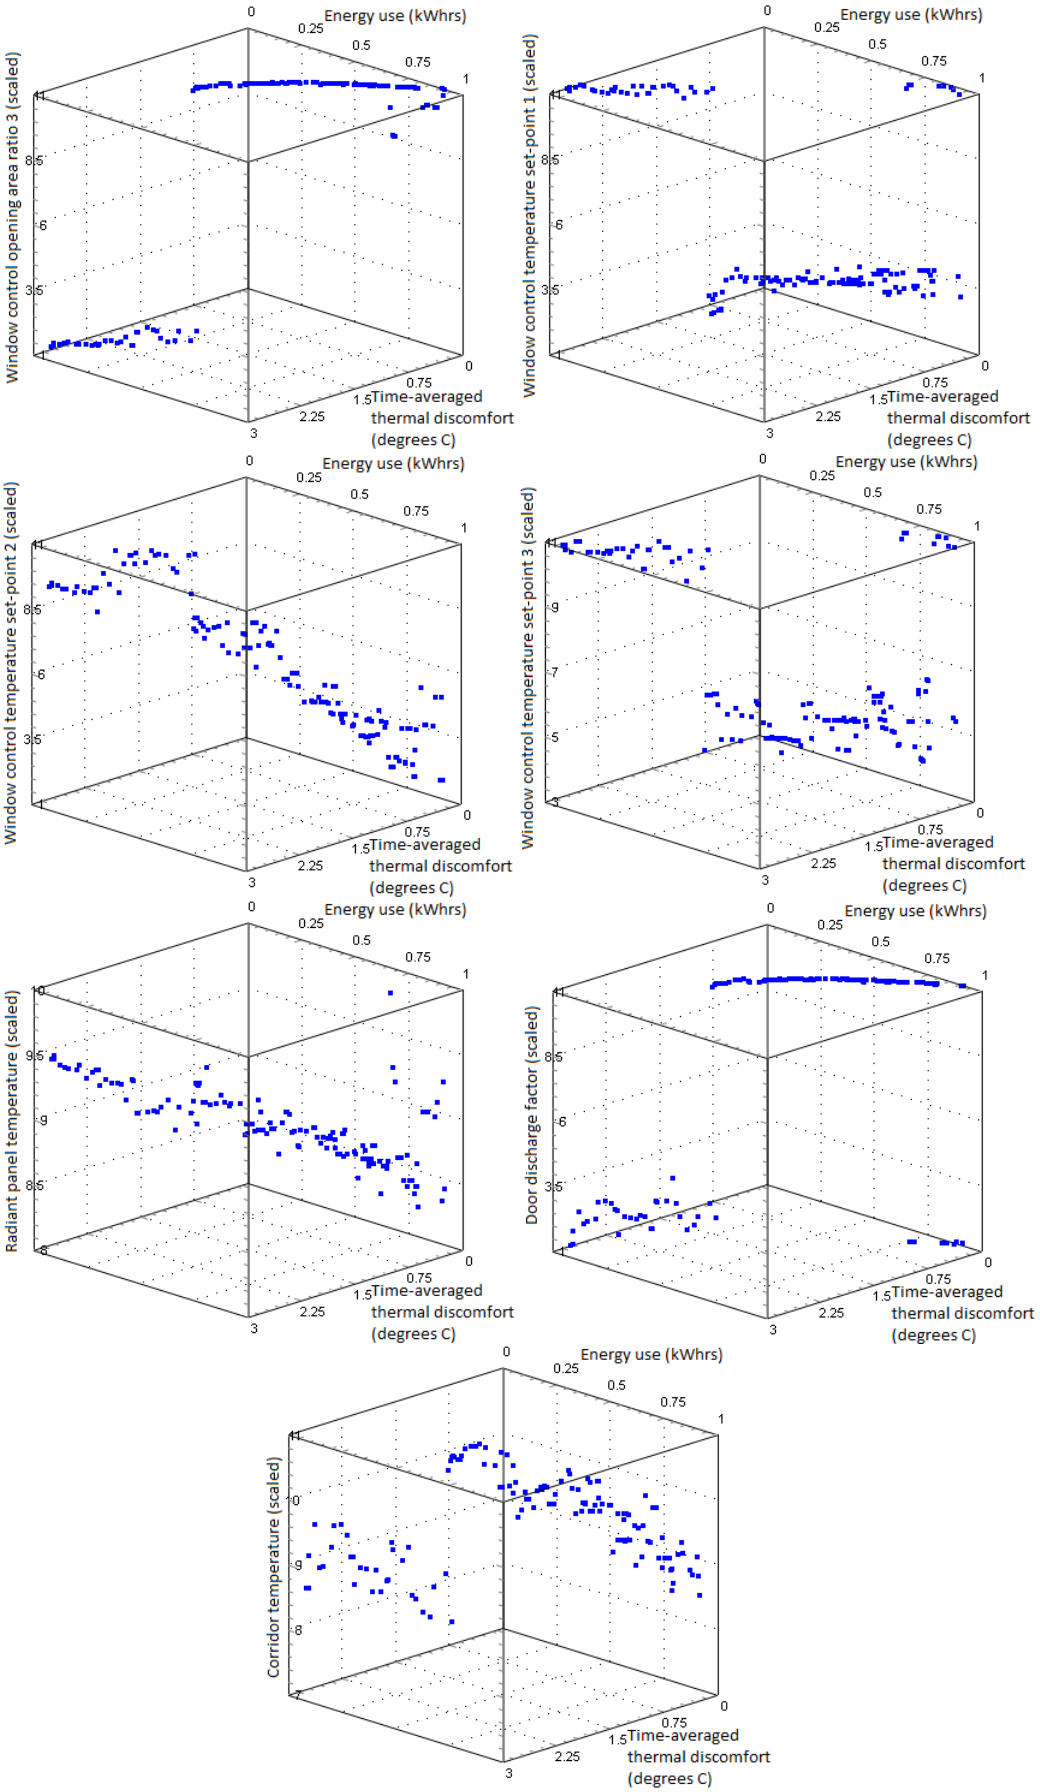

Figure B.18: Pareto front (as shown in Fig. B.17) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

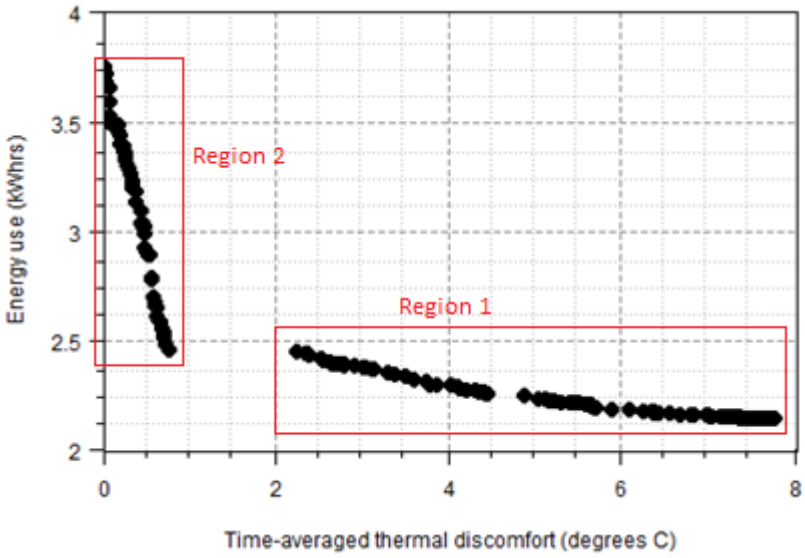

Figure B.19: Pareto front for the *SumHot1_15_7-12_dv9_ot* case shown as a trade-off between the two objectives, with distinct regions marked.

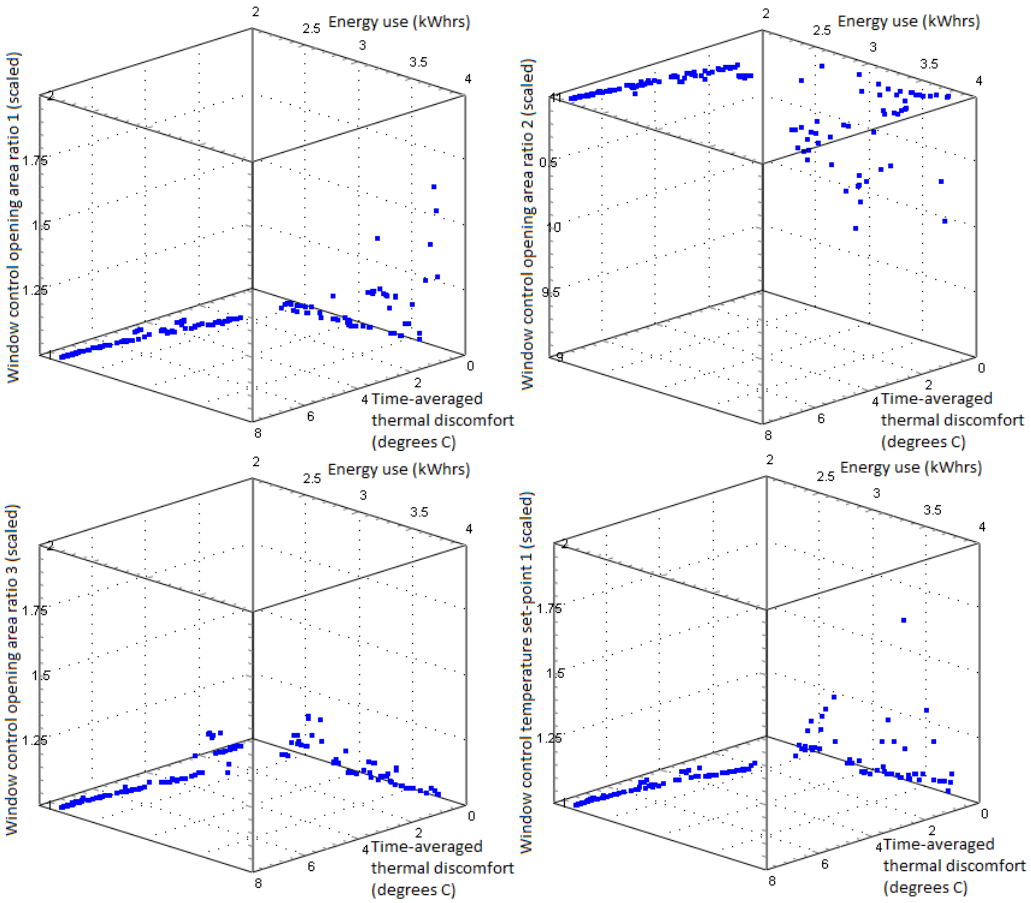

Figure B.20, continued overleaf

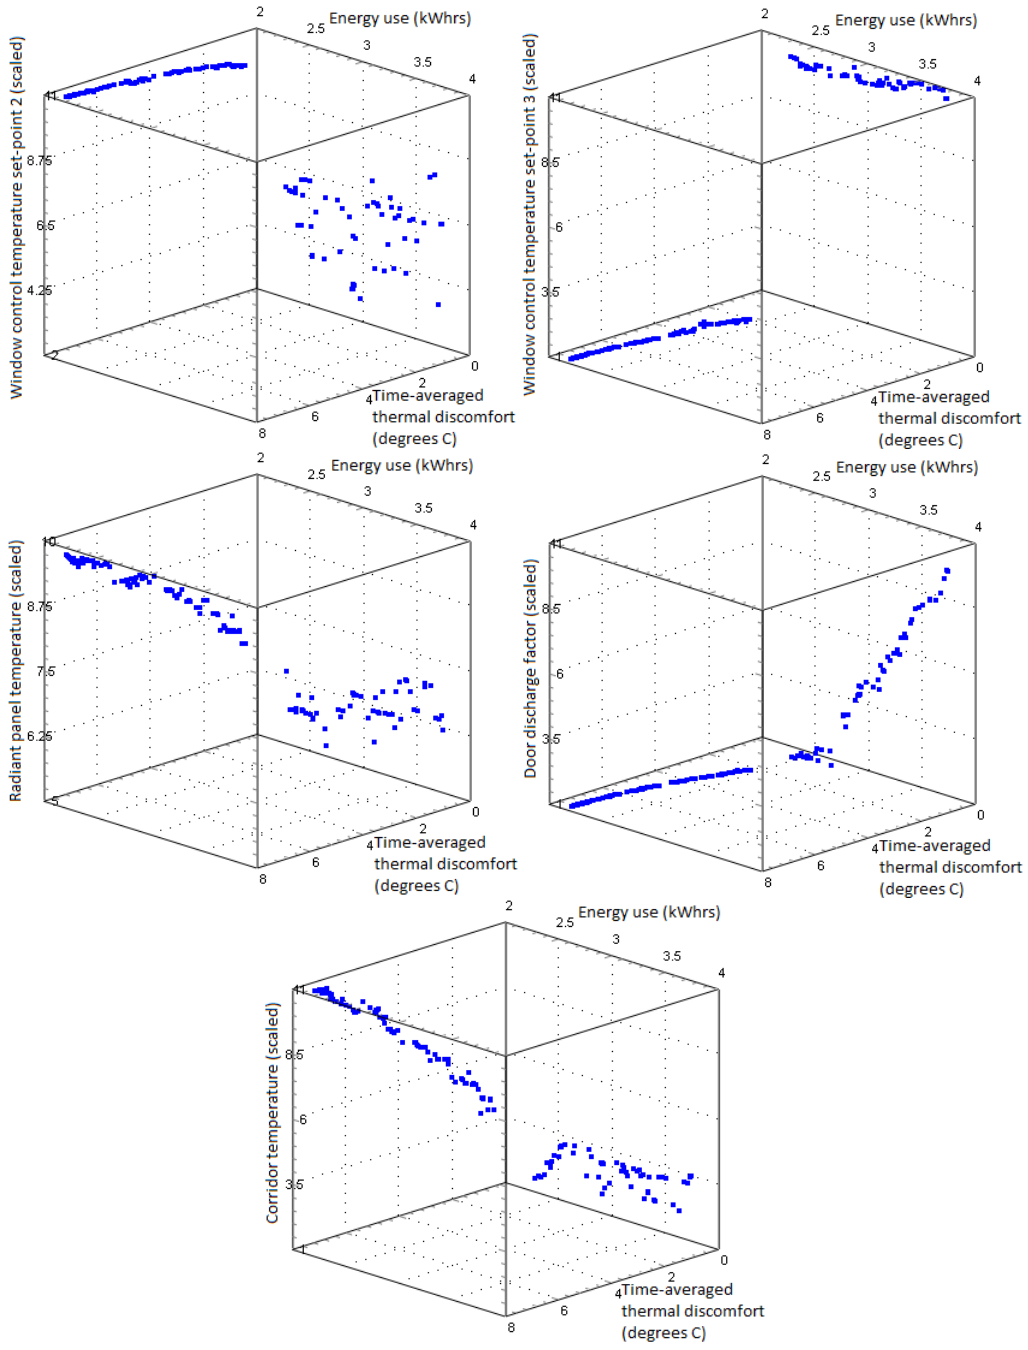

Figure B.20: Pareto front (as shown in Fig. B.19) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

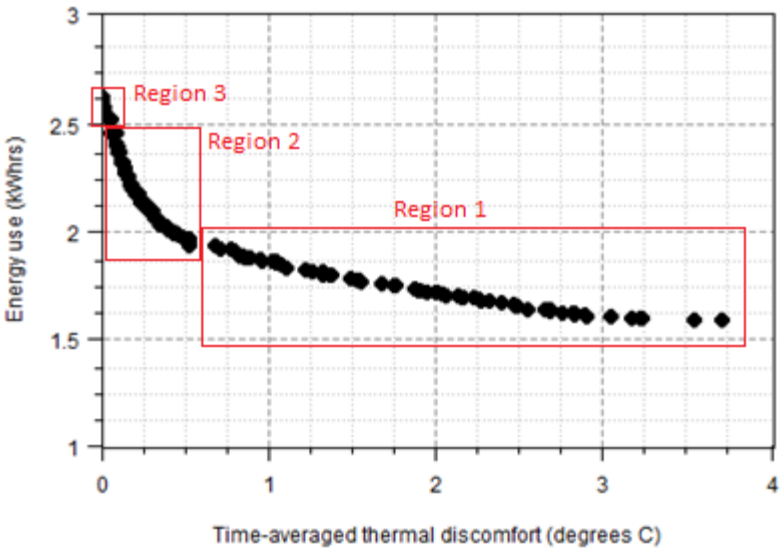

Figure B.21: Pareto front for the *SumHot1_15_13-18_dv9_ot* case shown as a trade-off between the two objectives, with distinct regions marked.

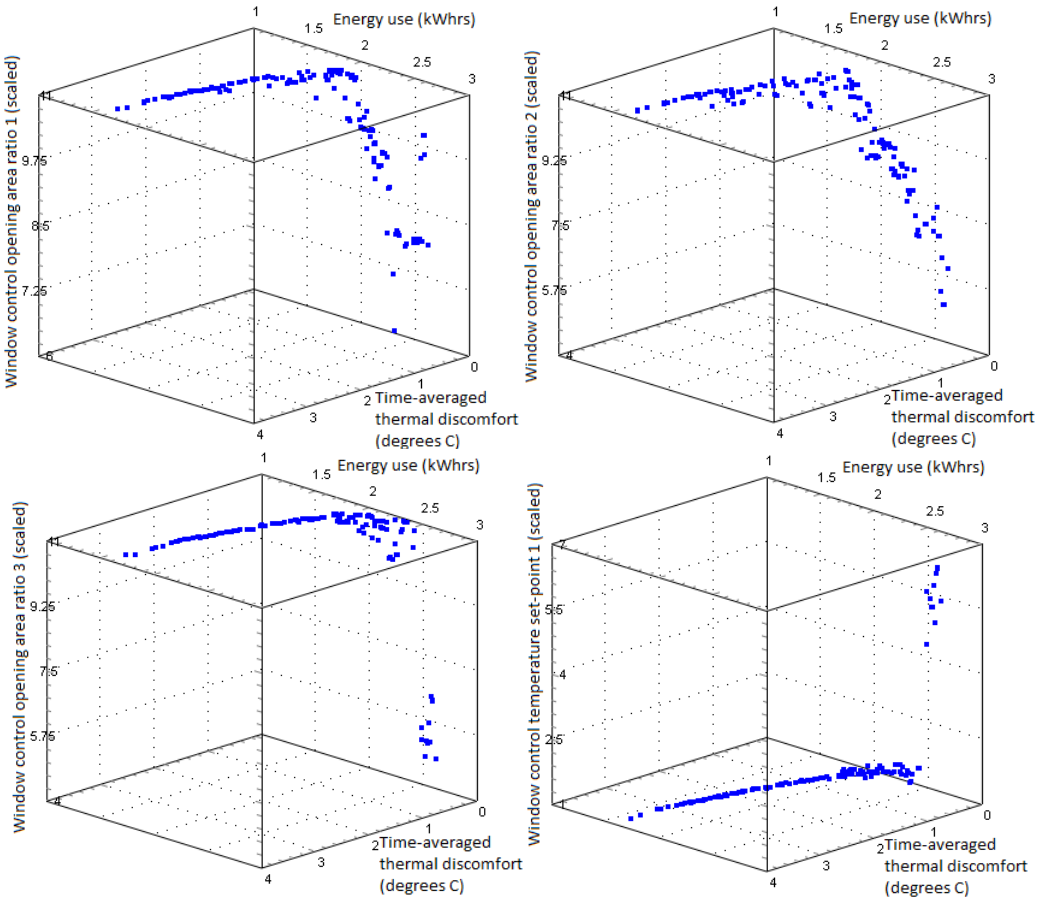

Figure B.22, continued overleaf

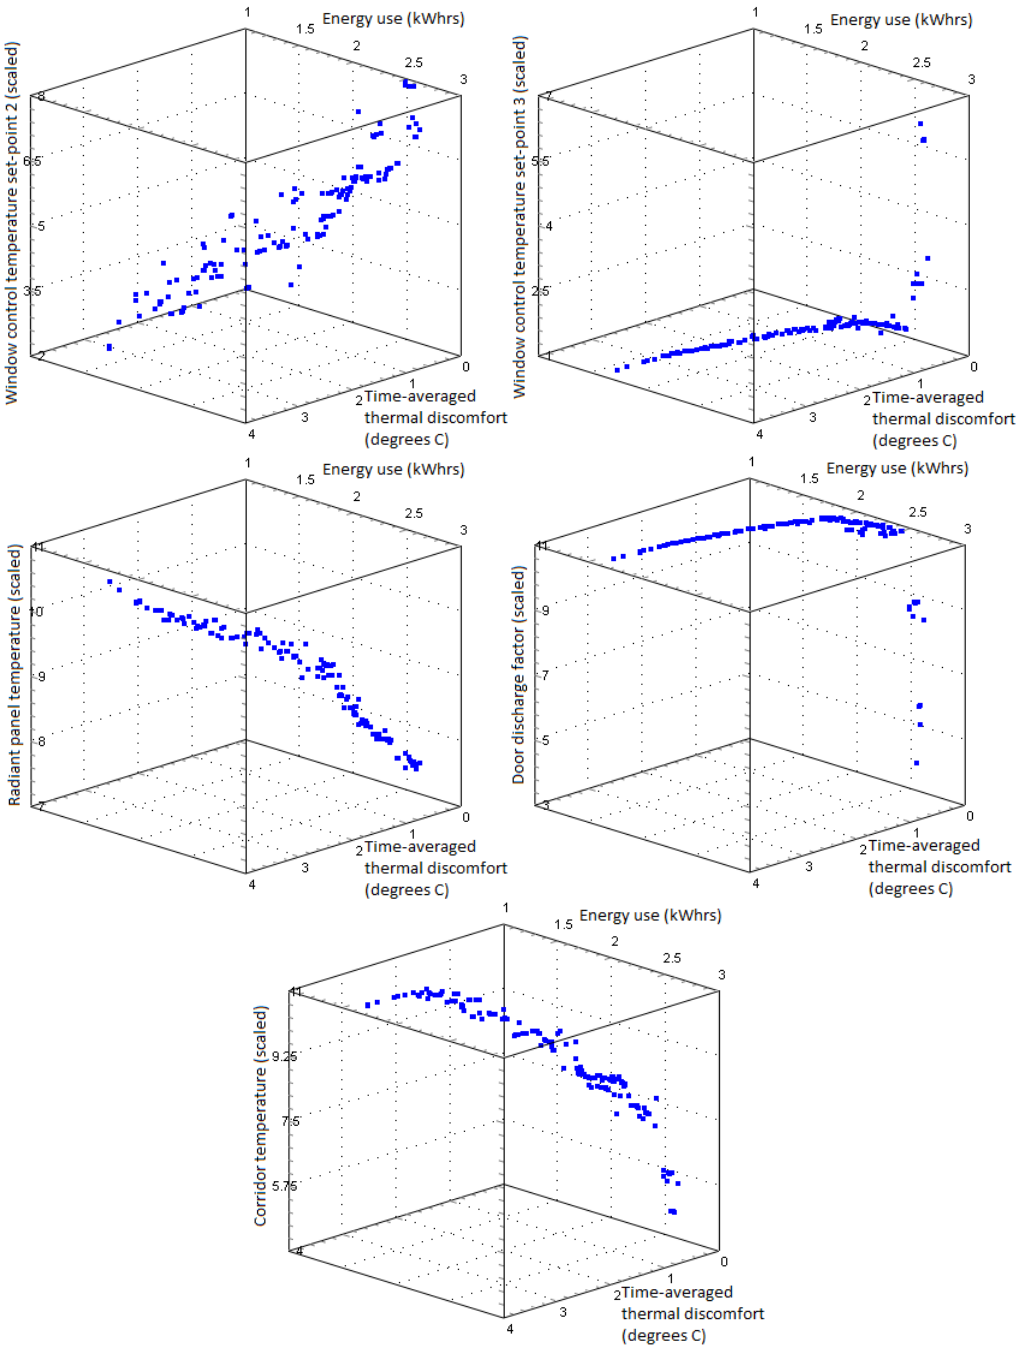

Figure B.22: Pareto front (as shown in Fig. B.21) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

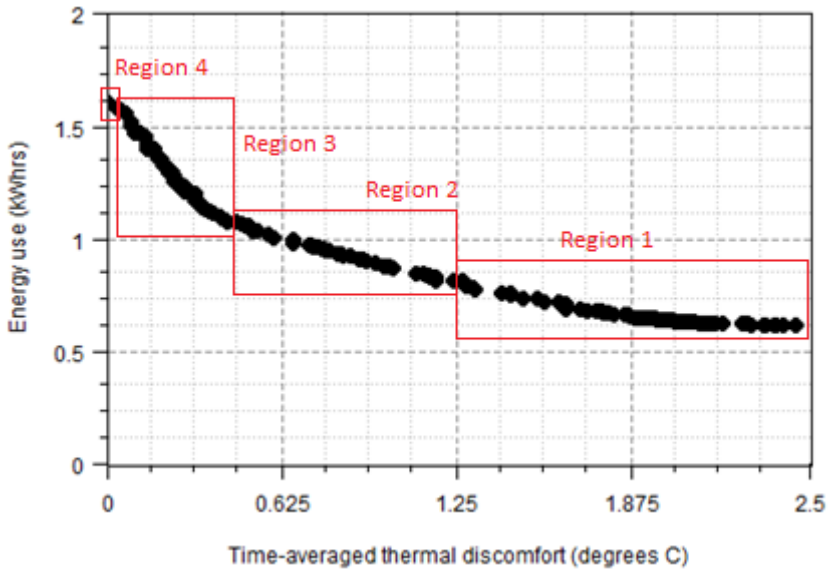

Figure B.23: Pareto front for the *SumHot1_15_19-24_dv9_ot* case shown as a trade-off between the two objectives, with distinct regions marked.

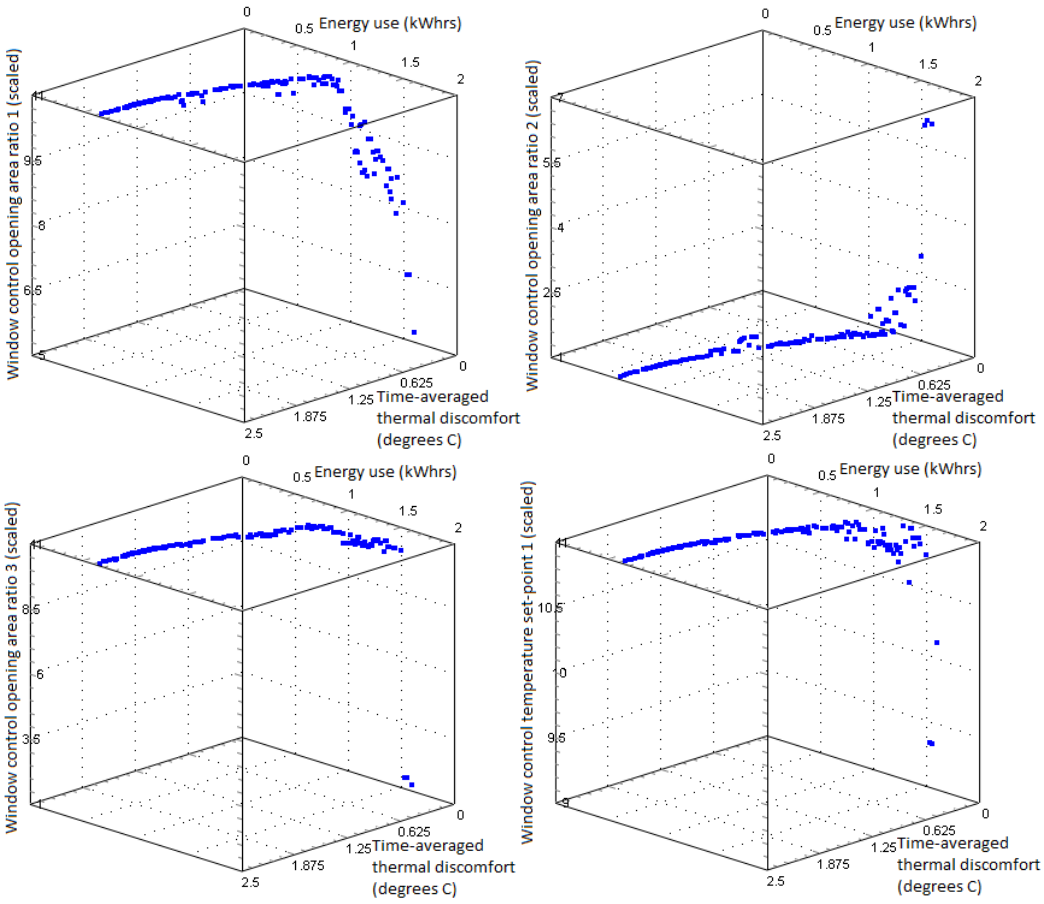

Figure B.24, continued overleaf

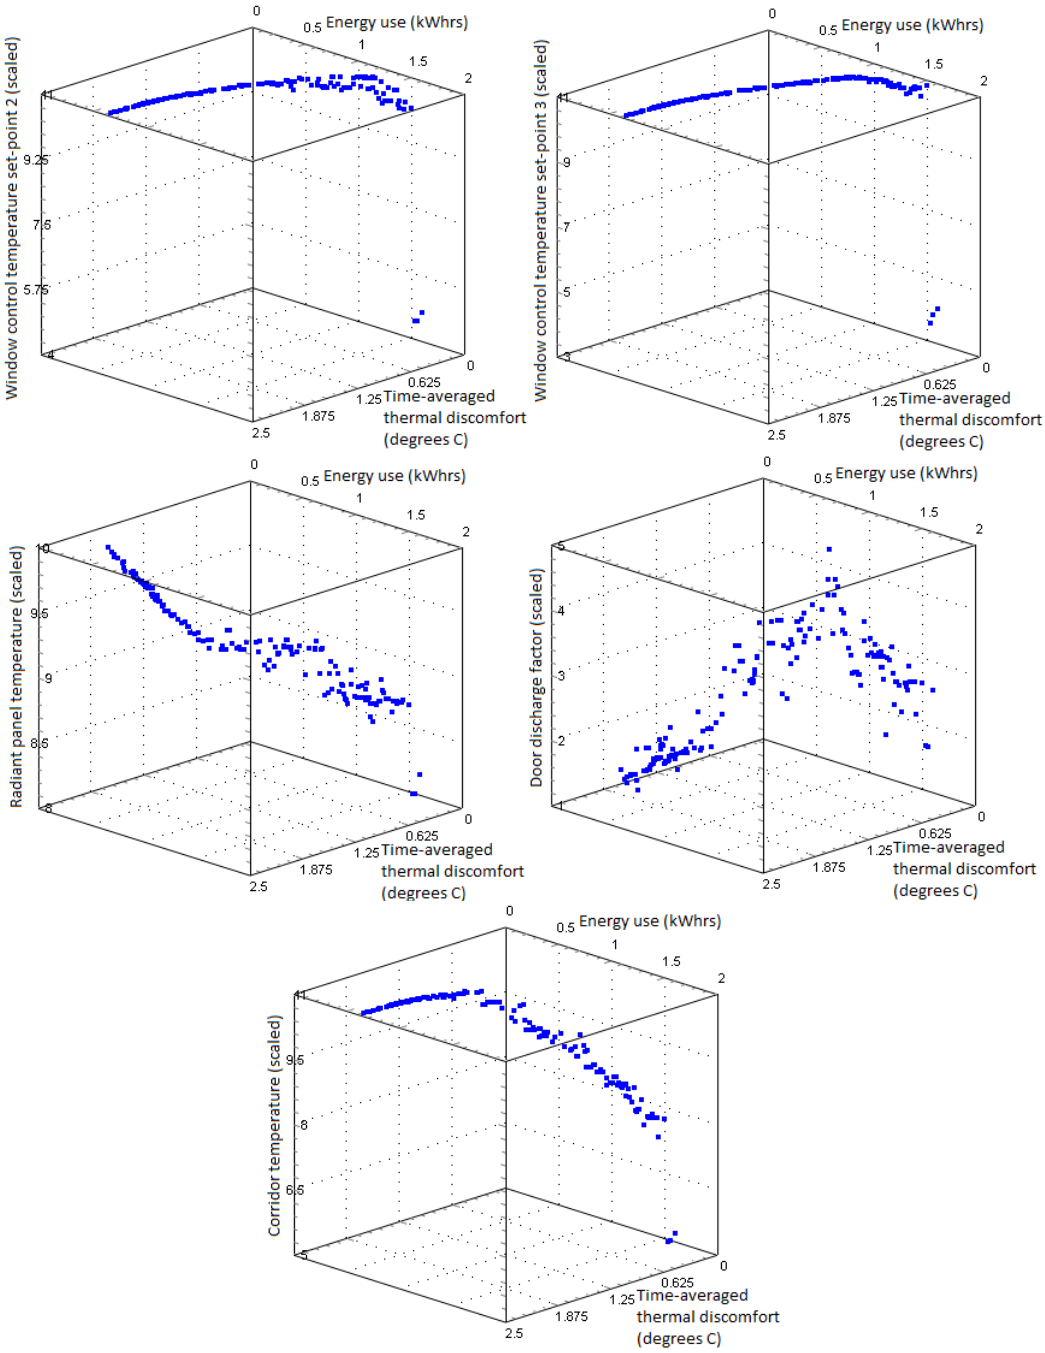

Figure B.24: Pareto front (as shown in Fig. B.23) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

7.3.2 Staff and visitor comfort optimisation

Case: SumHot1_13_1-24_dv5_pmvS

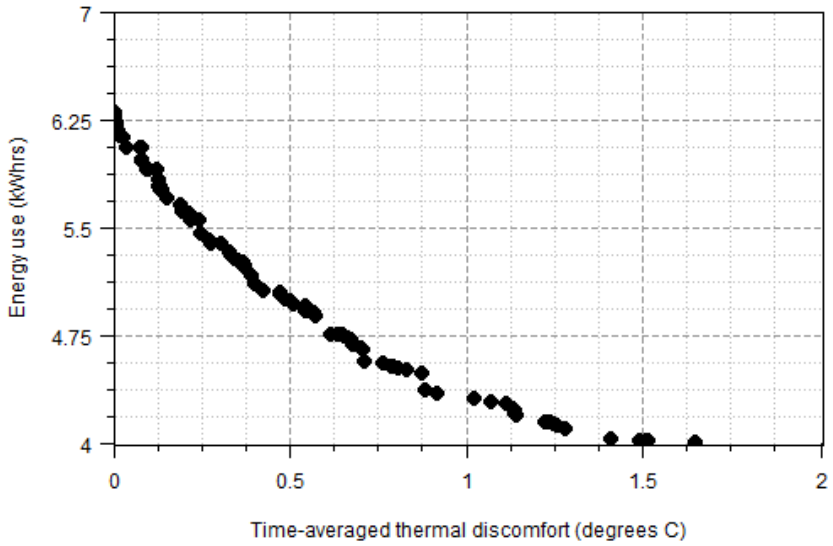

Figure B.25: Pareto front for the *SumHot1_13_1-24_dv5_pmvS* case shown as a trade-off between the two objectives.

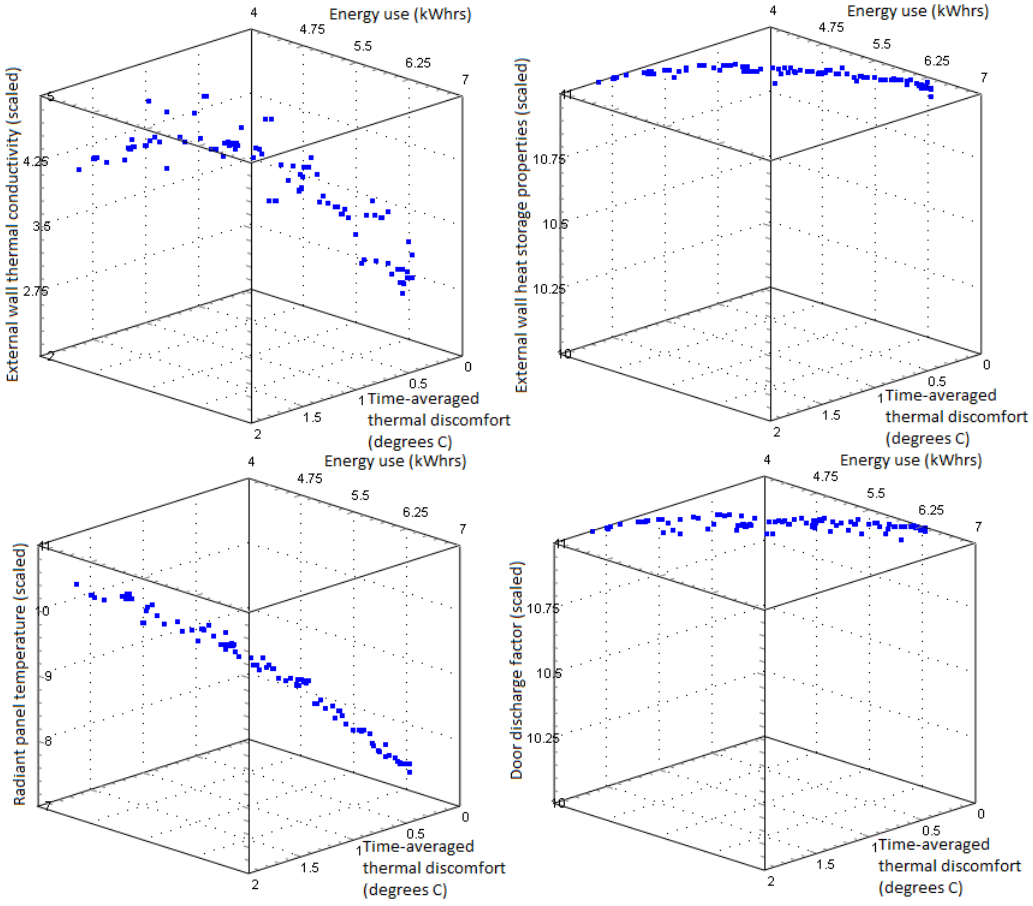

Figure B.26, continued overleaf

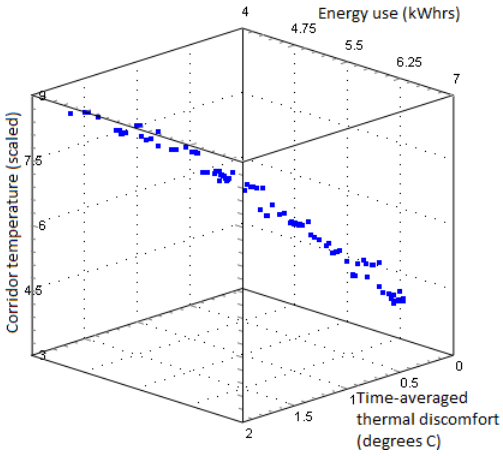

Figure B.26: Pareto front (as shown in Fig. B.25) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

Case: SumHot1_13_1-24_dv5_pmvV

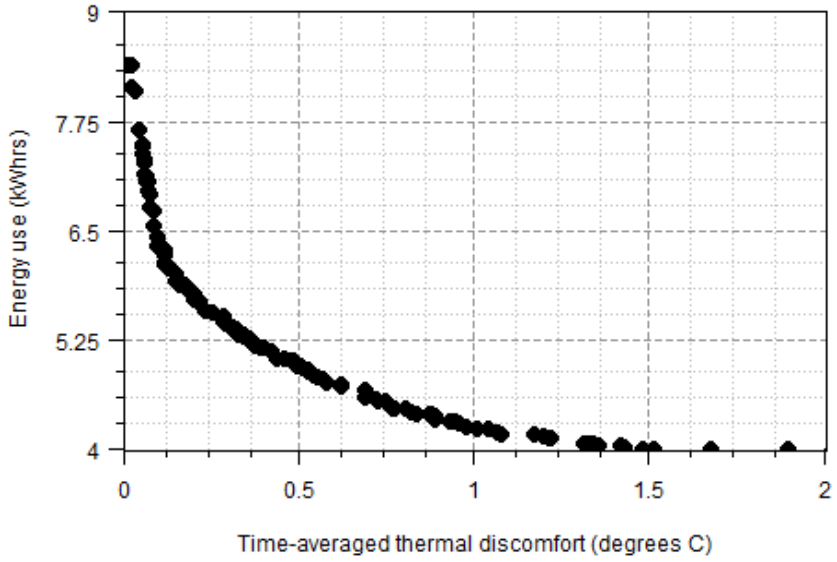

Figure B.27: Pareto front for the *SumHot1_13_1-24_dv5_pmvV* case shown as a trade-off between the two objectives.

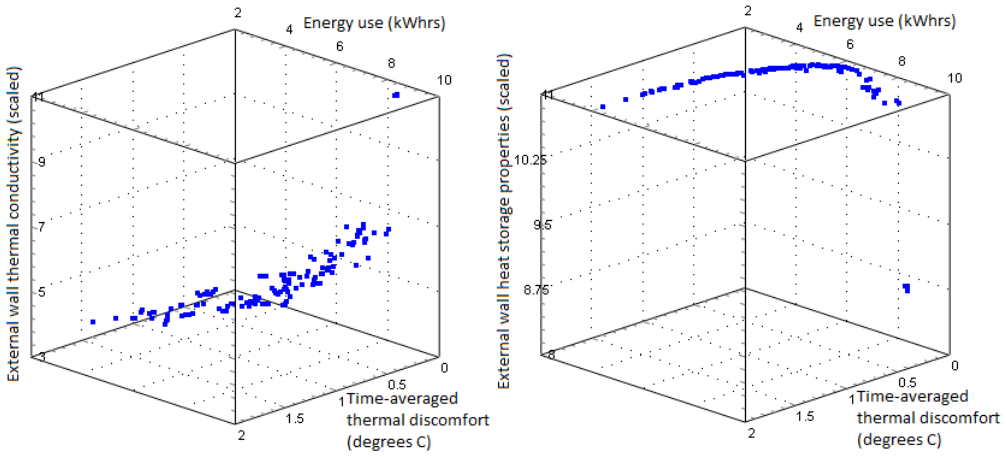

Figure B.28, continued overleaf

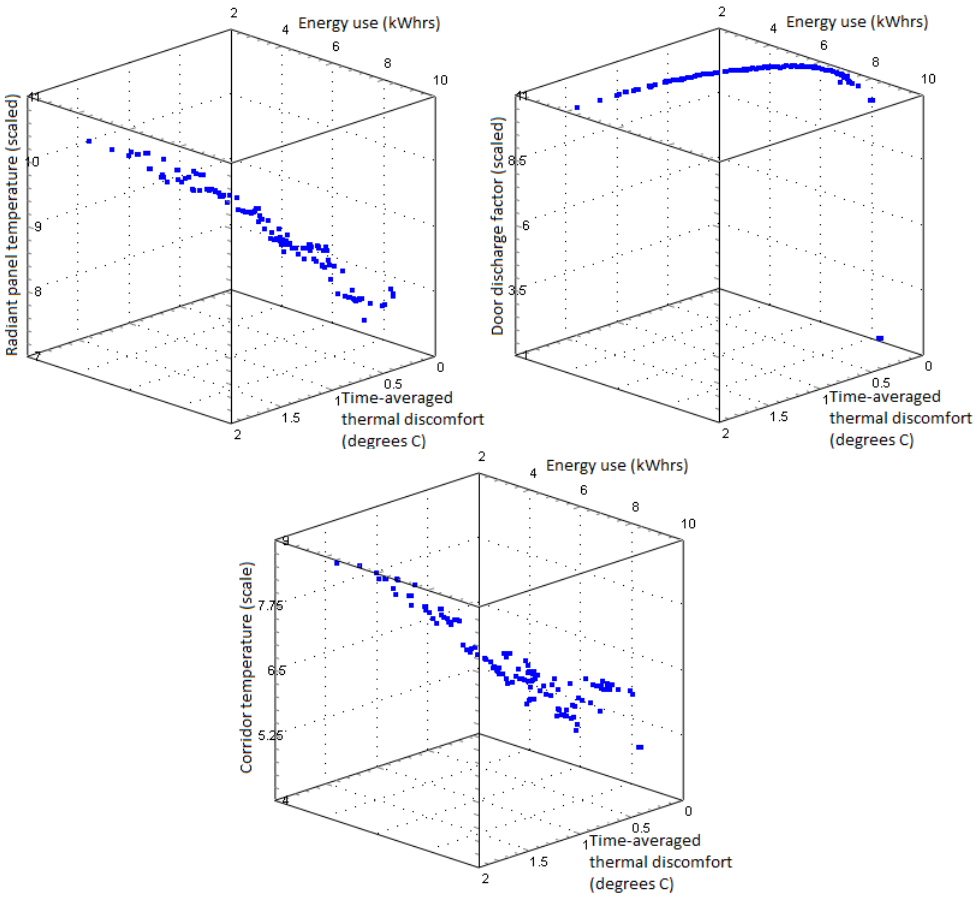

Figure B.28: Pareto front (as shown in Fig. B.27) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

Case: SumHot1_13_1-24_dv9_pmvS

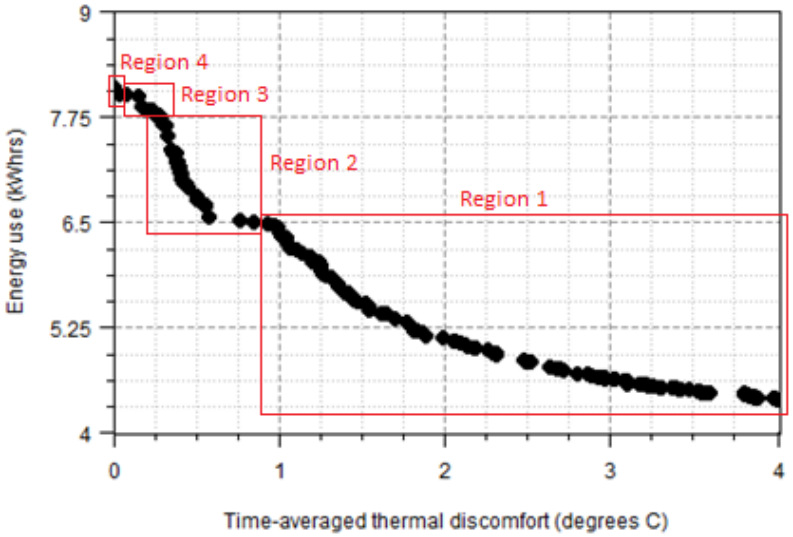

Figure B.29: Pareto front for the *SumHot1_13_1-24_dv9_pmvS* case shown as a trade-off between the two objectives.

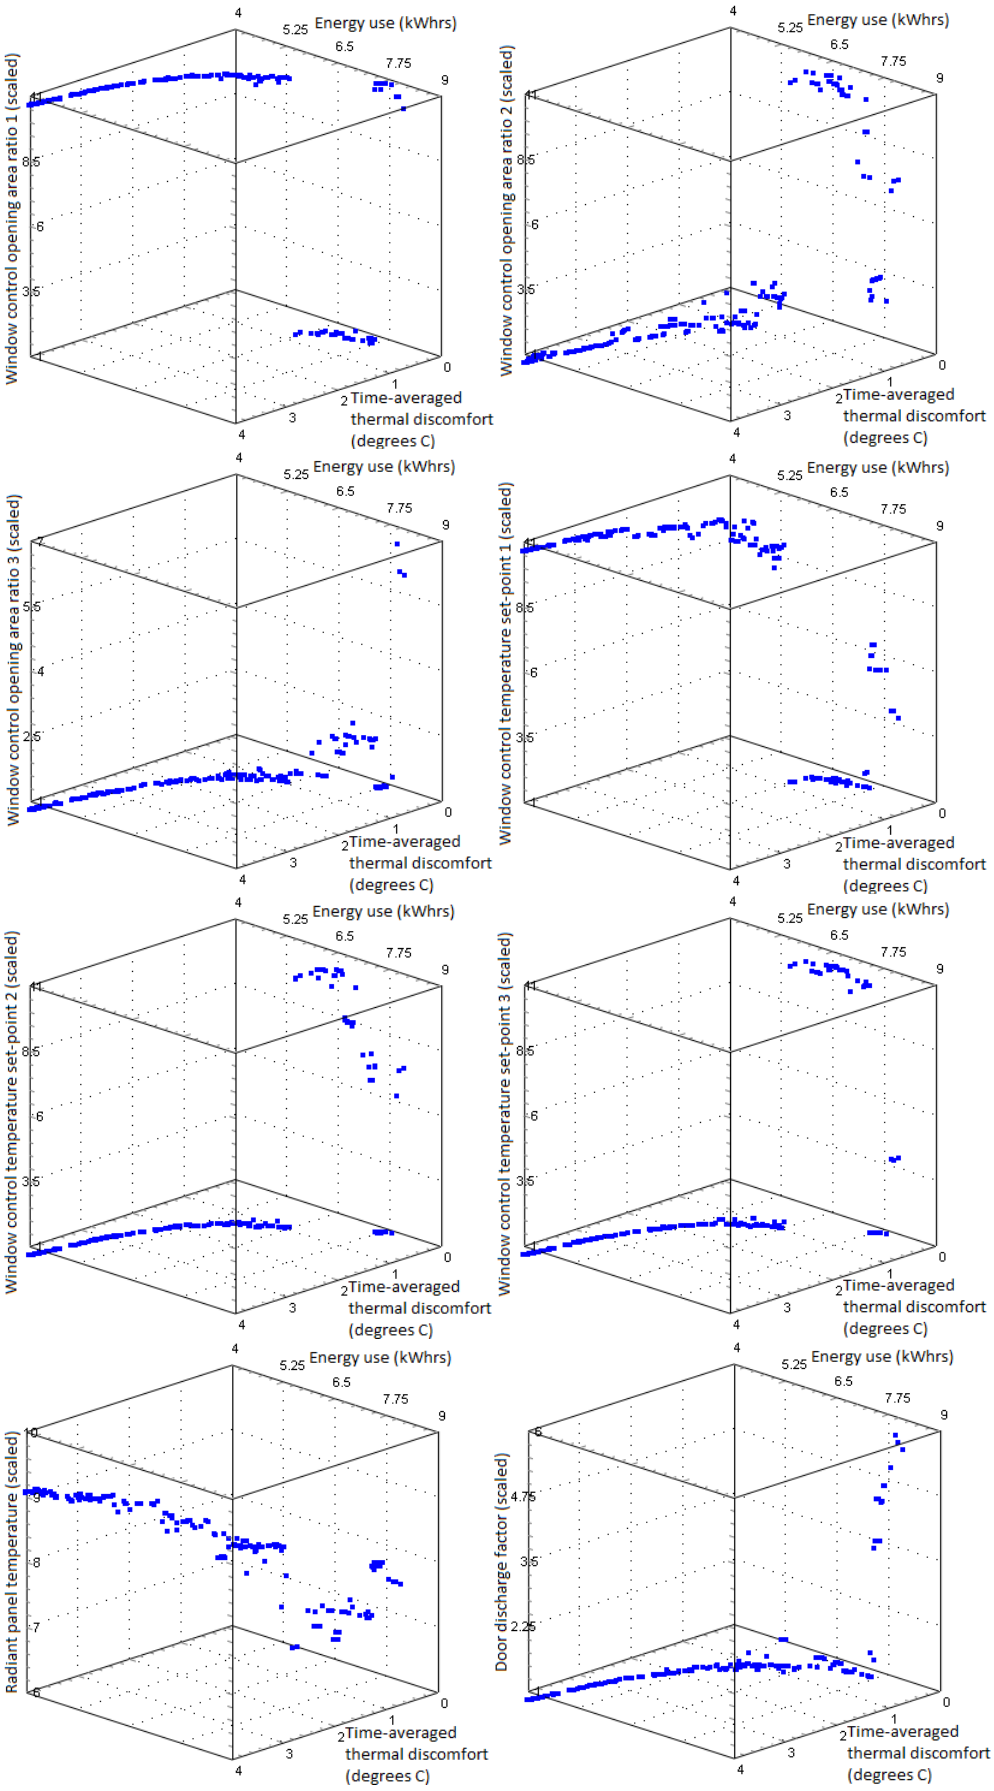

Figure B.30, continued overleaf

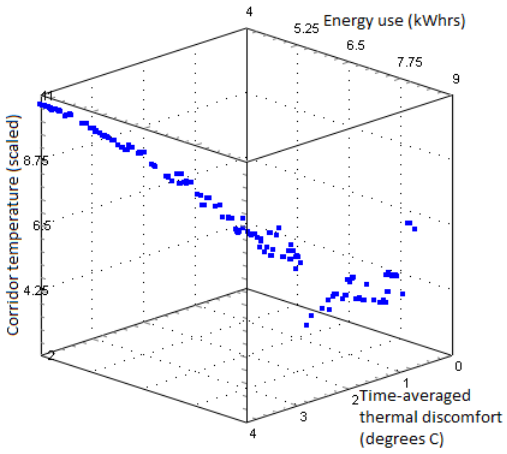

Figure B.30: Pareto front (as shown in Fig. B.29) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

Case: SumHot1_13_1-24_dv9_pmvV

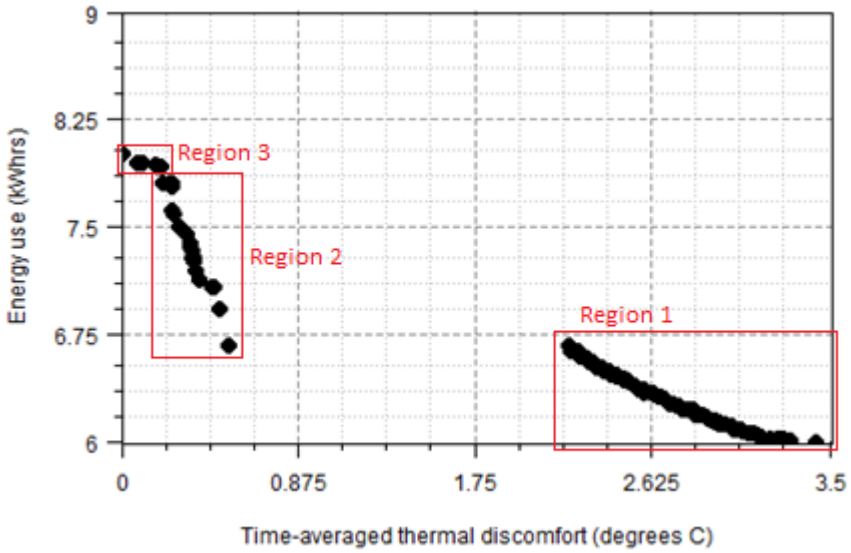

Figure B.31: Pareto front for the *SumHot1_13_1-24_dv9_pmvV* case shown as a trade-off between the two objectives.

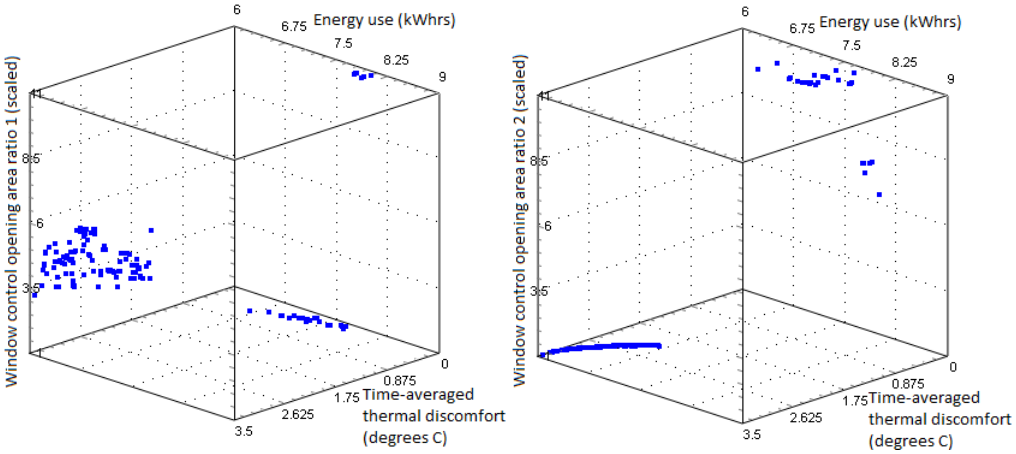

Figure B.32, continued overleaf

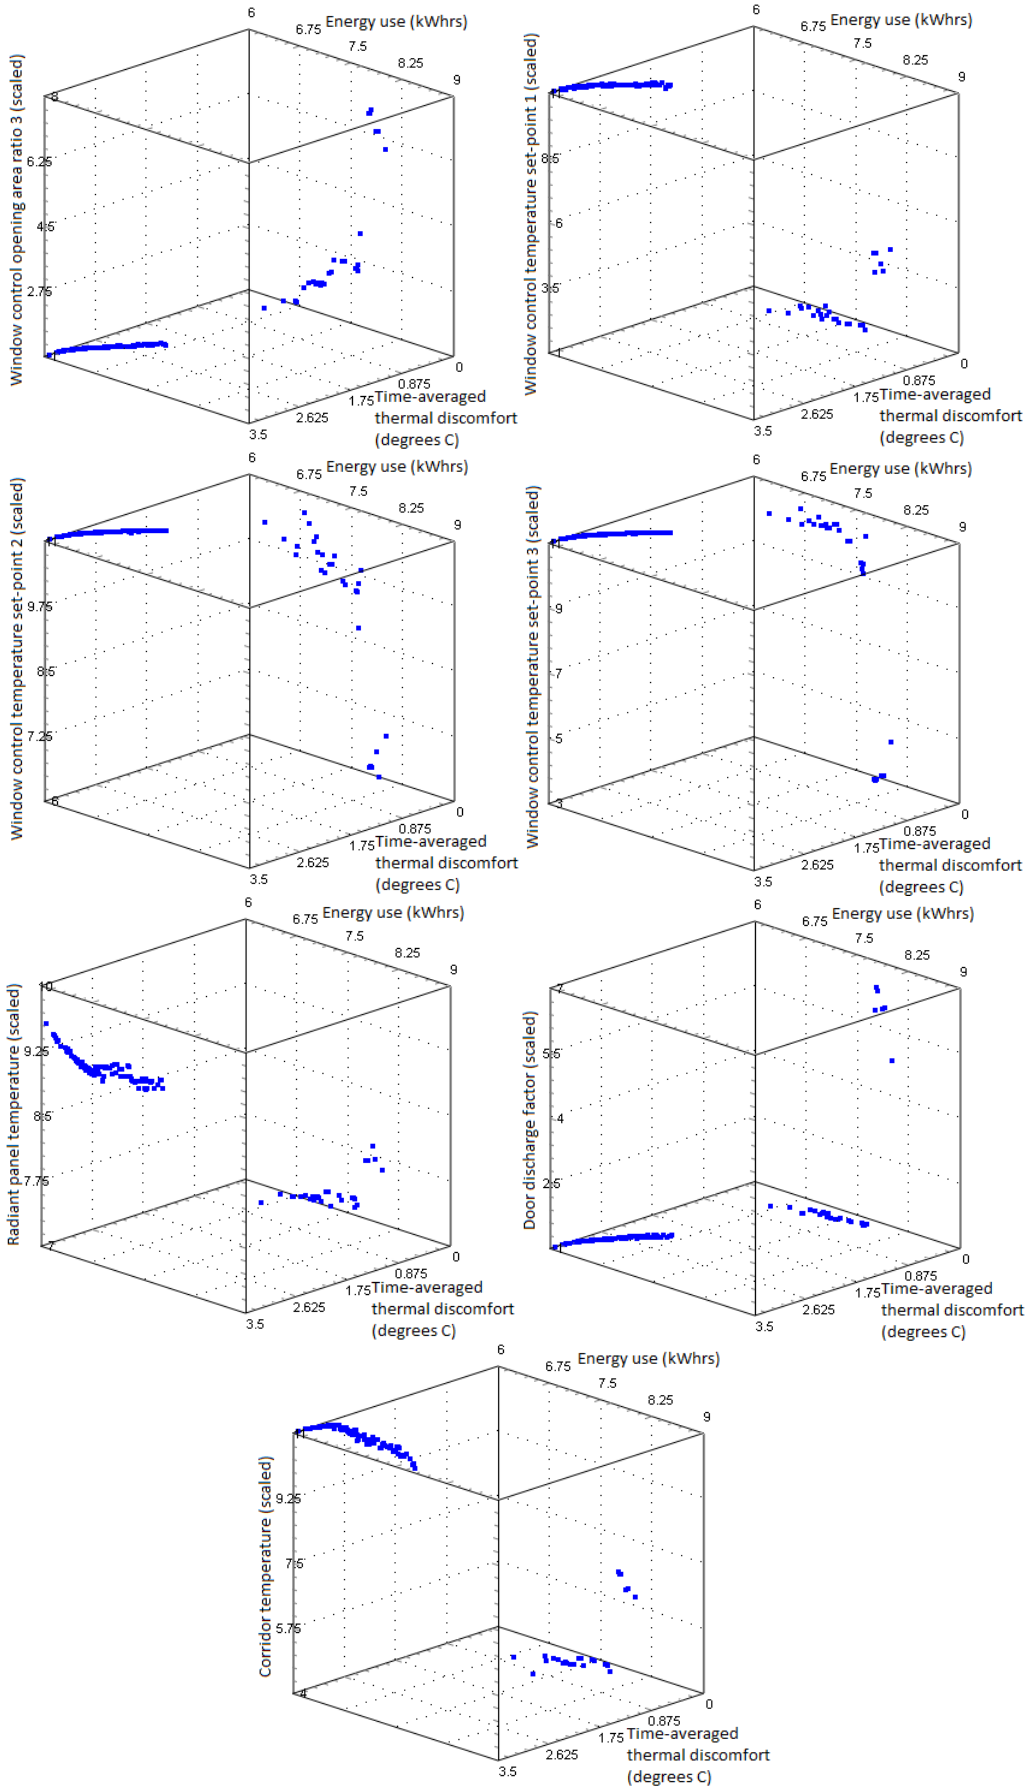

Figure B.32: Pareto front (as shown in Fig. B.31) plotted against each design variable in turn. All design variables were scaled to values of 1-11; for actual limits see beginning of section B.2.

B.3 Average Summer Conditions

B.3.1 Average summer conditions day 1

This set of results represents cases run under average summer conditions, day 1, as described in section 6.3.6.

Similarly to the extreme summer conditions, two separate samples were simulated for this climate case. The first was a sample of 5 design variables, of size 100 samples for the build set and 50 for the validation set. Design variables and ranges of values for this sample were:

- External wall thermal conductivity (0.1-1.5 W/mK)
- External wall thermal storage properties (500-2000 kg/m^3 & J/kgK)
- Radiant panel temperature (5-30 °C)
- Door discharge factor (0.1-0.9)
- Corridor air temperature (18-26

Initial analysis of metamodel performance for this sample revealed a consistent tendency for the optimisation to converge to the limits of the two external wall properties design variables, which was found to produce less reliable results due to increased inaccuracy at the edge of the metamodelling domains. As a result, all optimisations for this case were set to default values of external wall properties (as shown in Table 6.2), and were hence only in terms of the last three design variables:

- Radiant panel temperature (5-30 °C)
- Door discharge factor (0.1-0.9)
- Corridor temperature (18-26 °C)

The second was a sample of 9 design variables, of size 200 build samples and 100 validation samples. This sample considered the case of allowing the window to be opened. Design variables and ranges of values were:

- Window control opening area ratios (3 variables) (0.1-11)
- Window control temperature set-points (3 variables) (18-28 °C)
- Radiant panel temperature (5-30 °C)
- Door discharge factor (0.1-0.9)
- Corridor air temperature (18-26 °C)

Case: SumAve1_9_1-24_dv5_ot

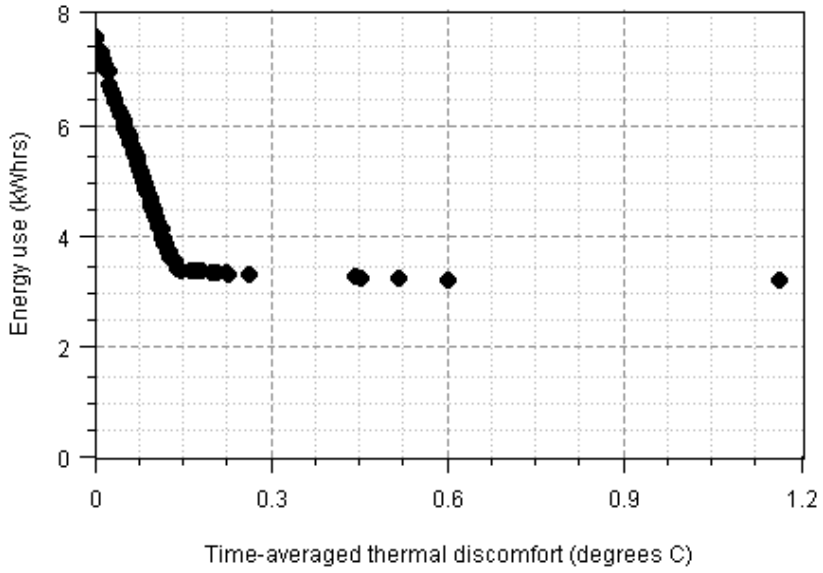

Figure B.33: Pareto front for the *SumAve1_9_1-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 discomfort:

- Energy use: 7.539 kWh
- Radiant panel temperature: 9.530 (26.33 °C)
- Door discharge factor: 1.150 (0.112)
- Corridor temperature: 2.909 (19.53 °C)

Case: SumAve1_13_1-24_dv5_ot

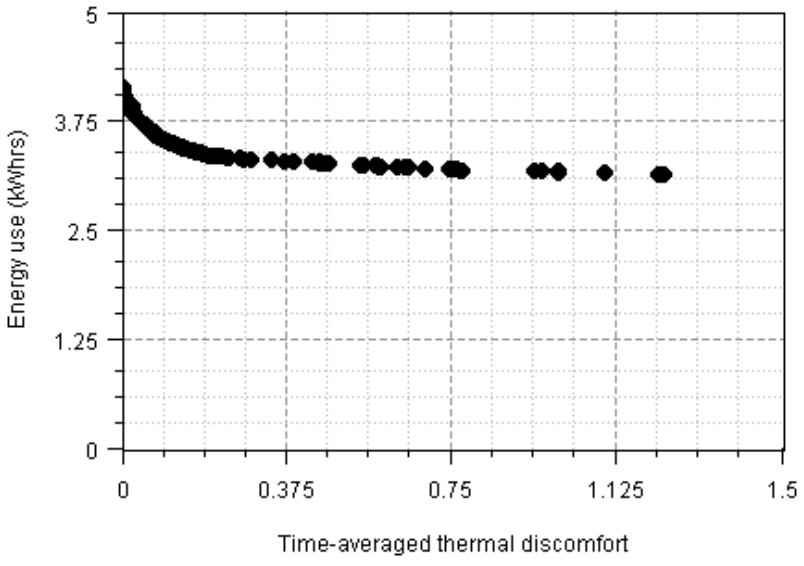

Figure B.34: Pareto front for the *SumAve1_13_1-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort:

- Energy use: 4.146 kWh
- Radiant panel temperature: 8.932 (24.83 °C)
- Door discharge factor: 1.001 (0.1)
- Corridor temperature: 5.429 (21.54 °C)

Case: SumAve1_9_1-6,7-12,13-18,19-24_dv5_ot

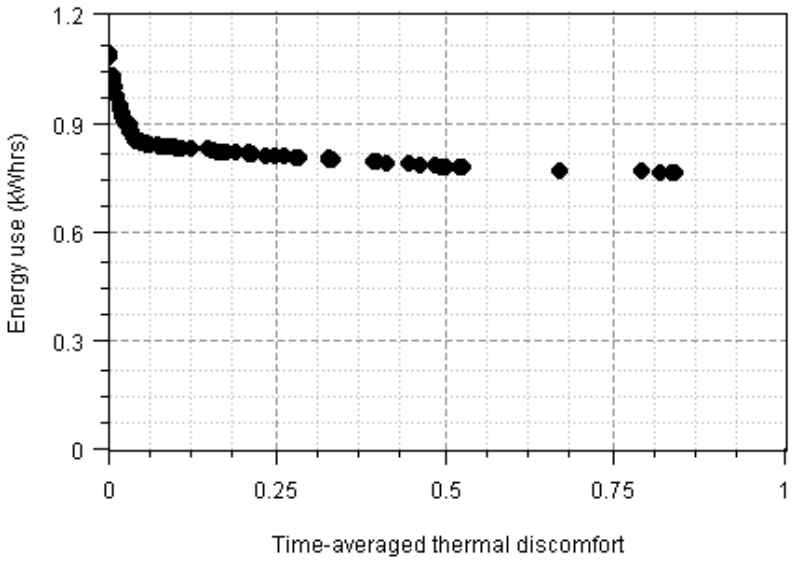

Figure B.35: Pareto front for the *SumAve1_9_1-6_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort for the 1-6 period:

- Energy use: 1.087 (kWh)
- Radiant panel temperature: 8.812 (24.53 °C)
- Door discharge factor: 1.144 (0.112)
- Corridor temperature: 5.000 (21.2 °C)

Objective function and design variable values for the solution at 0 °C discomfort for the 7-12 period:

- Energy use: 2.404 kWh
- Radiant panel temperature: 9.642 (26.61 °C)
- Door discharge factor: 1.016 (0.101)
- Corridor temperature: 1.249 (18.20 °C)

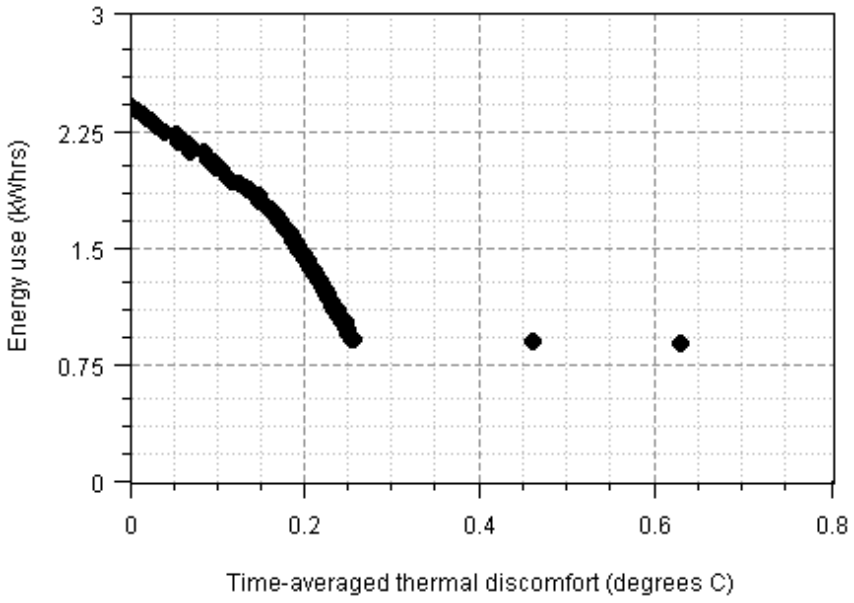

Figure B.36: Pareto front for the *SumAve1_9_7-12_dv5_ot* case shown as a trade-off between the two objectives.

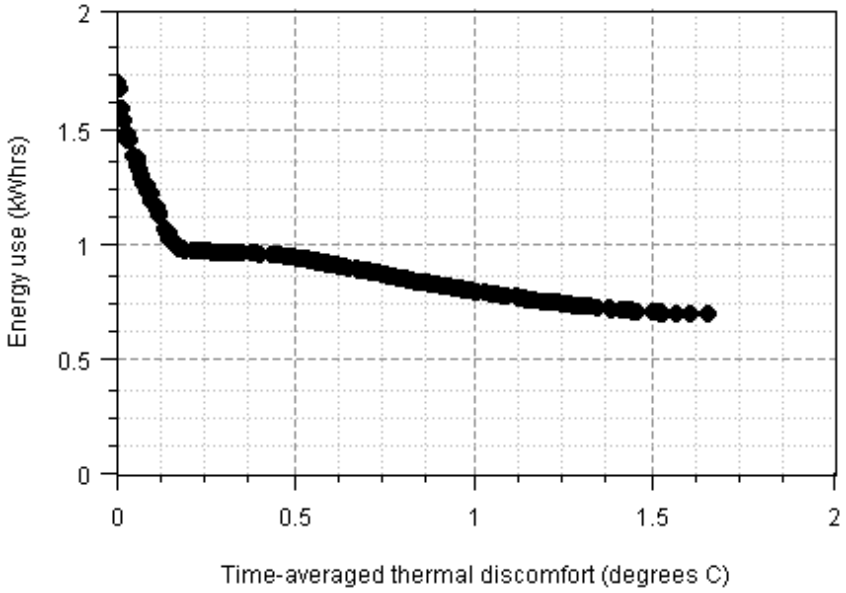

Figure B.37: Pareto front for the *SumAve1_9_13-18_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort for the 13-18 case:

- Energy use: 1.688 kWh
- Radiant panel temperature: 9.529 (26.32 °C)
- Door discharge factor: 1.039 (0.103)
- Corridor temperature: 2.714 (19.37 °C)

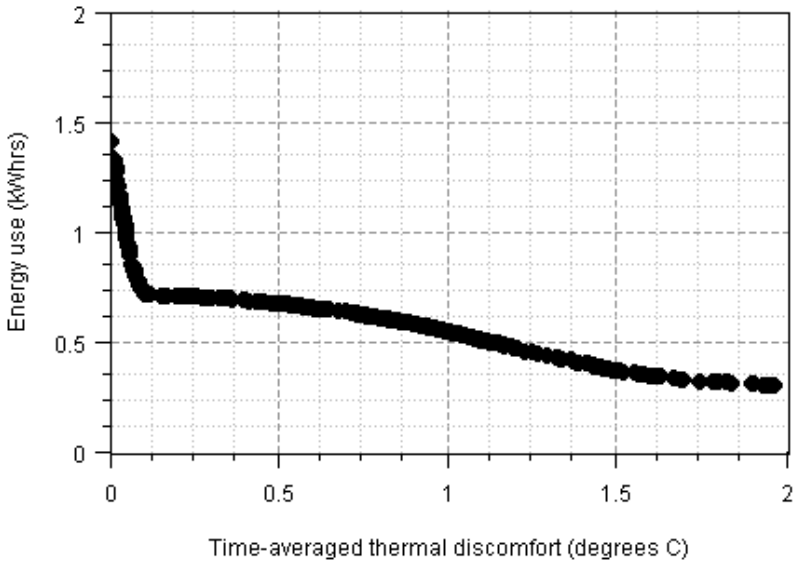

Figure B.38: Pareto front for the *SumAve1_9_19-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort for the 19-24 case:

- Energy use: 1.414 kWh
- Radiant panel temperature: 9.368 (25.92 °C)
- Door discharge factor: 1.266 (0.121)
- Corridor temperature: 4.345 (20.68 °C)

Case: SumAve1_13_1-24_dv9_ot

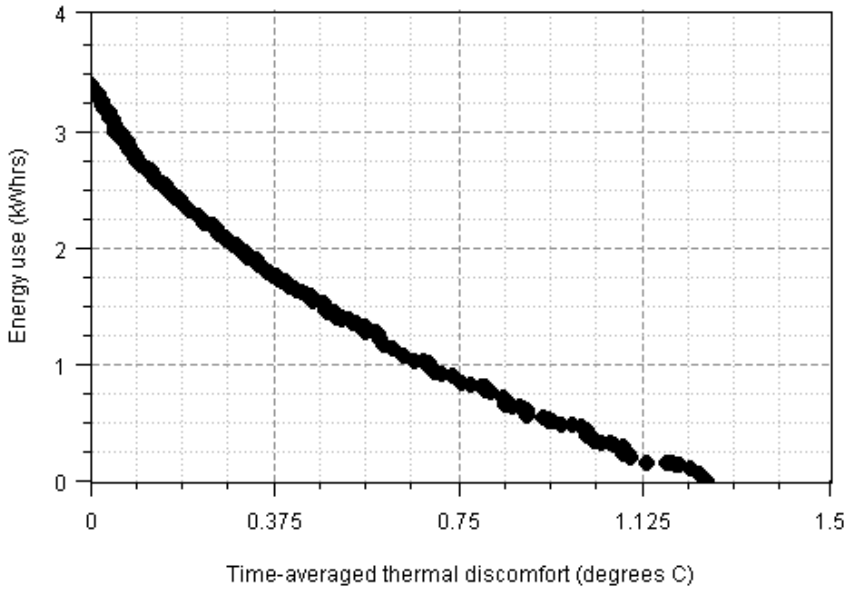

Figure B.39: Pareto front for the *SumAve1_13_1-24_dv9_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solutions at 0 °C discomfort and 0 kWh energy use:

- 0 °C thermal discomfort:
    - Energy use: 3.372 kWh
    - Window control opening area ratios: 1.361 / 5.774 / 10.986 (0.493 / 5.304 / 10.985)
    - Window control temperature set-points: 3.104 / 1.470 / 5.844 (20.10 / 20.48 / 24.12 °C)
    - Radiant panel temperature: 8.357 (23.39 °C)
    - Door discharge factor: 10.932 (0.895)
    - Corridor temperature: 7.606 (23.28 °C)
- 0 kWh energy use:
    - Thermal discomfort: 1.248 °C
    - Window control opening area ratios: 1.008 / 10.970 / 10.985 (0.109 / 10.97 / 10.98)
    - Window control temperature set-points: 2.632 / 8.464 / 5.111 (19.63 / 25.88 / 26.75 °C)
    - Radiant panel temperature: 9.888 (27.22 °C)
    - Door discharge factor: 10.871 (0.89)
    - Corridor temperature: 10.987 (25.99 °C)

B.3.2 Average summer conditions day 2

Case: SumAve2_9_1-24_dv5_ot

Objective function and design variable values were also extracted for the solution at 0 °C discomfort:

- Energy use: 9.175 kWh
- Radiant panel temperature: 9.653 (26.63 °C)
- Door discharge factor: 1.013 (0.101)
- Corridor temperature: 1.573 (18.46 °C)

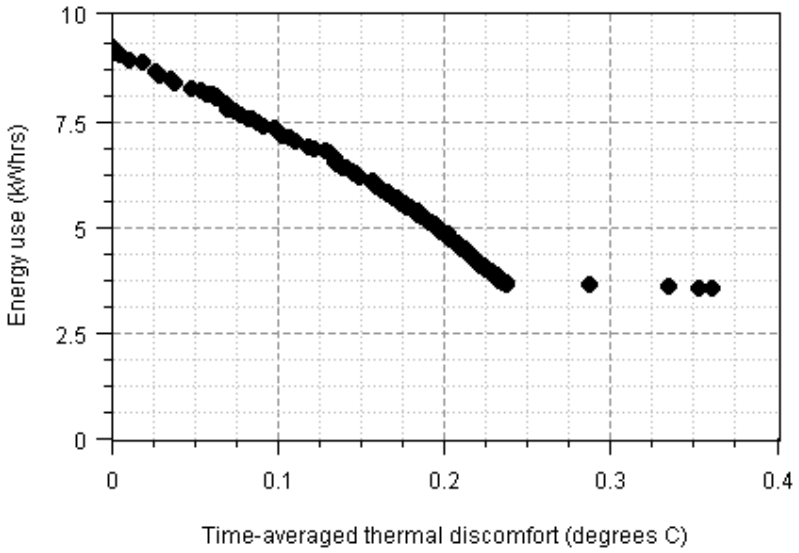

Figure B.40: Pareto front for the *SumAve2_9_1-24_dv5_ot* case shown as a trade-off between the two objectives.

Case: SumAve2_13_1-24_dv5_ot

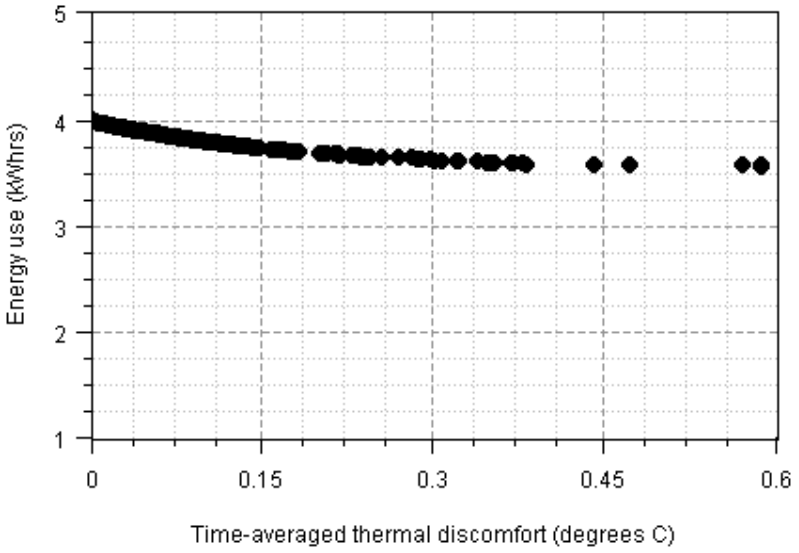

Figure B.41: Pareto front for the *SumAve2_13_1-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort:

- Energy use: 3.991 kWh
- Radiant panel temperature: 8.595 (23.99 °C)
- Door discharge factor: 1.001 (0.100)
- Corridor temperature: 6.189 (22.15 °C)

Case: SumAve2_9_1-6,7-12,13-18,19-24_dv5_ot

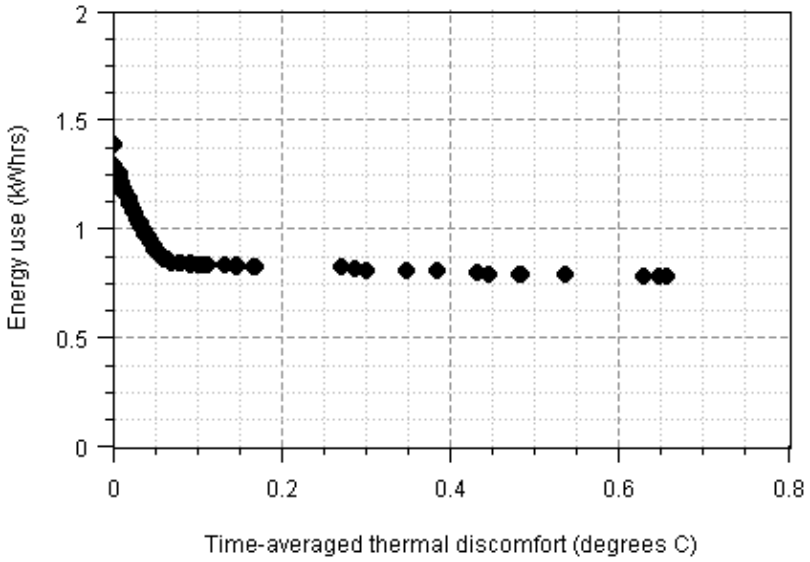

Figure B.42: Pareto front for the *SumAve2_9_1-6_dv5_ot* case shown as a trade-off between the two objectives.

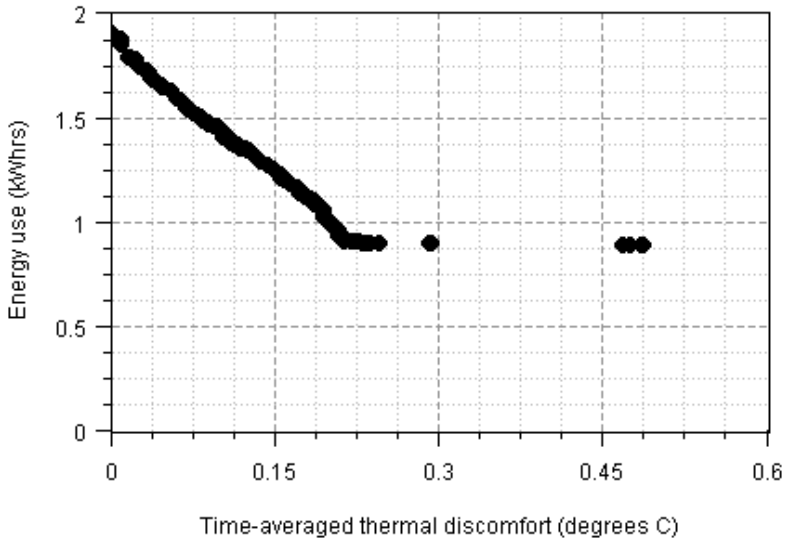

Figure B.43: Pareto front for the *SumAve2_9_7-12_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort for the 1-6 case:

- Energy use: 1.388 kWh
- Radiant panel temperature: 9.056 (25.14 °C)
- Door discharge factor: 1.386 (0.131)
- Corridor temperature: 4.186 (20.51 °C)

Objective function and design variable values for the solution at 0 °C discomfort for the 7-12 period:

- Energy use: 1.900 kWh
- Radiant panel temperature: 9.504 (26.26 °C)
- Door discharge factor: 1.026 (0.102)
- Corridor temperature: 2.269 (19.02 °C)

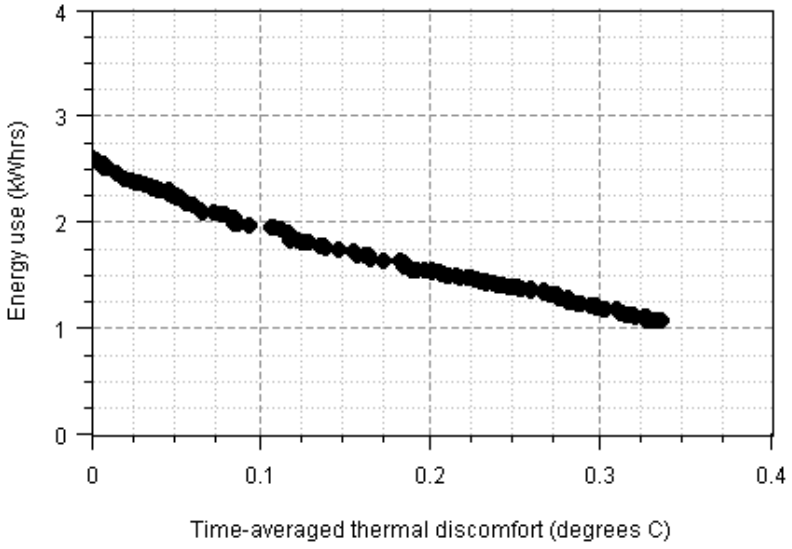

Figure B.44: Pareto front for the *SumAve2_9_13-18_dv5_ot* case shown as a trade-off between the two objectives.

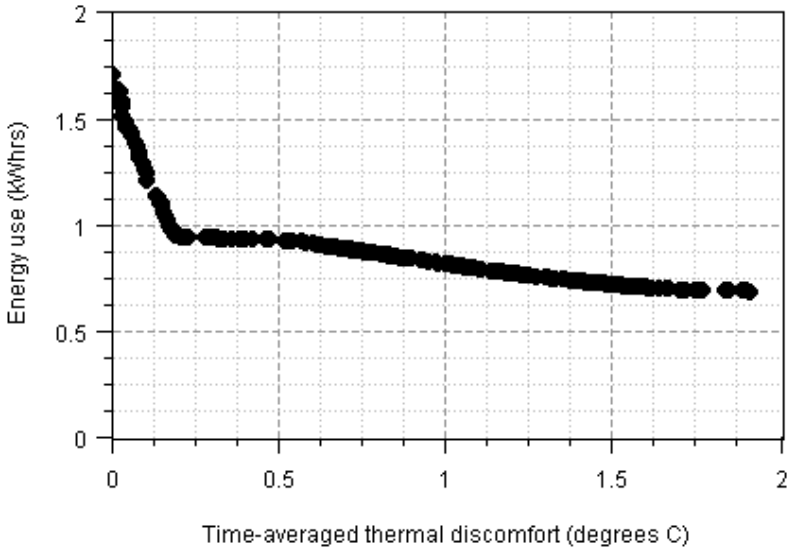

Fig. B.45: Pareto front for the *SumAve2_9_19-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort for the 13-18 case:

- Energy use: 2.599 kWh
- Radiant panel temperature: 9.976 (27.44 °C)
- Door discharge factor: 1.283 (0.123)
- Corridor temperature: 1.392 (18.31 °C)

Objective function and design variable values for the solution at 0 °C discomfort for the 19-24 case:

- Energy use: 1.706 kWh
- Radiant panel temperature: 9.553 (26.38 °C)
- Door discharge factor: 1.051 (0.104)
- Corridor temperature: 2.883 (19.51 °C)

Case: SumAve2_13_1-24_dv9_ot

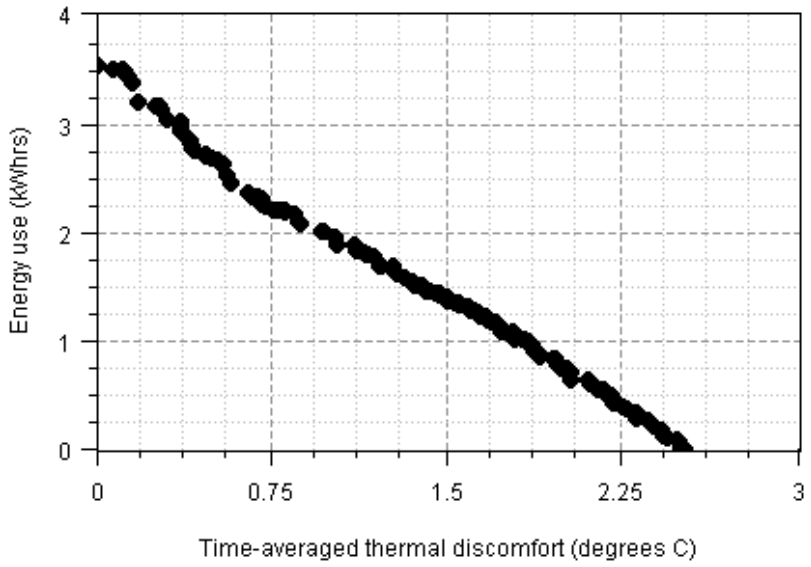

Figure B.46: Pareto front for the *SumAve2_13_1-24_dv9_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solutions at 0 °C discomfort and 0 kWh energy use:

- 0 °C thermal discomfort:
    - Energy use: 3.519 kWh
    - Window control opening area ratios: 10.938 / 1.000 / 10.921 (10.93 / 0.1 / 10.91)
    - Window control temperature set-points: 10.981 / 1.053 / 10.872 (27.98 / 27.98 / 28.00 °C)

    - Radiant panel temperature: 9.084 (25.21 °C)
    - Door discharge factor: 1.054 (0.104)
    - Corridor temperature: 9.252 (24.60 °C)
- 0 kWh energy use:
    - Thermal discomfort: 2.511 °C
    - Window control opening area ratios: 10.914 / 1.007 / 1.137 (10.91 / 0.11 / 0.25)
    - Window control temperature set-points: 10.836 / 10.881 / 10.953 (27.84 / 28.00 / 28.00 °C)
    - Radiant panel temperature: 9.515 (26.29 °C)
    - Door discharge factor: 2.146 (0.192)
    - Corridor temperature: 10.946 (25.96 °C)

B.4 Extreme winter conditions

This set of results represents cases run under extreme winter conditions, as described in section 6.3.6. Under these conditions it was assumed to be unlikely that the window would be open at all, so only the 5 design variable sample was simulated, with design variables and limits as follows:

- External wall thermal conductivity (0.1-1.5 W/mK)
- External wall thermal storage properties (500-2000 kg/m^3 & J/kgK)
- Radiant panel temperature (20-40 °C)
- Door discharge factor (0.1-0.9)
- Corridor air temperature (20-28 °C)

Given the relatively poor metamodel performance at the boundaries of the external wall design variables as already mentioned, and also the demonstrably more consistent results of optimising without these design variables, it was decided to run all optimisations for this climate case in terms of only the last three design variables, fixing the values of the external wall properties at their default values as given in Table 6.2.

Case: WinCol1_9_1-24_dv5_ot

Objective function and design variable values for the solution at 0 °C discomfort:

- Energy use: 4.005 kWh
- Radiant panel temperature: 3.004 (24.01 °C)
- Door discharge factor: 3.849 (0.328)

- Corridor temperature: 5.727 (23.78 °C)

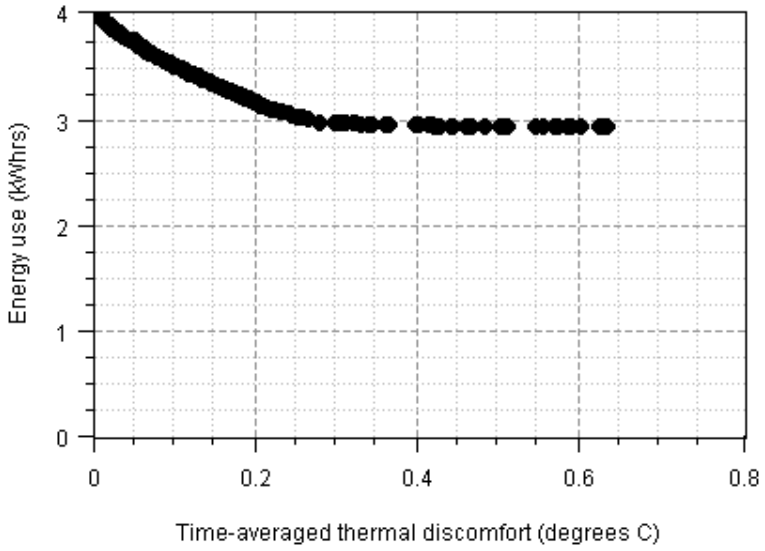

Figure B.47: Pareto front for the *WinCol1_9_1-24_dv5_ot* case shown as a trade-off between the two objectives.

Case: WinCol1_15_1-24_dv5_ot

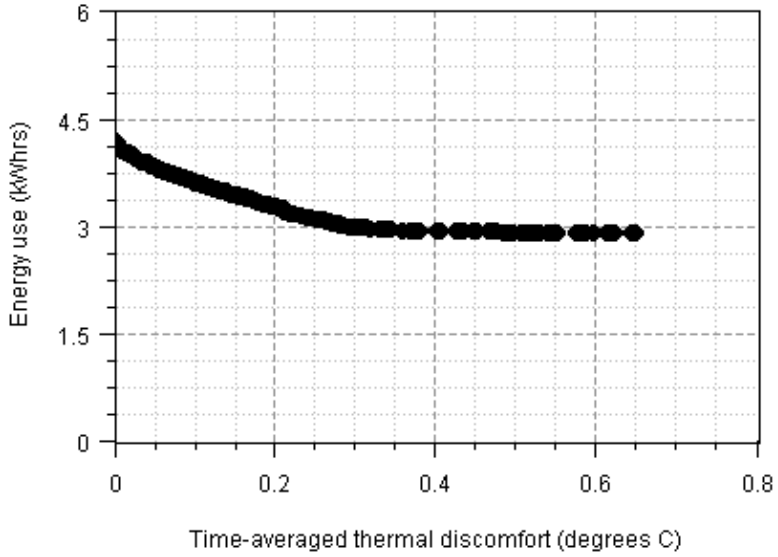

Figure B.48: Pareto front for the *WinCol1_15_1-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort:

- Energy use: 4.179 kWh
- Radiant panel temperature: 2.868 (23.74 °C)
- Door discharge factor: 4.268 (0.361)
- Corridor temperature: 6.110 (24.09 °C)

Case: WinCol1_15_1-24_dv5_pmvP

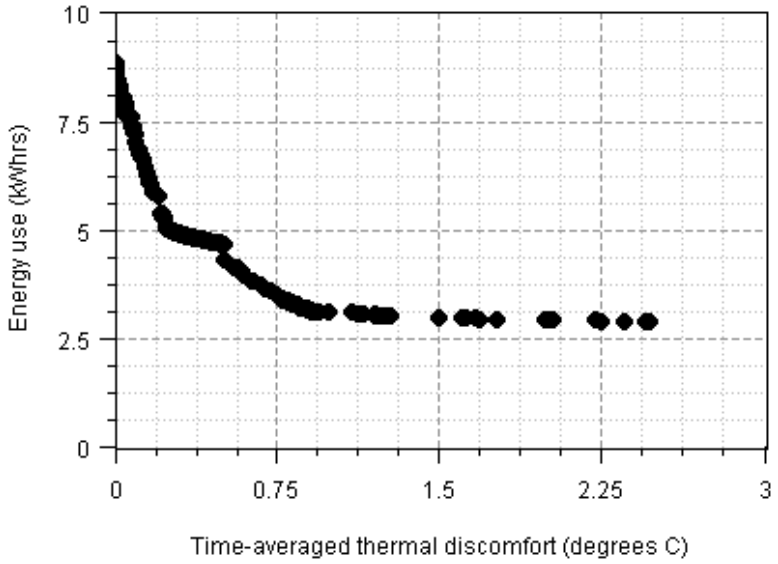

Figure B.49: Pareto front for the *WinCol1_15_1-24_dv5_pmvP* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort:

- Energy use: 8.808 kWh
- Radiant panel temperature: 1.735 (21.47 °C)
- Door discharge factor: 9.193 (0.755)
- Corridor temperature: 10.119 (27.30 °C)

Case: WinCol1_9_1-6,7-12,13-18,19-24_dv5_ot

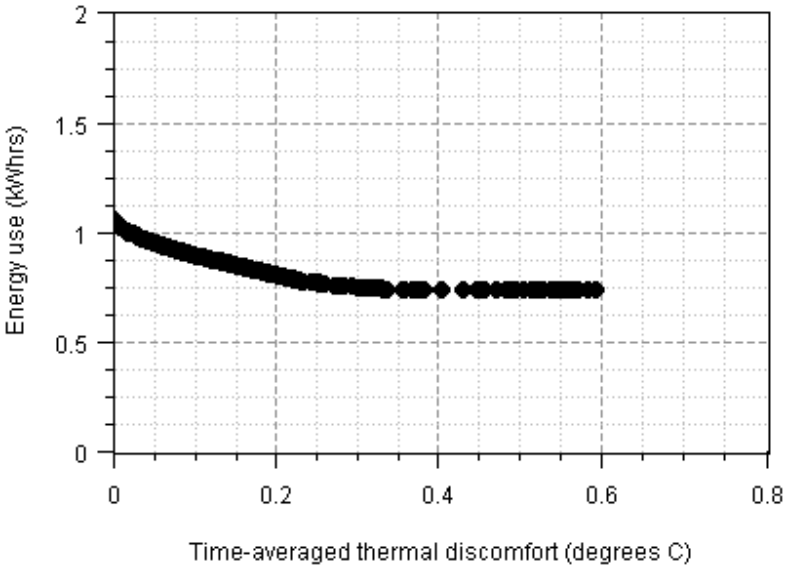

Figure B.50: Pareto front for the *WinCol1_9_1-6_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort for the 1-6 case:

- Energy use: 1.054 kWh
- Radiant panel temperature: 3.178 (24.36 °C)
- Door discharge factor: 4.422 (0.374)
- Corridor temperature: 5.549 (23.64 °C)

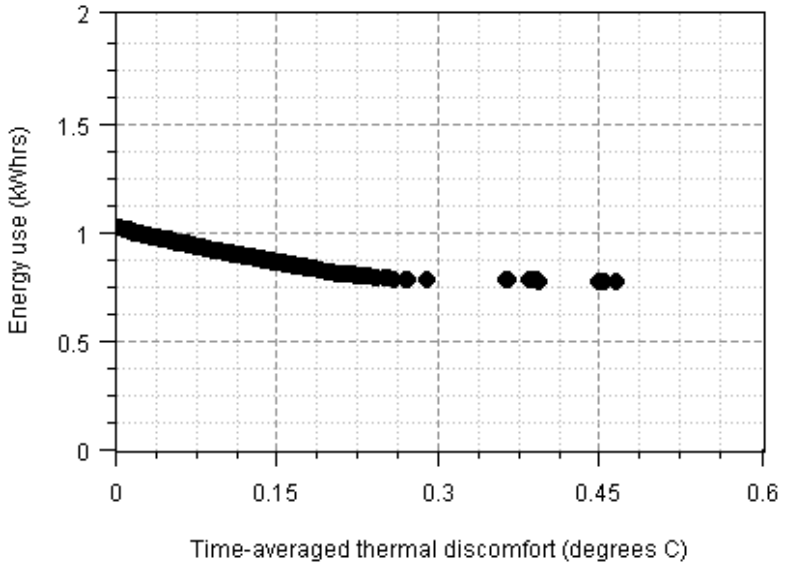

Figure B.51: Pareto front for the *WinCol1_9_7-12_dv5_ot* case shown as a trade-off between the two objectives.

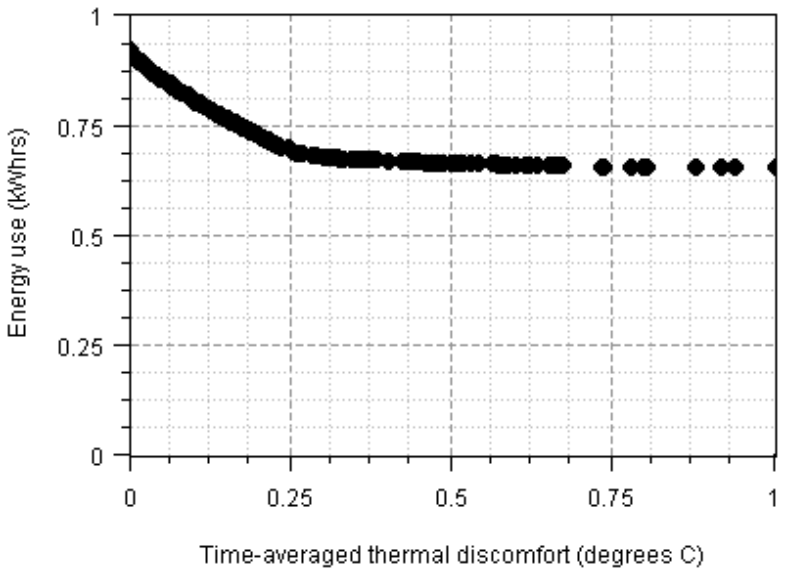

Figure B.52: Pareto front for the *WinCol1_9_13-18_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable results for the solution at 0 °C discomfort for the 9-12 case:

- Energy use: 1.022 kWh
- Radiant panel temperature: 2.881 (23.76 °C)
- Door discharge factor: 3.650 (0.312)
- Corridor temperature: 5.910 (23.93 °C)

Objective function and design variable results for the solution at 0 °C discomfort for the 13-18 case:

- Energy use: 0.916 kWh
- Radiant panel temperature: 2.844 (23.69 °C)
- Door discharge factor: 3.527 (0.302)
- Corridor temperature: 5.833 (23.87 °C)

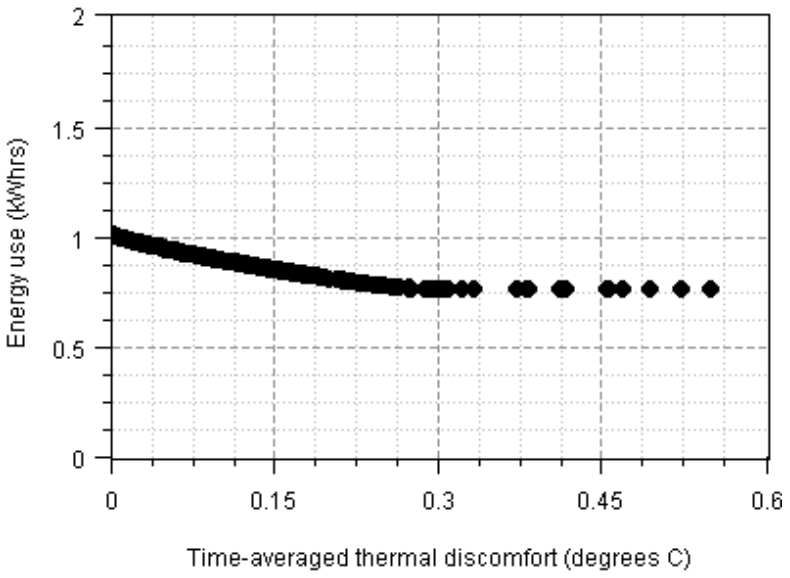

Figure B.53: Pareto front for the *WinCol1_9_19-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable results for the solution at 0 °C discomfort for the 19-24 case:

- Energy use: 1.009 kWh
- Radiant panel temperature: 3.000 (24.00 °C)
- Door discharge factor: 3.790 (0.323)
- Corridor temperature: 5.833 (23.78 °C)

B.5 Average Winter Conditions

B.5.1 Average winter conditions day 1

This set of cases was optimised using a sample run under average winter conditions, day 1, as described in section 6.3.6. As with extreme winter conditions, it was assumed that the window would not be open at all during this period so only one sample of 5 design variables was simulated. Design variables and their limits were:

- External wall thermal conductivity (0.1-1.5 W/mK)
- External wall thermal storage properties (500-2000 kg/m^3 & J/kgK)
- Radiant panel temperature (20-40 °C)
- Door discharge factor (0.1-0.9)
- Corridor air temperature (20-28 °C)

As with previous climate cases, values of the external wall properties design variables were fixed at default values as given in Table 6.2 in order to improve the quality and consistency of the results, so the optimisations were in terms of the last 3 design variables only.

Case: WinAve1_9_1-24_dv5_ot

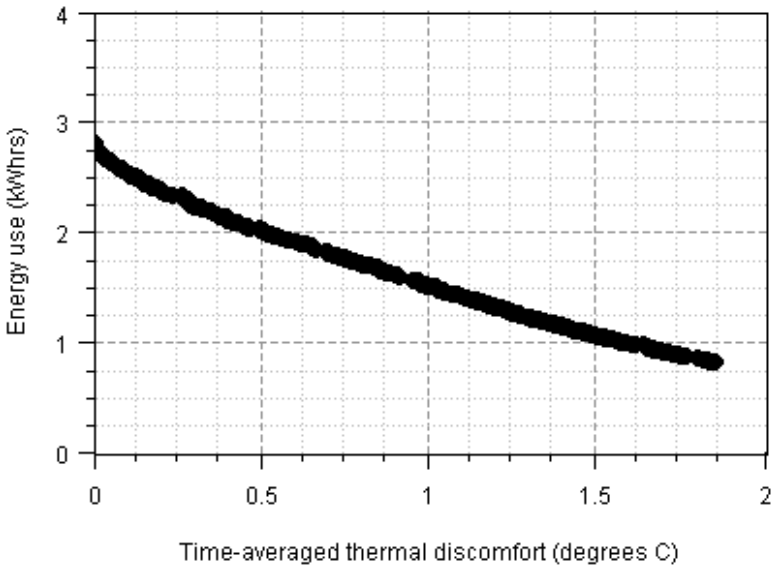

Figure B.54: Pareto front for the *WinAve1_9_1-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort:

- Energy use: 2.809 kWh
- Radiant panel temperature: 2.532 (23.06 °C)

- Door discharge factor: 3.622 (0.310)
- Corridor temperature: 5.899 (23.92 °C)

Case: WinAve1_17_1-24_dv5_ot

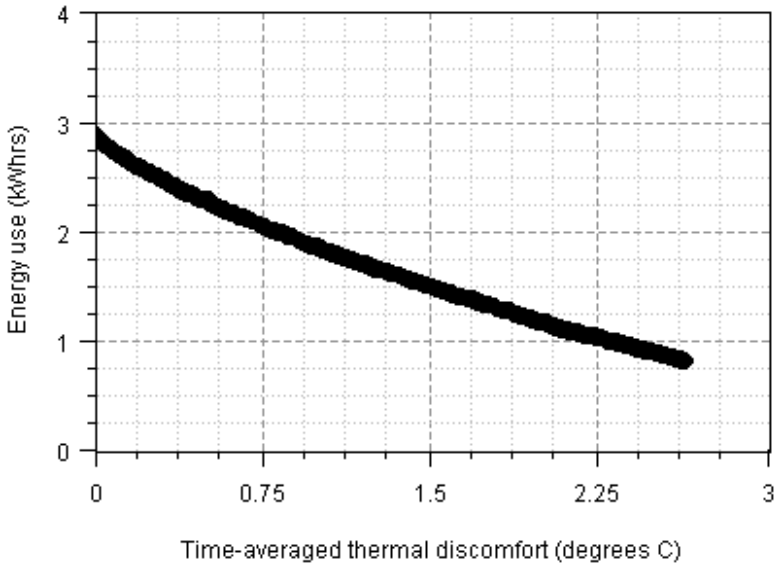

Figure B.55: Pareto front for the *WinAve1_17_1-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort:

- Energy use: 2.871 kWh
- Radiant panel temperature: 2.824 (23.65 °C)
- Door discharge factor: 4.252 (0.360)
- Corridor temperature: 5.944 (23.96 °C)

Case: WinAve1_9_1-6,7-12,13-18,19-24_dv5_ot

Objective function and design variable values for the solution at 0 °C discomfort for the 1-6 case:

- Energy use: 0.722 kWh
- Radiant panel temperature: 2.501 (23.00 °C)
- Door discharge factor: 3.881 (0.330)
- Corridor temperature: 6.070 (24.06 °C)

Objective function and design variable values for the solution near 0 °C discomfort for the 7-12 case:

- Thermal discomfort: 0.007 °C

- Energy use: 0.697 kWh
- Radiant panel temperature: 2.730 (23.46 °C)
- Door discharge factor: 3.722 (0.318)
- Corridor temperature: 5.508 (23.61 °C)

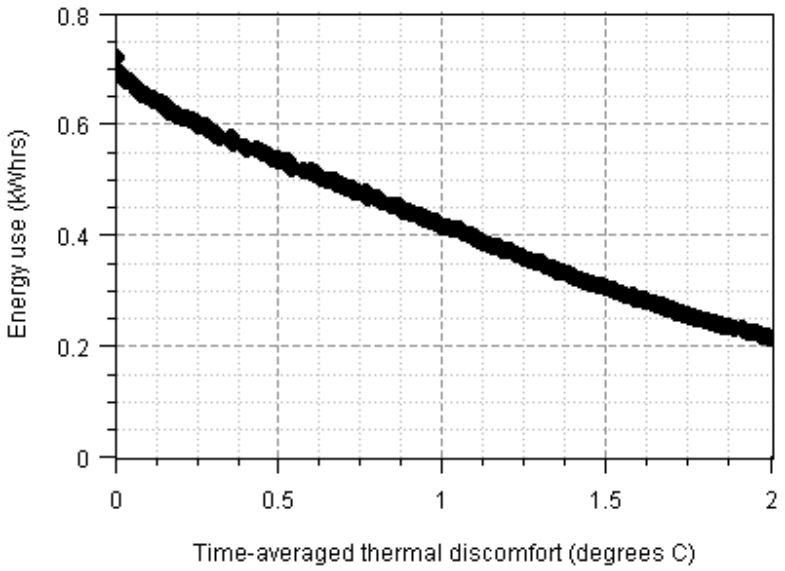

Figure B.56: Pareto front for the *WinAve1_9_1-6_dv5_ot* case shown as a trade-off between the two objectives.

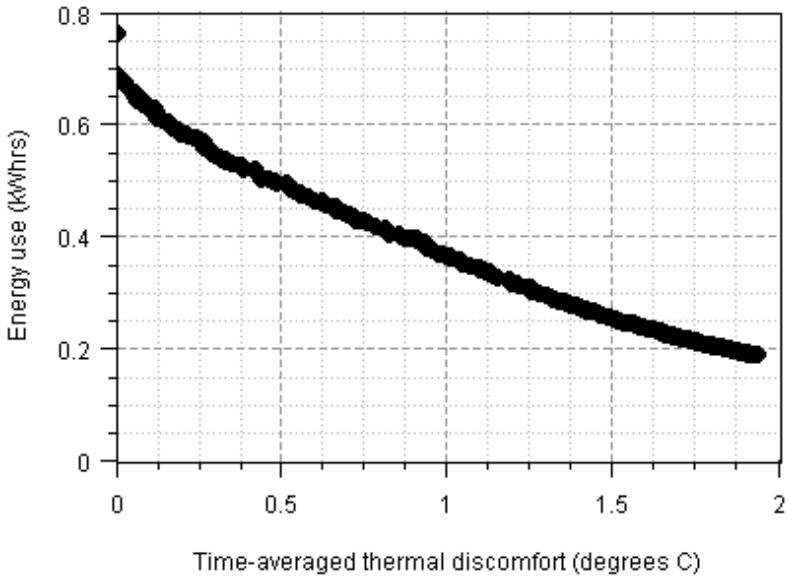

Figure B.57: Pareto front for the *WinAve1_9_7-12_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort for the 13-18 case:

- Energy use: 0.763 kWh

- Radiant panel temperature: 2.159 (22.32 °C)
- Door discharge factor: 3.100 (0.268)
- Corridor temperature: 6.406 (24.32 °C)

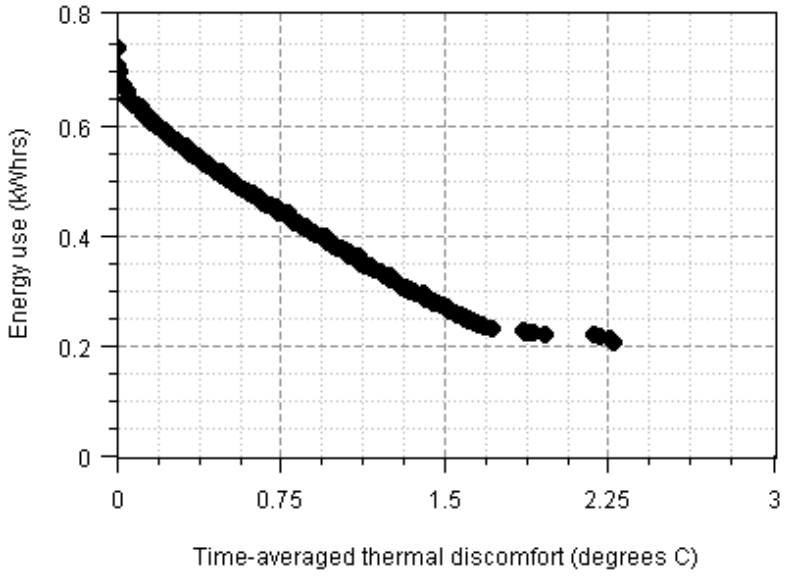

Figure B.58: Pareto front for the *WinAve1_9_13-18_dv5_ot* case shown as a trade-off between the two objectives.

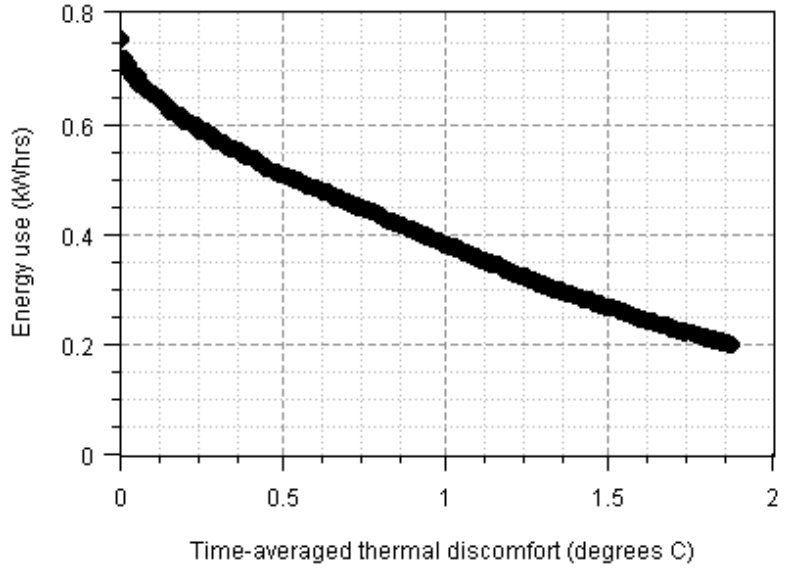

Figure B.59: Pareto front for the *WinAve1_9_19-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort for the 19-24 case:

- Energy use: 0.736 kWh
- Radiant panel temperature: 2.260 (22.52 °C)

- Door discharge factor: 3.470 (0.298)
- Corridor temperature: 6.316 (24.25 °C)

B.5.2 Average winter conditions day 2

This set of cases was optimised using a sample run with average winter conditions, day 2, as described in section 6.3.6. It was assumed that the window would not be open during the winter season, so only one sample of 5 design variables was simulated. Design variables and their limits were:

- External wall thermal conductivity (0.1-1.5 W/mK)
- External wall thermal storage properties (500-2000 kg/m^3 & J/kgK)
- Radiant panel temperature (20-40 °C)
- Door discharge factor (0.1-0.9)
- Corridor air temperature (20-28 °C)

As with previous cases all optimisations were performed in terms of only the last 3 design variables, the values of external wall properties being fixed at their default values as given in Table 6.2.

Case: WinAve2_9_1-24_dv5_ot

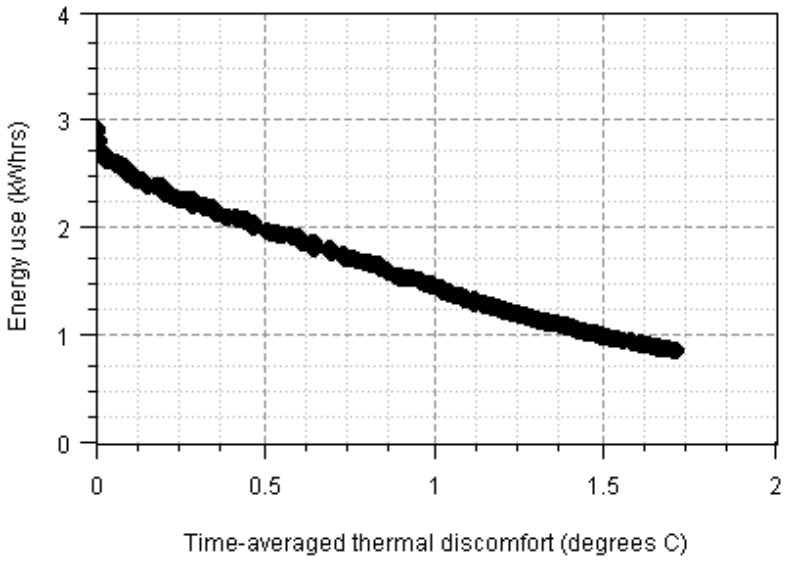

Figure B.60: Pareto front for the *WinAve2_9_1-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort:

- Energy use: 2.925 kWh

- Radiant panel temperature: 2.297 (22.59 °C)
- Door discharge factor: 3.534 (0.303)
- Corridor temperature: 6.243 (24.19 °C)

Case: WinAve2_17_1-24_dv5_ot

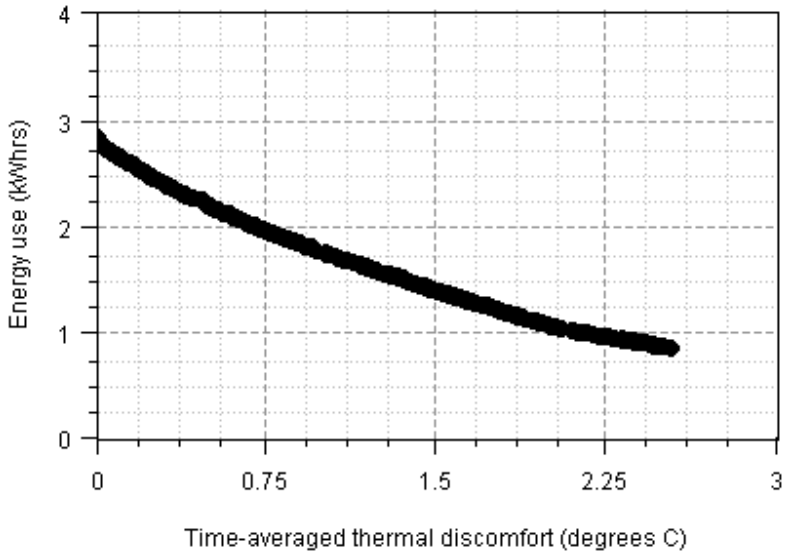

Figure B.61: Pareto front for the *WinAve2_17_1-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort:

- Energy use: 2.814 kWh
- Radiant panel temperature: 2.964 (23.93 °C)
- Door discharge factor: 4.333 (0.367)
- Corridor temperature: 5.654 (23.72 °C)

Case: WinAve2_9_1-6,7-12,13-18,19-24_dv5_ot

Objective function and design variable values for the solution at 0 °C discomfort for the 1-6 case:

- Energy use: 0.794 kWh
- Radiant panel temperature: 2.389 (22.78 °C)
- Door discharge factor: 4.365 (0.369)
- Corridor temperature: 6.603 (24.48 °C)

Objective function and design variable values for the solution at 0 °C discomfort for the 7-12 case:

- Energy use: 0.876 kWh
- Radiant panel temperature: 1.788 (21.58 °C)
- Door discharge factor: 2.714 (0.237)
- Corridor temperature: 6.742 (24.59 °C)

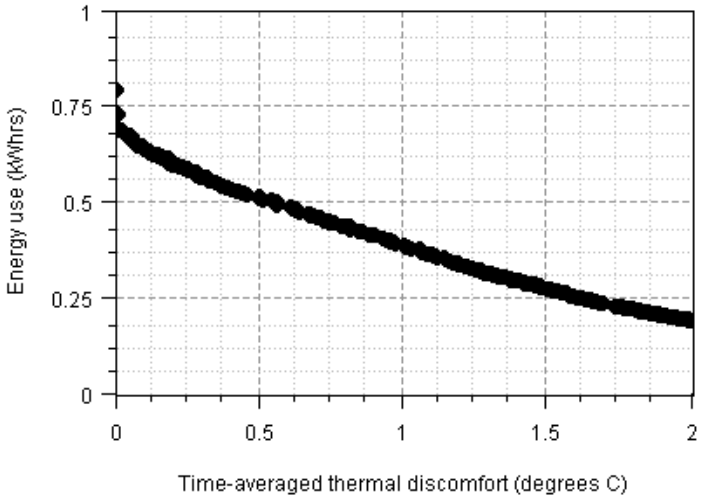

Figure B.62: Pareto front for the *WinAve2_9_1-6_dv5_ot* case shown as a trade-off between the two objectives.

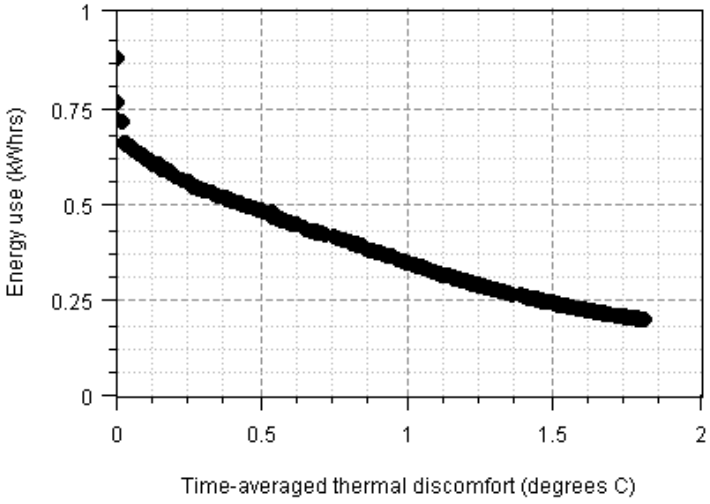

Figure B.63: Pareto front for the *WinAve2_9_7-12_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution closest to 0 °C discomfort for the 13-18 case:

- Thermal discomfort: 0.011 °C
- Energy use: 0.634 kWh
- Radiant panel temperature: 2.499 (23.00 °C)

- Door discharge factor: 4.040 (0.343)
- Corridor temperature: 5.811 (23.85 °C)

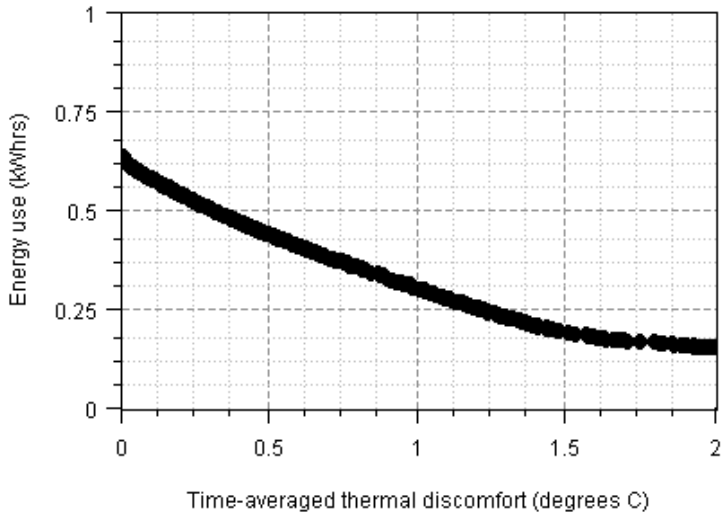

Figure B.64: Pareto front for the *WinAve2_9_13-18_dv5_ot* case shown as a trade-off between the two objectives.

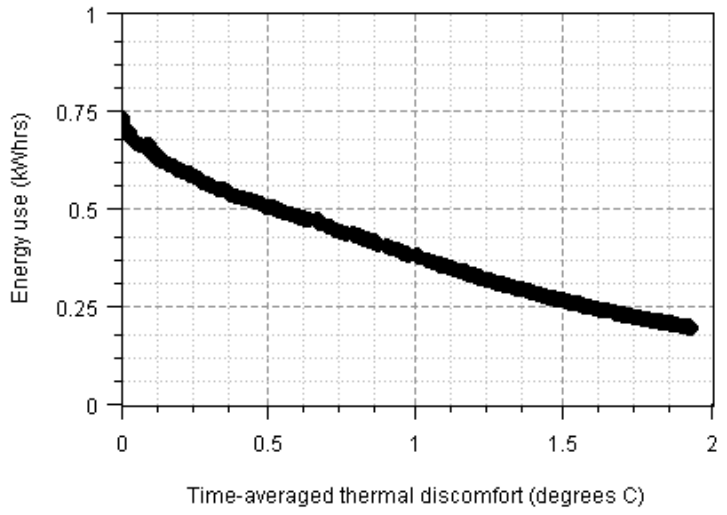

Figure B.65: Pareto front for the *WinAve2_9_19-24_dv5_ot* case shown as a trade-off between the two objectives.

Objective function and design variable values for the solution at 0 °C discomfort for the 19-24 case:

- Energy use: 0.731 kWh
- Radiant panel temperature: 2.426 (22.85 °C)
- Door discharge factor: 3.569 (0.306)
- Corridor temperature: 6.066 (24.05 °C)

References

Aasem, O.E. 1993. *Practical simulation of buildings and air-conditioning systems in the transient domain*. PhD thesis, University of Strathclyde, Glasgow.

Adamu, Z.A.; Price, A.D.F. and Cook, M.J. 2012. Performance evaluation of natural ventilation strategies for hospital wards - A case study of Great Ormond Street Hospital. *Building and Environment*; 56; 211-222.

Alcalá, R.; Casillas, J.; Cordón, O; González, A. and Herrera, F. 2005. A genetic rule weighting and selection process for fuzzy control of heating, ventilating and air conditioning systems. *Engineering Applications of Artificial Intelligence*; 18; 279-296.

Al-Homoud, M.S. 1994. *Design optimization of energy conserving building envelopes.* PhD thesis, Texas A&M University.

Al-Homoud, M.S. 1997a. Optimum thermal design of office buildings. *International Journal of Energy Research*; 21; pp. 941-957.

Al-Homoud, M.S. 1997b. Optimum Thermal Design of Air-Conditioned Residential Buildings. *Building and Environment*; 32(3); pp. 203-210.

Al-Homoud, M.S. 2001. Computer-aided building energy analysis techniques. *Building and Environment*; 36; pp. 421-433.

Al-Homoud, M.S. 2005. A Systematic Approach for the Thermal Design Optimization of Building Envelopes. *Journal of Building Physics*; 29(2); pp. 95-119.

Al-Homoud, M.S. 2009. Envelope Thermal Design Optimization of Buildings with Intermittent Occupancy. *Journal of Building Physics*; 33(1); pp. 65-82.

Al-Othmani, M.; Ghaddar, N. and Ghali, K. 2008. A multi-segmented human bioheat model for transient and asymmetric radiative environments. *International Journal of Heat and Mass Transfer*; 51; 5522-5533.

Al-Sanea, S.A. and Zedan, M.F. 2011. Improving thermal performance of building walls by optimizing insulation layer distribution and thickness for same thermal mass. *Applied Energy*; 88; pp. 3113–3124.

Amissah, P.K. 2005. *Indoor air quality – combining air humidity with construction moisture.* PhD thesis, University of Strathclyde, Glasgow.

Anastaselos, D.; Oxizidis, S. and Papadopoulos, A.M. 2011. Energy, environmental and economic optimization of thermal insulation solutions by means of an integrated decision support system. *Energy and Buildings*; 43; pp. 686–694.

Arens, E. and Zhang, H. 2006. The skin's role in human thermoregulation and comfort. *Chapter 16 in: Thermal and Moisture Transport in Fibrous Materials*; eds. N. Pan and P. Gibson; Woodhead Publishing Ltd., pp 560-602.

Arens, E.; Turner, S.; Zhang, H. and Paliaga, G. 2009. Moving air for comfort. *ASHRAE Journal*; 51; 8-18.

Arens, E.; Zhang, H. and Huizenga, C. 2006a. Partial- and whole-body thermal sensation and comfort: Part I: Uniform environmental conditions. *Journal of Thermal Biology*; 31; 53-59.

Arens, E.; Zhang, H. and Huizenga, C. 2006b. Partial- and whole-body thermal sensation and comfort: Part II: Non-uniform environmental conditions. *Journal of Thermal Biology*; 31; 60-66.

Ari, S.; Wilcoxen, P.; Khalifa, H. E.; Dannenhoffer, J. F. and Isik, C. 2008. A practical approach to individual thermal comfort and energy optimization problem. In: *Proceedings of the Annual Meeting of the North American Fuzzy Information Processing Society* (New York). Pp. 1-6.

Arumi, F. 1977. Day lighting as a factor in optimizing energy performance of buildings. *Energy and Buildings*; 1(2); pp. 175-182.

ASHRAE. 2004. *ASHRAE Standard 55: thermal environmental conditions for human occupancy*. American Society of Heating, Refrigerating and Air-conditioning Engineers; Atlanta.

ASHRAE. 2009. *2009 ASHRAE Handbook – Fundamentals.* American Society of Heating, Refrigerating and Air-conditioning Engineers; Atlanta.

Autodesk. 2014. *Autodesk Ecotect Analysis*. Website accessed 14th April 2013, available on the World Wide Web: http://usa.autodesk.com/ecotect-analysis/

Avolio, A. P. 1980. Multi-branched model of the human arterial system. *Medical and Biological Engineering and Computing*; 18; 709-718.

Azizpour, F.; Moghimi, S.; Salleh, E.; Mat, S.; Lim, C.H. and Sopian, K. 2013. Thermal comfort assessment of large-scale hospitals in tropical climates: A case study of University Kebangsaan Malaysia Medical Centre (UKMMC). *Energy and Buildings*; 64; 317–322.

Backus, G. and Gilbert, F. 1968. The resolving power of gross earth data. *Geophysical Journal of the Royal Astronomical Society*; 16; pp. 169-205.

Barricelli, N. A. 1954. Esempi numerici di processi di evoluzione. *Methodos*; 45-68.

Bates, S.J.; Sienz, S.J.; Toropov, V.V. 2004. Formulation of the optimal Latin hypercube design of experiments using a permutation genetic algorithm. In: *Proceedings of the 45th AIAA Structures, Structural Dynamics & Materials Conference*, Palm Springs, USA.

Beausoleil-Morrison, I. 2000. *The adaptive coupling of heat and air flow modeling within dynamic whole-building simulation*. PhD Thesis, University of Strathclyde, Glasgow.

Bellman, R. E. 2003 (reprint of 1957 original). *Dynamic Programming*. Dover Publications Inc.

Ben-Nakhi, A. E. and Mahmoud, M. A. 2002. Energy conservation in buildings through efficient A/C control using neural networks. *Applied Energy*; 73; 5-23.

Berndt, E, K.; Hall, B. H.; Hall, R. E. and Hausman, J. A. 1974. Estimation and inference in nonlinear structural models. *Annals of Economic and Social Measurement*; 3/4; 653-665.

Bichiou, Y. and Krarti, M. 2011. Optimization of envelope and HVAC systems selection for residential buildings. *Energy and Buildings*; 43; pp. 3373–3382.

Birge, J. 1994. Citation for George Dantzig prize. In: *Optima*; 44; 4-5.

Blum, J. R. 1954. Multidimensional stochastic approximation methods. *Annals of Mathematical Statistics*; 25; 737-744.

Bos, L.P. and Salkauskas, K. 1989. Moving Least-Squares Are Backus-Gilbert Optimal. *Journal of Approximation Theory*; 59; pp. 267-275.

Bouchlaghem, N. 2000. Optimising the design of building envelopes for thermal performance. *Automation in Construction*; 10; pp. 101–112.

Bouchlaghem, N.M. and Letherman, K.M. 1990. Numerical Optimization Applied to the Thermal Design of Buildings. *Building and Environment*; 25(2); pp. 117-124.

BRE. 2014a. The Government's Standard Assessment Procedure for Energy Rating of Dwellings. Website accessed 10/10/14. Available online: http://www.bre.co.uk/filelibrary/SAP/2012/SAP-2012_9-92.pdf

BRE. 2014b. Energy in the NHS. Website accessed 23/9/2014. Available online: http://www.bre.co.uk/page.jsp?id=2627

BRE. 2014c. Carbon Footprint. Website accessed 1/10/2014. Available online: http://www.bre.co.uk/page.jsp?id=2637

Bremermann, H. J. 1962. Optimization through evolution and recombination. In: *Self-Organizing Systems; proceedings of conference on self-organizing systems* (Chicago). Spartan Book, Washington D. C. pp. 93-106.

Broyden, C. G. 1965. A class of methods for solving nonlinear simultaneous equations. *Mathematics of Computation*; 19; 577-593.

BSi. 2005(1). *BS EN ISO 15927- 4 Hygrothermal performance of buildings. Calculation and presentation of climatic data. Hourly data for assessing the annual energy use for heating and cooling.* British Standards Institution; Brussels.

BSi. 2005(2). *BS EN ISO 15927-5:2004+A1:2011 Hygrothermal performance of buildings. Calculation and presentation of climatic data. Data for design heat load for space heating.* British Standards Institution; Brussels.

BSi. 2006. *BS EN ISO 7730:2005 Ergonomics of the thermal environment — Analytical determination and interpretation of thermal comfort using calculation of the PMV and PPD indices and local thermal comfort criteria*. British Standards Institution; Brussels.

BSi. 2008. *BS EN 15251:2007 Indoor environmental input parameters for design and assessment of energy performance of buildings addressing indoor air quality, thermal environment, lighting and acoustics*. British Standards Institution; Brussels.

Chantrelle, F.P.; Lahmidi, H.; Keilholz, W.; El Mankibi, M. and Michel, P. 2011. Development of a multicriteria tool for optimizing the renovation of buildings. *Applied Energy*; 88; 1386–1394.

Chen, T. Y. 2001. Real-time predictive supervisory operation of building thermal systems with thermal mass. *Energy and Buildings*; 33; pp. 141-150.

Chen, Q. and Xu, W. 1998. A Zero-Equation Turbulence Model for Indoor Airflow Simulation. *Energy and Buildings*; 28; pp. 137-144.

CIBSE. 2006. *Environmental Design CIBSE Guide A*. Chartered Institution of Building Services Engineers; London.

Cioppa, A. D.; de Stefano, C. and Marcelli, A. 2004. On the role of population size and niche radius in fitness sharing. *IEEE Transactions on Evolutionary Conputation*; 8; pp. 580-592.

Clarke, J.A. 1977. *Environmental Systems Performance*. PhD Thesis, University of Strathclyde, Glasgow.

Clarke, J.A. 2001. *Energy simulation in building design*. Butterworth-Heinemann, Oxford.

Coley, D.A. 1999. *An introduction to Genetic Algorithms for Scientists and Engineers.* World Scientific Publishing.

Cortes, C. and Vapnik, V. 1995. Support-vector networks. *Machine Learning*; 20(3); pp. 273-297.

Cowie, A.C; Noakes, C.J; Sleigh, P.A. and Toropov, V.V. 2013. Meta-model based Optimization of Building Thermal Performance incorporating Local Comfort Analysis. In: *Proceedings of the Fourteenth International Conference on Civil, Structural and Environmental Engineering Computing, Cagliari, Sardinia*; 3 – 6 September 2013.

Cowie, A.C; Noakes, C.J.; Sleigh, P.A. and Toropov, V.V. 2015. A meta-model based method for thermal comfort and energy use optimisation at a local resolution. In: *Proceedings of the Fourtheenth International Conference of the International Building Performance Simulation Association, Hyderabad, India*; 7 – 9 December 2015.

Cressie, N. 1990. The origins of kriging. *Mathematical Geology*; 22(3); pp. 239-252.

Csébfalvi, A. 2013. Hybrid Metaheuristic Methods in Truss Optimiation: A Review. *Computational Technology Reviews*; 8; 63-92.

D'cruz, N.A. and Radford, A.D. 1987. A Multicriteria Model for Building Performance and Design. *Building and Environment*; 22(3); pp. 167-179.

D'cruz, N.; Radford, A.D. and Gero, J.S. 1983 A pareto optimization problem formulation for building performance and design. *Engineering Optimization*; 7(1); pp. 17-33.

Dantzig, G. B. 1951a. Maximization of a linear function of variables subject to linear inequalities. In: *Activity Analysis of Production and Allocation* (proceedings of Conference on Linear Programming, Chicago, Illinois, 1949). Wiley, New York. Pp 339-347.

Dantzig, G. B. 1951b. Application of the simplex method to a transportation problem. In: *Activity Analysis of Production and Allocation* (proceedings of Conference on Linear Programming, Chicago, Illinois, 1949). Wiley, New York. Pp 359-373.

Davidon, W. C. 1959. Variable metric method for minimization. Rep. ANL-5990 Rev, Argonne National Laboratories, Argonne, Illinois.

de Dear, R. J. and Brager, G. S. 2002. Thermal comfort in naturally ventilated buildings: revisions to ASHRAE Standard 55. *Energy and Buildings*; 34; 549-561.

de Giuli, V.; Zecchin, R.; Salmaso, L.; Corain, L. and de Carli, M. 2013. Measured and perceived indoor environmental quality: Padua Hospital case study. *Building and Environment*; 59; pp. 211-226.

Deb, C. and Ramachandraiah, A. 2010. Evaluation of thermal comfort in a rail terminal location in India. *Building and Environment*; 45; 2571-2580.

Deb, K.; Pratap, A.; Agarwal, S. and Meyarivan, T. 2002. A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. *IEEE transactions on evolutionary computation*; 6(2); pp. 182-197.

Dennis, J. E. and Moree, J. J. 1977. Quasi-newton methods, motivation and theory. *SIAM Review*; 19; 46-89.

Djuric, N.; Novakovic, V.; Holst, J. and Mitrovic, Z. 2007. Optimization of energy consumption in buildings with hydronic heating systems considering thermal comfort by use of computer-based tools. *Energy and Buildings*; 39; pp. 471-477.

DoH. 1998. *Health Technical Memorandum 55 – 'Windows'*. The Stationery Office; London.

DoH. 2006. *Health technical memorandum 07e02: EnCO_2de-, making energy work in healthcare*. The Stationery Office; London.

DoH. 2007. *Heating and ventilation systems. Health Technical Memorandum 03-01: Specialised ventilation for healthcare premises. Part A: Design and validation.* The Stationary Office; London.

Dorigo, M. and Caro, G. D. 1999. Ant algorithms for discrete optimization. *Artificial Life*; 5; pp. 137-172.

Dorigo, M. and Gambardella, L. M. 1997. Ant colony system: a cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computing*; 1; pp. 53-66.

Dorigo, M.; Maniezzo, V. and Colorni, A. 1996. Ant system: optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man and Cybernetics – Part B: Cybernetics*; 26; pp. 29-41.

Eberhart, R. and Kennedy, J. 1995. A new optimizer using particle swarm theory. In: *Proceedings of the Sixth International Symposium on Micro Machine and Human Science* (Nagoya). IEEE Service Center, Piscataway, New Jersey; pp. 39-43.

EDSL. 2012. *Tas: A complete solution for the thermal simulation of new or existing buildings*. Website accessed 21st April 2013, available on the World Wide Web: http://www.edsl.net/main/

EIA. 2008. *Table E1A. Major Fuel Consumption (Btu) by End Use for All Buildings, 2003*. Website accessed 2/10/2014. Available online: http://www.eia.gov/consumption/commercial/data/2003/pdf/e1a-e11a.pdf.

EIA. 2013. *Heating and cooling no longer majority of U.S. home energy use*. Website accessed 2/10/2014. Available online: http://www.eia.gov/todayinenergy/detail.cfm?id=10271&src=%E2%80%B9%20Consumption%20%20%20%20%20%20Residential%20Energy%20Consumption%20Survey%20%28RECS%29-b1

EIA. 2014. *How much energy is consumed in residential and commercial buildings in the United States?* Website accessed 2/10/2014. Available online: http://www.eia.gov/tools/faqs/faq.cfm?id=86&t=1

Eisenhower, B.; O'Neill, Z.; Narayanan, S.; Fonoberov, V.A. and Mezic, I. 2012. A methodology for meta-model based optimization in building energy models. *Energy and Buildings*; 47; 292–301.

Fanger, P. O. 1972. *Thermal Comfort: Analysis and Applications in Environmental Engineering*. McGraw-Hill Book Company, New York.

Fanger, P. O. 1989. New comfort equation for indoor air quality. *ASHRAE Journal*; 31(10).

Fanger, P. O. 2000. Indoor air quality in the 21st century: search for excellence. *Indoor Air*; 10; 68-73.

Fanger, P. O. and Toftum, J. 2002. Extension of the PMV model to non-air-conditioned buildings in warm climates. *Energy and Buildings*; 34; 533-536.

Fanger, P.O.; Ipsen, B. M.; Langkilde, G.; Olesen, B. W.; Christensen, N. K. and Tanabe, S. 1985. Comfort limits for asymmetric thermal radiation. *Energy and Buildings*; 8; 225-236.

Fanger, P.O.; Melikov, A. K.; Hanzawa, H. and Ring, J. 1988. Air turbulence and sensation of draught. *Energy and Buildings*; 12; 21-39.

Farwig, R. 1986. Multivariate interpolation of arbitrarily spaced data by moving least squares methods. *Journal of Computational and Applied Mathematics*; 16; pp. 79-93.

Farzaneh, Y. and Tootoonchi, A. A. 2008. Controlling automobile thermal comfort using optimized fuzzy controller. *Applied Thermal Engineering*; 28; 1906-1917.

Fengzhi, L. and Yi, L. 2005. Effect of clothing material on thermal response of the human body. *Modelling and Simulation in Materials Science and Engineering*; 13; 809–827.

Feustel, H.E. and Dieris, J. 1992. A Survey of Airflow Models for Multizone Structures. *Energy and Buildings*; 18; 79-100.

Fiala, D.; Havenith, G.; Bröde, P.; Kampmann, B. and Jendritzky, G. 2011. UTCI-Fiala multi-node model of human heat transfer and temperature regulation. *International Journal of Biometeorology*; not in print, accessed online 25th April 2011, available on the World Wide Web: http://www.springerlink.com/content/buk216088ml12233/

Fiala, D.; Lomas, K. J. and Stohrer, M. 2001. Computer predictions of human thermoregulatory and temperature responses to a wide range of environment conditions. *International Journal of Biometeorology*; 45; 143–59.

Fletcher, R and Powell, M. J. D. 1963. A rapidly convergent descent method for minimization. *Computer Journal*; 6; 163-168.

Fletcher, R. 1982. A model algorithm for composite nondifferentiable optimization problems. In: *Mathematical Programming Study.* pp. 67-76.

Fletcher, R. and Reeves, C. M. 1964. Function minimization by conjugate gradients. *Computer Journal*; 7; 149-154.

Fogel, D. B.; Fogel, G. B. and Ohkura, K. 2001. Multiple-vector self-adaptation in evolutionary algorithms. *Biosystems*; 61; 155-162.

Fogel, L. J.; Owens, A. J. and Walsh, M. J. 1966. *Artificial intelligence through simulated evolution.* John Wiley, New York.

Fong, K. F.; Hanby, V. I. and Chow, T. T. 2006. HVAC system optimization for energy management by evolutionary programming. *Energy and Buildings*; 38; pp. 220-231.

Ford, L. R. and Fulkerson, D. R. 1956. Maximal flow through a network. *Canadian Journal of Mathematics*; 8; 399-404.

Freire, R. Z.; Oliveira, G. H. C. and Mendes, N. 2008. Predictive controllers for thermal comfort optimization and energy savings. *Energy and Buildings*; 40; 1353-1365.

Fu, G. 1995. A transient, 3-D mathematical thermal model for the clothed human. *PhD thesis*; Department of Mechanical Engineering, Kansas State University, Manhattan.

Gagge, A. P.; Burton, A. C. and Bazett, H. C. 1941. A practical system of units for the description of the heat exchange of man with his environment. *Science*; 94; 428-430.

Gagge, A. P.; Fobelets, A. P. and Berglund, A. G. 1986. A standard predictive index of human response to the thermal environment. *ASHRAE Transactions*; 92; 709-731.

Gagge, A. P.; Hardy, J. D. and Rapp, G. M. 1965(1). Exploratory study of comfort for high temperature sources of radiant heat. *ASHRAE Transactions*; 71; 19-26.

Gagge, A. P.; Stolwijk, J. A. J. and Hardy, J. D. 1965(2). A novel approach to measurement of man's heat exchange with a complex radiant environment. *Aerospace Medical*; 36; 431-435.

Gagge, A. P.; Stolwijk, J. A. J. and Nishi, Y. 1971. An effective temperature scale based on a simple model of human physiological regulatory response. *ASHRAE Transactions*; 77; 247-262.

Gagge, A. P.; Stolwijk, J. A. J. and Saltin, B. 1969. Comfort and thermal sensations and associated physiological responses during exercise at various ambient temperatures. *Environmental Research*; 3; 209-229.

Gagge, A. P.; Stolwijk, J.A. J. and Hardy, J. D. 1967. Comfort and thermal sensations and associated physiological responses at various ambient temperatures. *Environmental Research*; 1;1-20.

Garg, A.; Chaffin, D. B. and Herrin, G. D. 1978. Prediction of metabolic rates for manual materials handling jobs. *American Industrial Hygiene Association Journal*; 39; 661-674.

General American Transportation Corporation (GATC). 1970. *Computer program for analysis of energy utilization in postal facilities.* GATX, USA.

Gengembre, E.; Ladevie, B.; Fudym, O. and Thuillier, A. 2013. A Kriging constrained efficient global optimization approach applied to low-energy building design problems. *Inverse Problems in Science and Engineering*; 20(7); pp. 1101-1114.

Gero, J.S. 1980. Computer-aided design by optimization in architecture. *Design Studies*; 1(4); pp. 227-230.

Gero, J.S.; D'cruz, N. and Radford, A.D. 1983. Energy in Context: A Multicriteria Model for Building Design. *Building and Environment*; 18(3); pp. 99-107.

Gilkeson, C.A.; Toropov, V.V.; Thompson, H.M.; Wilson, M.C.T.; Foxley, N.A. and Gaskell, P.H. 2013. Multi-objective aerodynamic shape optimization of small livestock trailers. *Engineering Optimization*; 45; pp. 1309-1330.

Gill, P. E. and Wong, E. 2010. *Sequential quadratic programming methods*; Technical report NA-10-03. UC San Diego, Department of Mathematics.

Google. 2014. *Google Maps*. Available online: https://maps.google.co.uk

Gov.uk. 2014a. *Improving the energy efficiency of buildings and using planning to protect the environment*. Website accessed 15/10/2014. Available online: https://www.gov.uk/government/policies/improving-the-energy-efficiency-of-buildings-and-using-planning-to-protect-the-environment/supporting-pages/energy-performance-of-buildings

Gov.uk. 2014b. Reducing the UK's greenhouse gas emissions by 80% by 2050. Website accessed 2/10/2014. Available online: https://www.gov.uk/government/policies/reducing-the-uk-s-greenhouse-gas-emissions-by-80-by-2050

Gupta, C.L. 1970. A Systematic Approach to Thermal Design. *Building Science*; 5; pp. 165- 173.

Gustafsson, S. 1998. Mixed integer linear programming and building retrofits. *Energy and Buildings*; 28; pp. 191-196.

Gustafsson, S. 2000. Optimization of insulation measures on existing buildings. *Energy and Buildings*; 33; pp. 49-55.

Gustafsson, S. 2001. Optimal fenestration refits by use of MILP programming technique. *Energy and Buildings*; 33; pp. 843-851.

Gustafsson, S. and Karlsson, B. G. 1989. Life-cycle cost minimization considering retrofits in multi-family residences. *Energy and Buildings*; 14; pp. 9-17.

Gustafsson, S. and RÖnnqvist, M. 2008. Optimal heating of large block of flats. *Energy and Buildings*; 40; pp. 1699-1708

Haldi, F. and Robinson, D. 2010. On the unification of thermal perception and adaptive actions. *Building and Environment*; 45; pp. 2440-2457.

Hall, J. F. and Klemm, F. K. 1967. Thermoregulatory responses in disparate thermal environments. *Journal of Applied Physiology*; 23; pp. 540-544.

Hamdy, M.; Hasan, A. and Siren, K. 2011. Impact of adaptive thermal comfort criteria on building energy use and cooling equipment size using a multi-objective optimization scheme. *Energy and Buildings*; 43; pp. 2055-2067.

Han, S. P. 1976. Superlinearly convergent variable metric algorithms for general nonlinear programming problems. *Mathematical Programming*; 11; pp. 263-282.

Hancock, H. 1917. *Theory of Maxima and Minima*. Ginn and company, Boston, New York.

Hand, J.W. 1998. *Removing barriers to the use of simulation in the building design professions*. PhD thesis, University of Strathclyde, Glasgow.

Hardy, J. D. and Stolwijk, J. A. J. 1966. Partial calorimetric studies of responses of man to thermal transients. *Journal of Applied Physiology*; 21; 1799-1806.

Hashiguchi, N.; Hirakawa, M.; Tochihara, Y.; Kaji, Y. and Karaki, C. Thermal Environment and Subjective Responses of Patients and Staff in a Hospital during Winter. *Journal of Physiological Anthropology and Applied Human Science*; 24(1); pp. 111-115.

Hensen, J.L.M. 1991. *On the thermal interaction and heating and ventilation system.* PhD thesis, University of Strathclyde, Glasgow.

Henze, G. P.; Felsmann, C. and Knabe, G. 2004. Evaluation of optimal control for active and passive building thermal storage. *International Journal of Thermal Sciences*; 43; pp. 173-183.

Henze, G.P.; Pfafferott, J.; Herkel, S. and Felsmann, C. 2007. Impact of adaptive comfort criteria and heat waves on optimal building thermal mass control. *Energy and Buildings*; 39; pp. 221–235.

Hestenes, M. R. and Stiefel, E. 1952. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*; 49; 409-436.

Hinton, G.E. 2007. Learning multiple layers of representation. *Trends in Cognitive Sciences*; 11(10); pp. 428-434.

Ho, S.L.; Yang, S.; Ni, P. and Wong, H.C. 2002. Developments of an efficient global optimal design technique – a combined approach of MLS and SA algorithm. *International Journal for Computation and Mathematics in Electrical and Electronic Engineering*; 21(4); pp. 604-614.

Holland, J.H. 1975. *Adaption in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, USA.

Holmes M.J.; Lam J.K.-W.; Ruddick K.G. and Whittle, G.E. 1990. Computation of Conduction, Convection, and Radiation in the Perimeter Zone of an Office Space. *Proc. ROOMVENT '90*, Oslo Norway.

Hong, T.; Chou, S.K. and Bong, T.Y. 2000. Building simulation: an overview of developments and information sources. *Building and Environment*; 35; pp. 347-361.

House, J. M. and Smith, T. F. 1995. Optimal control of building and HVAC systems. In: *Proceedings of The American Control Conference* (Seattle). pp. 4326-4330.

House, J. M.; Smith, T. F. and Arora, J. S. 1991. Optimal control of a thermal system. *ASHRAE Transactions*; 97; pp. 991-1001.

HSE (Health and Safety Executive). 2012. *Health Services Information sheet number no. 6: Managing the risks from hot water and surfaces in health and social care*. Website accessed 21/5/2016. Available online: http://www.hse.gov.uk/pubns/hsis6.pdf

HSCIC. 2014. Hospital Estates and Facilities Statistics, ERIC (Estates Return Information Collection). Website accessed 4/10/2014. Available online: http://hefs.hscic.gov.uk/ERIC.asp

Huizenga, C.; Zhang, H and Arens, E. 2001. A model of human physiology and comfort for assessing complex thermal environments. *Building and Environment*; 36; pp. 691-699.

Humphreys, M. and Nicol, J. F. 1998. Understanding the adaptive approach to thermal comfort. A*SHRAE Transactions* 104(1), pp. 991–1004.

Hwang, R.-L.; Lin, T.-P.; Cheng, M.-J. and Chien, J.-H. 2007. Patient thermal comfort requirement for hospital environments in Taiwan. *Building and Environment*; 42; pp. 2980–2987.

IBISWorld. 2011. Heating and Air-Conditioning Equipment Manufacturing: U.S. Industry Report (free preview). Website accessed 20th April 2011, available on the World Wide Web: http://www.ibisworld.com.cn/industry/default.aspx?indid=698&rcid=1

IES Limited. 2014. *IESVE for Engineers*. Website accessed 19th March 2013, available on the World Wide Web: http://www.iesve.com/software/ve-pro

Imre, L.; Bitai, A.; Horvath, C. D.; Banhidi, L. and Pammer, Z. 1988. Thermal analysis of human body–clothing–environment system. *International Journal for Numerical Methods in Engineering*; 25; pp. 357–371.

Jin, M.; Zuo, W. and Chen, Q. 2013. Simulating Natural Ventilation in and Around Buildings by Fast Fluid Dynamics. *Numerical Heat Transfer, Part A: Applications: An International Journal of Computation and Methodology*; 64(4); pp. 273-289.

John, F. 1948. Extremum problems with inequalities as subsidiary conditions. In: *Studies and Essays, Presented to R. Courant on his 60th Birthday January 8, 1948*. pp. 187–204. Interscience, New York.

Jones, B. W. and Ogawa, Y. 1992. Transient interaction between the human and the thermal environment. *ASHRAE Transactions*; 98; pp. 189–195.

Jones, B. W. and Ogawa, Y. 1993. Transient response of the human–clothing system. *Journal of Thermal Biology*; 18; pp. 413–416.

Jones, P.J. and Whittle, G.E. 1992. Computational Fluid Dynamics for Building Air Flow Prediction—Current Status and Capabilities. *Building and Environment*; 27(3); pp. 321-38.

Juan, Y.; Gao, P. and Wang, J. 2010. A hybrid decision support system for sustainable office building renovation and energy performance improvement. *Energy and Buildings*; 42; pp. 290–297.

Judkoff, R. and Neymark, J. 1995. *International Energy Agency Building Energy Simulation Test (BESTEST) and Diagnostic Method.* NREL Report No. TP-472-6231.

Jurovics, S.A. 1978. Optimization applied to design of an energy-efficient building. *IBM Journal of Research and Development*; 22(4); pp. 378-385.

Kantoriovich, L. K. 1960 (English reprint of 1939 Russian original). Mathematical methods of organizing planning production. *Management Science*; 6; pp. 366-422.

Karmarkar, N. 1984. A new polynomial-time algorithm for linear programming. *Combinatorica*; 4; pp. 373-395.

Karush, W. 1939. *Minima of functions of several variables with inequalities as side constraints.* M.Sc. Dissertation. Department of Mathematics, University of Chicago, Chicago, Illinois.

Kaynakli, O. and Kilic, M. 2005. Investigation of indoor thermal comfort under transient conditions. *Building and Environment*; 40; pp. 165–174.

Kaynakli, O.; Unver, U. and Kilic, M. 2003. Evaluating thermal environments for sitting and standing posture. *International Communications in Heat and Mass Transfer*; 30; pp. 1179–1188.

Kelly, N.J. 1998. *Towards a design environment for building integrated energy systems: The integration of electrical power flow modelling with building simulation.* PhD thesis, University of Strathclyde, Glasgow.

Kennedy, J. and Eberhart, R. 1995. Particle swarm optimization. In: *Proceedings of the IEEE International Conference on Neural Networks* (Perth). IEEE Service Center, Piscataway, New Jersey; pp. 1942-1948.

Khan, M.A.I.; Noakes, C.J. and Toropov, V.V. 2012. Development of a numerical optimization approach to ventilation system design to control airborne contaminant dispersion and occupant comfort. *Building Simulation*; 5; pp. 39-50.

Kiefer, J. and Wolfowitz, J. 1952. Stochastic estimation of the maximum of a regression function. *Annals of Mathematical Statistics*; 23; pp. 462-466.

Kim, C.H. and Arora, J.S. 2003. Development of simplified dynamic models using optimization: Application to crushed tubes. *Computational Methods in Applied Mechanical Engineering*; 192; pp. 2073–2097.

King, M.-F.; Noakes, C. J.; Sleigh, P.A. and Camargo-Valero, M.A. 2013. Bioaerosol deposition in single and two-bed hospital rooms: A numerical and experimental study. *Building and Environment*; 59; pp. 436-447

Kintner-Meyer, M. and Emery, A. F. 1995. Optimal control of an HVAC system using cold storage and building thermal capacitance. *Energy and Buildings*; 23; pp. 19-31.

Kjeldsen, T. H. 2000. A contextualized historical analysis of the Kuhn–Tucker theorem in nonlinear programming: The impact of World War II. *Historia Mathematica*; 27; pp. 331-361.

Klemm, K; Wojciech Marks, W. and Klemm, A.J. 2000. Multicriteria optimisation of the building arrangement with application of numerical simulation. *Building and Environment*; 35; pp. 537-544.

Kolokotsa, D.; Stavrakakis, G. S.; Kalaitzakis, K. and Agoris, D. 2002. Genetic algorithms optimized fuzzy controller for the indoor environmental management in buildings implemented using PLC and local operating networks. *Engineering Application of Artificial Intelligence*; 15; pp. 417-428.

Kuhn, H. W. and Tucker, A. W. 1951. Nonlinear Programming. In: *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability* (Conference held 1950, UC Berkeley). University of California Press, California.

Kusuda, T. 1976. *NBSLD: The Computer Program for Heating and Cooling Loads in Buildings*, NBS Building Science Series No. 69, National Bureau of Standards, Washington USA.

Lancaster, P. and Salkauskas, K. 1981. Surfaces Generated by Moving Least Squares Methods. *Mathematics of Computation*; 37(155); pp. 141-158.

Launder, B.E. and Spalding, D.B. 1974. The Numerical Computation of Turbulent Flows. *Computer Methods in Applied Mechanics and Engineering*; 3; pp. 269-289.

Levin, D. 1998. The Approximation Power of Moving Least-Squares. *Mathematics of Computation*; 67; pp. 1517-1531.

Lin, Z. and Deng, S. 2008. A study on the thermal comfort in sleeping environments in the subtropics—Measuring the total insulation values for the bedding systems commonly used in the subtropics. *Building and Environment*; 43; pp. 905–916

Lomas, K.J. Ed. 1994. *Empirical Validation of Thermal Building Simulation Programs Using Test Room Data, Volume 3: Working Reports*. IEA Energy Conservation in Buildings and Community Systems Programme Annex 21 and IEA Solar Heating and Cooling Programme Task 12.

Lomas, K.J.; Eppel, H.; Martin, C. and Bloomfield D. 1994. *Empirical Validation of Thermal Building Simulation Programs Using Test Room Data, Volume 1: Final Report*. IEA Energy Conservation in Buildings and Community Systems Programme Annex 21 and IEA Solar Heating and Cooling Programme Task 12.

Lomas, K.J.; Martin, C.; Eppel, H.; Watson, M. and Bloomfield D. 1994. *Empirical Validation of Thermal Building Simulation Programs Using Test Room Data, Volume 2: Empirical Validation Package*. IEA Energy Conservation in Buildings and Community Systems Programme Annex 21 and IEA Solar Heating and Cooling Programme Task 12.

Lomas, K.J. and Giridharan, R. 2012. Thermal comfort standards, measured internal temperatures and thermal resilience to climate change of free-running buildings: A case-study of hospital wards. *Building and Environment*; 55; pp. 57-72.

Lomas, K.J.; Giridharan, R.; Short, C.A. and Fair, A.J. 2012. Resilience of 'Nightingale' hospital wards in a changing climate. *Building Services Engineering Research and Technology*; 33(1); pp. 81-103.

Lu, L.; Cai, W.; Xie, L.; Li, S. and Soh, Y. C. 2005. HVAC system optimization—in-building section. *Energy and Buildings*; 37; pp. 11-22.

Mackenzie, C.A. and Gero, J.S. 1987. Learning design rules from decisions and performances *Artificial Intelligence*; 2(1); pp. 2-10.

MacQueen, J. 1997. *The modelling and simulation of energy management control systems*. PhD thesis, University of Strathclyde, Glasgow.

Magnier, L. and Haghighat, F. 2010. Multiobjective optimization of building design using TRNSYS simulations, genetic algorithm, and artificial neural network. *Building and Environment*; 45; pp. 739-746.

Maratos, N. 1978. *Exact penalty function algorithms for finite-dimensional and control optimization problems.* PhD thesis, Department of Computing and Control, University of London.

Markowitz, H. 1952. Portfolio selection. *The Journal of Finance*; 7; pp. 77-91.

Martin, R. D. and Masreliez, C. J. 1975. Robust estimation via stochastic approximation. *IEEE Transaction on Information Theory*; 21; pp. 263-271.

McCartney, K.J. and Nicol, J.F. 2002. Developing an adaptive control algorithm for Europe. *Energy and Buildings*; 34; pp. 623–635

McCulloch, W.S. and Pitts, W.H. 1943. A Logical Calculus of the Ideas Immanent in Nervous Activity. *Bulletin of Mathematical Biophysics*; 7; pp. 115–133.

McNall Jr, P. E. and Biddison, R. E. 1970. Thermal and comfort sensations of sedentary persons exposed to asymmetric radiant fields. *ASHRAE Transactions*; 76; pp. 123-136.

Megiddo, N. 1989. Pathways to the optimal set in linear programming. In: *Progress in Mathematical Programming: Interior-Point and Related Methods*. Springer-Verlag, New York.

Mehrotra, S. 1992. On the implementation of a primal-dual interior point method. *SIAM Journal on Optimization*; 2; pp. 575-601.

Menger, K. 1932. "Das botenproblem", in Ergebnisse eines Mathematischen Kolloquiums; 2. Teubner, Leipzig, pp. 11-12.

Miyanaga, T.; Urabe, W. and Nakano, Y. 2001. Simplified human body model for evaluating thermal radiant environment in a radiant cooled space. *Building and Environment*; 36; pp. 801-808.

Moroşan, P.; Bourdais, R.; Dumur, D. and Buisson, J. 2010. Building temperature regulation using a distributed model predictive control. *Energy and Buildings*; 42; pp. 1445–1452.

Murakami, S.; Kato, S. and Zeng, J. 2000. Combined simulation of airflow, radiation and moisture transport for heat release from a human body. *Building and Environment*; 35; pp. 489–500.

Nakhi, A.E. 1995. *Adaptive construction modelling within whole building dynamic simulation*. PhD thesis, University of Strathclyde, Glasgow.

Negrão, C.O.R. 1995. *Conflation of computational fluid dynamics and building thermal simulation.* PhD thesis, University of Strathclyde, Glasgow.

Nelder, J. A. and Mead, R. 1965. A simplex method for function minimization. *Computer Journal*; 7; pp. 308-313.

Newton, I. (translated by Motte, A.). 1729; first published in Latin in 1687. *The Mathematical Principles of Natural Philosophy (In two volumes)*. Middle-Temple-Gate, London; digitized by Google. Website accessed 20th April 2011, available on the World Wide Web:

volume I: http://books.google.com/books?id=Tm0FAAAAQAAJ&pg=PA1#v=onepage&q&f=false

volume II: http://books.google.com/books?id=6EqxPav3vIsC&pg=PA1#v=onepage&q&f=false

Nicol, F. and Humphreys, M. 2007. Maximum temperatures in European office buildings to avoid heat discomfort. *Solar Energy*; 81; pp. 295–304.

Nicol, F. and Humphreys, M. 2010. Derivation of the adaptive equations for thermal comfort in free-running buildings in European standard EN15251. *Building and Environment*; 45; pp. 11–17.

Nobelprize.org. 2013. The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 1975. Accessed 10/8/2011. Available online: http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/1975/

Olufsen, M. S.; Peskin, C. S.; Kim, W. Y.; Pedersen, E. M.; Nadim, A. and Larsen, J. 2000. Numerical simulation and experimental validation of blood flow in arteries with structured tree outflow conditions. *Annals of Biomedical Engineering*; 28; pp. 1281–1299.

Otto, M.; Uwe, H. and Wegsheider, W. 1989. Multi-dimensional interpolation by the moving least squares approach for modelling of chromatographic retention data. *Journal of Chromatography*; 485; pp. 453-460.

Ozel, M. and Pihtili, K. 2007. Optimum location and distribution of insulation layers on building walls with various orientations. *Building and Environment*; 42; pp. 3051–3059.

Pantelic, J.; Raphael, B. and Tham, K.W. 2012. A preference driven multi-criteria optimization tool for HVAC design and operation. *Energy and Buildings*; 55; pp. 118–126.

Parameshwaran, R.; Karunakaran, R.; Kumar, C. V. R. and Iniyan, S. 2010. Energy conservative building air conditioning system controlled and optimized using fuzzy-genetic algorithm. *Energy and Buildings*; 42; pp. 745-762.

Park, C. 2003. *Occupant responsive optimal control of smart façade systems.* PhD thesis, College of Architecture, Georgia Institute of Technology.

Pennes, H.H. 1948. Analysis of tissue and arterial blood temperatures in the resting human forearm. *Journal of Applied Physiology*; 1; pp. 93-122.

Piccinini, G. 2004. The first computational theory of mind and brain: a close look at Mcculloch and Pitts's ''logical calculus of ideas immanent in nervous activity''. *Synthese*; 141; pp. 175-215.

Polyak, B.T. 1969. The conjugate gradient method. *Proceedings of the 2nd Winter School on Mathematical Programming and Allied Problems* (Drogobych; in Russian); pp. 1.

Pourshaghaghy, A. and Omidvari, M. 2012. Examination of thermal comfort in a hospital using PMV-PPD model. *Applied Ergonomics*; 43; pp. 1089-1095.

Powell, M. J. D. 1978. A fast algorithm for nonlinearly constrained optimization calculations. In: *Numerical Analysis, no. 630 in Lecture Notes in Mathematics.* Springer Verlag, pp. 144-157.

Radford, A.D. and Gero, J.S. 1980. On optimization in computer aided architectural design. *Building and Environment*; 15; pp. 73-80.

Robbins, H. and Monro, S. 1951. A stochastic approximation method. *Annals of Mathematical Statistics*; 22; pp. 400-407.

Rodi, W. 1980. *Turbulence Models and their Applications in Hydraulics—A State of the Art Review*. International Association for Hydraulic Research; Delft; The Netherlands.

Rollins, D.; Bhandar, N. and Hulting, S. 2006. System identification of the human thermoregulatory system using continuous-time block-oriented predictive modeling. *Chemical Engineering Science*; 61; pp. 1516-1527.

Ross, D.E. 1977. Building automation - energy optimization by computer. *Architectural record*; 161(5); pp. 143-144.

Rudolph, G. 2001. Self-adaptive mutations may lead to premature convergence. *IEEE Transactions on Evolutionary Computation*; 5; pp. 410-414.

Salloum, M.; Ghaddar, N. and Ghali, K. 2007. A new transient bioheat model of the human body and its integration to clothing models. *International Journal of Thermal Sciences*; 46; pp. 371-384.

Sambou, V.; Lartigue, B.; Monchoux, F. and Adj, M. 2009. Thermal optimization of multilayered walls using genetic algorithms. *Energy and Buildings*; 41; pp. 1031–1036.

Samuel, A.A. 2006. *On the conflation of contaminant behaviour prediction within whole building performance simulation.* PhD thesis, University of Strathclyde, Glasgow.

Schaelin, A.; Doror, V.; van der Maas, J. and Moser, A. 1994. Improvement of multizone model predictions by detailed flow path values from CFD calculations. *ASHRAE Transactions*; 100(2); pp. 709-720.

Schellen, L.; van Marken Lichtenbelt, W. D.; Loomans, M.G.L.C.; Toftum, J. and de Wit, M.H. 2010. Differences between young adults and elderly in thermal comfort, productivity, and thermal physiology in response to a moderate temperature drift and a steady-state condition. *Indoor Air*; 20; pp. 273-283.

Schlegel, J. C. and McNall, P. E. 1968. The effect of asymmetric radiation on the thermal and comfort sensations of sedentary subjects. *ASHRAE Transactions*; 74; pp. 144-154.

Schrijver, A. 2005. On the history of combinatorial optimization (till 1960). *Handbooks in Operations Research and Management Science*; 12; pp. 1-68.

Shewchuk, J.R. 1994. *An introduction to the conjugate gradient method without the agonizing pain, edition 1¼.* School of Computer Science, Carnegie Mellon University, Pittsburgh.

Shor, N. Z. 1985. *Minimization methods for non-differentiable functions*. Springer-Verlag, New York.

Shor, N. Z. and Zhurbenko, N. G. 1971. A minimization method using the operation of extension of the space in the direction of the difference of two successive gradients. *Cybernetics and Systems Analysis (translated from Kibernetica)*; 7; pp. 450-459.

Short, C.A.; Lomas, K.J.; Giridharan, R. and Fair, A.J. 2012. Building resilience to overheating into 1960's UK hospital buildings within the constraint of the national carbon reduction target: Adaptive strategies. *Building and Environment*; 55; pp. 73-95.

Simmonds, P. 1994. The utilization of optimal design and operation strategies in lowering the energy consumption in office buildings. *Renewable Energy*; 5(2); pp. 1193-1201.

Skoog, J.; Fransson, N. and Jagemar, L. 2005. Thermal environment in Swedish hospitals: Summer and winter measurements. *Energy and Buildings*; 37; pp. 872–877.

Smith, C. E. 1991. A transient three-dimensional model of the human thermal system. *PhD thesis*; Department of Mechanical Engineering, Kansas State University, Manhattan.

Sofrata, H. and Salmeen, B. 1993. Optimization of insulation thicknesses using micros. *Energy Convers. Mgmt*; 34(6); pp. 471-479.

Spall, J. C. 1992. Multivariate stochastic approximation using a simultaneous gradient approximation. *IEEE Transactions on Automatic Control*; 37; pp. 332-341.

Spendley, W.; Hext, G. R. and Himsworth, F. R. 1962. Sequential application of simplex designs in optimisation and evolutionary operation. *Technometrics*; 4; pp. 441-461.

Srebric, J.; Chen, Q. and Glicksman, L. R. 1999. Validation of a Zero-Equation Turbulence Model for Complex Indoor Airflow Simulation. *ASHRAE Transactions*; 105 (2); pp. 414-427.

Stavrakakis, G.M.; Zervas, P.L.; Sarimveis, H. and Markatos, N.C. 2012. Optimization of window-openings design for thermal comfort in naturally ventilated buildings. *Applied Mathematical Modelling*; 36; pp. 193–211.

Stephenson, D.G. and Mitalas, G.P. 1967. Cooling Load Calculations by Thermal Response Factor Method, *ASHRAE Transactions*, 73(1); pp. 508-515.

Stolwijk, J. A. J. 1971. *NASA contractor report: A mathematical model of physiological temperature regulation in man.* Yale University School of Medicine.

Stolwijk, J. A. J. and Hardy, J. D. 1966. Partial calorimetric studies of responses of man to thermal transients. *Journal of Applied Physiology*; 21; pp. 967-977.

Tanabe, S.; Kobayashi, K.; Nakano, J.; Ozeki, Y. and Konishi, M. 2002. Evaluation of thermal comfort using combined multi-node thermoregulation (65MN) and radiation models and computational fluid dynamics (CFD). *Energy and Buildings*; 34; pp. 637–646.

Toropov, V.V.; Filatov, A.A. and Polynkin, A.A. 1993. Multiparameter structural optimization using FEM and multipoint explicit approximations. *Structural Optimization*; 6; pp. 7-14.

Toy, S. and Yilmaz, S. 2010. Thermal sensation of people performing recreational activities in shadowy environment: a case study from Turkey. *Theoretical Applied Climatology*; 101; pp. 329-343.

Turing, A.M. 1936–37 [1965]. On Computable Numbers, with an Application to the Entscheidungsproblem. In M. Davis (ed.), *The Undecidable*, Raven, Ewlett, pp. 116–154.

UKCP. 2009. *UK Climate Projections*. Website accessed 3/10/2012, available on the World Wide Web: http://ukclimateprojections.metoffice.gov.uk/.

Verheyen, J.; Theys, N.; Allonsius, L. and Descamps, F. 2011. Thermal comfort of patients: Objective and subjective measurements in patient rooms of a Belgian healthcare facility. *Building and Environment*; 46; pp. 1195-1204.

Versteeg, H. K. and Malalasekera, W. 1995. *An Introduction to Computational Fluid Dynamics: The Finite Volume Method*. Longman Group.

Von Neumann, J. and Morgenstern, O. 2007 (reprint of 1944 original). *Theory of Games and Economic Behaviour (Commemorative Edition)*. Princeton University Press.

Vose, M.D. 1999. *The Simple Genetic Algorithm: Foundations and Theory*. MIT Press.

Walton, G.N. 1983. *Thermal Analysis Research Program Reference Manual*, NBSIR 83-2655, National Bureau of Standards, Washington.

Wang, S. and Jin, X. 2000. Model-based optimal control of VAV air-conditioning system using genetic algorithm. *Building and Environment*; 35; pp. 471-487.

Wang, W.; Rivard, H. and Zmeureanu, R. 2005. An object-oriented framework for simulation-based green building design optimization with genetic algorithms. *Advanced Engineering Informatics*; 19; pp. 5–23.

Werbos, P.J. 1975. *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*. PhD thesis, Harvard University, USA.

Wetter, M., 2000. Design optimization with GenOpt. Building Energy Simulation, User News, Lawrence Berkeley National Laboratory, vol. 21.

Wilson, R. B. 1963. *A Simplicial Method for Convex Programming*. PhD thesis, Harvard University, USA.

Wissler, E. 1964. A mathematical model of the human thermal system. *Bulletin of Mathematical Biophysics*; 26**;** pp. 147–167.

Wright, J. A.; Loosemore, H. A. and Farmani, R. 2002. Optimization of building thermal design and control by multi-criterion genetic algorithm. *Energy and Buildings*; 34; pp. 959-972.

Xiong, G.; Lu, M.; Chent, C.L.; Wangs, B.P. and Kehlt, D. 2001. Numerical Optimization of a Power Electronics Cooling Assembly. In: *Annual IEEE Conference on Applied Power Electronics Conference and Exposition (APEC)* (held March 4-8, 2001, London); pp. 1068-1073.

Xu, X. and Werner, J. 1997. A dynamic model of the human/clothing/environment system. *Applied Human Science*; 16; pp. 61-75.

Xu, X.; Wang, S.; Sun, Z. and Xiao, F. 2009. A model-based optimal ventilation control strategy of multi-zone VAV air-conditioning systems. *Applied Thermal Engineering*; 29; pp. 91-104.

Xu, Y.; Chen, X. and Zhang, C. 2008. Optimization method for indoor thermal comfort based on interactive numerical calculation. *Journal of Beijing Institute of Technology*; 17; pp. 495-500.

Yang, B.; Sekhar, S. C. and Melikov, A. K. 2010. Ceiling-mounted personalized ventilation system integrated with a secondary air distribution system – a human response study in hot and humid climate. *Indoor Air*; 20; pp. 309-319.

Yang, Z.; Li, X.; Bowers, C.P.; Schnier, T.; Tang, K. and Yao, X. 2012. An Efficient Evolutionary Approach to Parameter Identification in a Building Thermal Model. *IEEE transactions on systems, man, and cybernetics—Part C: applications and reviews*; 42(6); pp. 957-969.

Yang, X.; Zhao, L.; Bruse, M. and Meng, Q. 2012. An integrated simulation method for building energy performance assessment in urban environments. *Energy and Buildings*; 54; pp. 243-251.

Yau, Y.H. and Chew, B.T. 2009. Thermal comfort study of hospital workers in Malaysia. *Indoor Air*; 19: pp. 500–510.

Yi, L.; Fengzhi, L.; Yingxi, L. and Zhongxuan L. 2004. An integrated model for simulating interactive thermal processes in human-clothing system. *Journal of Thermal Biology*; 29; pp. 567–575.

Yigit, A. 1998. The computer-based human thermal model. *International communications in heat and mass transfer*; 25; pp. 969-977.

Zhang, H.; Arens, E.; Huizenga, C. and Han, T. 2010a. Thermal sensation and comfort models for non-uniform and transient environments: part I: local sensation of individual body parts. *Building and Environment*; 45; pp. 380-388.

Zhang, H.; Arens, E.; Huizenga, C. and Han, T. 2010b. Thermal sensation and comfort models for non-uniform and transient environments: part II: local comfort of individual body parts. *Building and Environment*; 45; pp. 389-398.

Zhang, H.; Arens, E.; Huizenga, C. and Han, T. 2010c. Thermal sensation and comfort models for non-uniform and transient environments: part III: whole-body sensation and comfort. *Building and Environment*; 45; pp. 399-410.

Zhang, Y. and Zhao, R. 2008. Overall thermal sensation, acceptability and comfort. *Building and Environment*; 43; pp. 44-50.

Zhang, Y.; Wang, J.; Chen, H.; Zhang, J. and Meng, Q. 2010. Thermal comfort in naturally ventilated buildings in hot-humid area of China. *Building and Environment*; 45; pp. 2562-2570.

Zhou, L. and Haghighat, F. 2009a. Optimization of ventilation system design and operation in office environment, Part I: Methodology. *Building and Environment*; 44; pp. 651-656.

Zhou, L. and Haghighat, F. 2009b. Optimization of ventilation system design and operation in office environment, Part II: Results and discussions. *Building and Environment*; 44; pp. 657-665.

Zolfaghari, A. and Maerefat, M. 2010. A new simplified model for evaluating non-uniform thermal sensation caused by wearing clothing. *Building and Environment*; 45; pp. 776-783.

Zolfaghari, A. and Maerefat, M. 2010. A new simplified thermoregulatory bioheat model for evaluating thermal response of the human body to transient environments. *Building and Environment*; 45; pp. 2068-2078.

Zoutendijk, G. 1960. *Methods of feasible directions: A study in linear and nonlinear programming.* PhD thesis, University of Amsterdam.

Zuo, W. and Chen, Q. 2010. Fast and informative flow simulations in a building by using fast fluid dynamics model on graphics processing unit. *Building and Environment*; 45; pp. 747-757.