



**TÉCNICO**  
LISBOA

**UNIVERSIDADE DE LISBOA**  
**INSTITUTO SUPERIOR TÉCNICO**



**Context-Based Thermodynamic Modeling and  
Optimal Management of Energy and Occupant  
Comfort for Building Spaces**

**Pedro Viçoso Fazenda**

**Supervisor:** Doctor Pedro Manuel Urbano de Almeida Lima  
**Co-Supervisor:** Doctor Una-May O'Reilly

**Thesis approved in public session to obtain the PhD Degree in  
Sustainable Energy Systems**

**Jury final classification: Pass with Distinction**

**Jury**

**Chairperson:** Chairman of the IST Scientific Board

**Members of the Committee:**

Doctor Pedro Manuel Urbano de Almeida Lima  
Doctor José Manuel Prista do Valle Cardoso Igreja  
Doctor Una-May O'Reilly  
Doctor Vasco Miguel Moreira Amaral  
Doctor Paulo Jorge Fernandes Carreira  
Doctor Carlos Augusto Santos Silva

**2016**





**UNIVERSIDADE DE LISBOA**  
**INSTITUTO SUPERIOR TÉCNICO**  
**MIT Portugal Program**

**Context-Based Thermodynamic Modeling and  
Optimal Management of Energy and Occupant  
Comfort for Building Spaces**

**Pedro Viçoso Fazenda**

**Supervisor:** Doctor Pedro Manuel Urbano de Almeida Lima

**Co-Supervisor:** Doctor Una-May O'Reilly

**Thesis approved in public session to obtain the PhD Degree in  
Sustainable Energy Systems**

**Jury final classification: Pass with Distinction**

**Jury**

**Chairperson:** Chairman of the IST Scientific Board

**Members of the Committee:**

Doctor Pedro Manuel Urbano de Almeida Lima, Professor Associado (com Agregação) do Instituto Superior Técnico da Universidade de Lisboa

Doctor José Manuel Prista do Valle Cardoso Igreja, Professor Coordenador do Instituto Superior de Engenharia de Lisboa – Instituto Politécnico de Lisboa

Doctor Una-May O'Reilly, Principal Research Scientist, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, USA

Doctor Vasco Miguel Moreira Amaral, Professor Auxiliar da Faculdade de Ciências e Tecnologia da Universidade Nova de Lisboa

Doctor Paulo Jorge Fernandes Carreira, Professor Auxiliar do Instituto Superior Técnico da Universidade de Lisboa

Doctor Carlos Augusto Santos Silva, Professor Auxiliar do Instituto Superior Técnico da Universidade de Lisboa

**Funding Institution**

FCT - Fundação para a Ciência e a Tecnologia

**2016**



## Resumo

Os edifícios são responsáveis por uma parte significativa da fatura energética global. Prevê-se, porém, que a integração de novas tecnologias na gestão de sistemas e serviços, tenha impactos significativos na promoção da eficiência energética neste sector. Contribuindo nesse sentido, esta dissertação começa por apresentar um método para melhorar o desempenho do sistema de aquecimento, ventilação e ar condicionado. Como a operação deste sistema não está, em muitos casos, devidamente regulada e ajustada ao perfil de ocupação dos edifícios, o resultado é o desperdício de energia. Como tal, a solução aqui apresentada inclui a utilização de um controlador, com aprendizagem por reforço, capaz de aprender as regularidades estatísticas da ocupação de um espaço e o correto agendamento das temperaturas para esse espaço. Por simulação demonstra-se que o controlador proposto é capaz de regular o aquecimento por forma a minimizar os custos energéticos garantindo, contudo, o conforto dos ocupantes.

No seguimento deste estudo, a linha de investigação centrou-se na busca de soluções para representar o ambiente do edifício com foco particular no comportamento térmico dos espaços. Isto porque existem eventos, associados por exemplo a abertura e fecho de porta e janelas, que alteram esse comportamento. Os modelos termodinâmicos tradicionalmente utilizados são normalmente invariantes no tempo e não contemplam estas alterações. Como tal, ficam desajustados à realidade desses espaços. Para modelar uma zona térmica esta dissertação propõe a utilização de um modelo termodinâmico baseado num sistema híbrido. Este modelo é capaz de representar os diferentes comportamentos térmicos associados aos diferentes contextos do espaço, definidos como os modos discretos do sistema. Apresenta-se uma aplicação-exemplo e para validar o conceito utiliza-se os resultados do simulador EnergyPlus como referência de comparação. Os resultados demonstram que o modelo proposto consegue representar, com exatidão aceitável, as temperaturas de uma zona térmica através de diferentes contextos.

**Palavras Chave:** HVAC, Aprendizagem por Reforço, Raciocínio baseado em Contextos, Sistemas Híbridos, Comportamento Térmico, EnergyPlus, Edifício Inteligente, Inteligência Ambiente, Q-Learning, Domótica, Eficiência Energética.



# Abstract

Buildings are responsible for a considerable amount of the global energy bill and energy savings can be accomplished by integrating more efficient building technologies. Occupants tend to forget to adjust systems appropriately and, in many spaces, the conditioning requirements are not adjusted to the occupancy of those spaces. As a result, energy is unnecessarily wasted. This dissertation presents a method to enhance the management of the heating, ventilation, and air conditioning systems in buildings. Since these systems are some of the most energy-demanding services in buildings, we describe the application of a reinforcement-learning-based supervisory control approach that actively learns how to appropriately schedule thermostat temperature set points. The result is a learning controller that learns the statistical regularities in the occupants' behavior, allowing them to achieve comfort requirements while optimizing energy costs.

The study then proceeds towards finding suitable representations for building environments, with a particular focus on how to represent the thermal behavior of building spaces. The occurrence of events such as doors, windows and blinds being opened or closed, can drastically affect the underlying processes that govern the dynamics of temperature evolution of building spaces rendering standard thermodynamic models less effective for control and prediction. Therefore, a framework is presented for model structure and parameter selection that takes these events into account, based on the notion of context. Contexts are modeled as discrete states of a hybrid system and depending on how context changes, the thermodynamic model transitions through a set of different linear time-invariant sub-models. Each sub-model is effective in representing the thermal behavior of a building space in its associated context. We present an application example and use the outputs of EnergyPlus as reference for model performance evaluation. We show, through different context changes, how a context-based model can be used to represent, with reasonable accuracy, the evolution of temperatures in a simulated thermal zone.

**Keywords:** HVAC, Reinforcement Learning, Context-Awareness, Thermodynamic Modeling, Hybrid Systems, EnergyPlus, Smart Buildings, Ambient Intelligence, Q-Learning, Home Automation.



# Preface and Acknowledgements

First and foremost, I would like to express my sincere gratitude to my advisor, Professor Pedro Lima, from *Instituto Superior Técnico* (IST), and co-adviser, Principal Research Scientist Una-May O'Reilly, from the *Massachusetts Institute of Technology* (MIT). It has been an honor and an inspiration to work with both of them. The completion of my dissertation has been a long journey and I appreciate all of their contributions of time, ideas, and guidance to make my PhD experience productive and stimulating.

I am thankful to Pedro Lima for the excellent example he has provided as a group leader. I started working with him in 2004 when I was completing my master's degree in Robotics and Control. Pedro was my adviser, and at that time, we were working on a novel approach for the formation control of robot swarms. The *Intelligent Robots and Systems group* (IRSg), of which we are members, is composed of enthusiastic researchers working in machine learning and other control algorithms with several different applications to robotics. Some of the work developed by the group could be extended to energy systems and it was in the advent of a research collaboration with the *Center for Innovation, Technology and Policy Research* (IN<sup>+</sup>) that I was challenged to pursue a doctoral degree in the MIT Portugal Program.

I would like to express my gratitude to all of the collaborators of the MIT Portugal Program, especially to Prof. Paulo Ferrão and Prof. Carlos Silva for encouraging me in this challenge. I am grateful to our Project Manager, Ana Quaresma, the Director of Education, João Fumega, and the rest of the staff for their support. I also want to express my gratitude to all of the teachers and fellow students whom I met in the program.

The completion of my dissertation and subsequent PhD has been a long journey that has opened new doors, led to many new friends, and given me the opportunity to work with Una-May and the *AnyScale Learning For All* (ALFA) group at MIT's *Computer Science and Artificial Intelligence Laboratory* (CSAIL). During my stay

at MIT, the ALFA group was a source of friendship, as well as good advice and collaboration. I am grateful for my time spent with Dr. Kalyan Veeramachaneni, Dr. James McDermott, and Una-May.

My experience at MIT has been nothing short of amazing. My time at CSAIL was made enjoyable in large part due to the exceptional working and stimulating environment, in which members are encouraged to participate in activities, communicate, share and debate ideas, and create long-lasting relationships. I will forever remember and miss this wonderful period of my life.

I am also thankful to the *Institute for Systems and Robotics* (ISR) at IST and to all of my colleagues for providing a similar environment to the one I had at CSAIL here in Portugal. At ISR, I have met like-minded students who share the same passion for engineering and science. They have contributed, over the years, countless ideas and insightful questions. During these years at ISR, I have also appreciated the camaraderie and local expertise of many scholars. In particular, I must acknowledge the relational influence of Professor Isabel Ribeiro, as well as her rigorous scientific attitude and critical thinking.

I am indebted to many colleagues and friends who have supported and encouraged me, making the realization of this work possible. Besides my advisors, I am especially grateful to Professor Paulo Carreira for his research contribution, insightful comments, and rigorous reviews. I am grateful to all of my colleagues from the *Área Departamental de Engenharia Electrónica e de Telecomunicações e de Computadores* at *Instituto Superior de Engenharia de Lisboa* (ADEETC/ISEL), where I lecture as an Assistant Teacher. Many colleagues at ADEETC assumed part of my workload and teaching duties so that I could pursue this research, particularly António Teófilo and Diogo Remédios. Without their contribution, it would have been very hard to complete this thesis. I am also indebted to my fellow co-founders of ET-Concept and FI-Sonic João Casaleiro, Tiago Oliveira, Ricardo Reis and Joel Paulo for managing and executing our company projects while I was completing my PhD. I would like to express my gratitude to all of my colleagues from *Centro de Estudos e Desenvolvimento de Electrónica e de Telecomunicações* (CEDET) at ISEL, and especially to António Couto Pinto, Helena Sousa Ramos, Miguel Campilho Gomes, Carlos Carvalho, José Rocha and Vítor Costa for their encouragement and support.

I gratefully acknowledge the funding sources that made my PhD work and my stay in the U.S. possible. I received financial support from the Portuguese National Science and Technology Foundation (Fundação para a Ciência e Tecnologia – FCT),

under the strategic project FCT [UID/EEA/50009/2013] and by the grant number SFRH/BD/60481/2009.

Last but not least, I would like to thank my family for all of their love and encouragement. For my mother, who raised me with a love of science and supported me in all of my pursuits. Lastly, I would like to thank my partner for life, Yvonne Weber, for all of her love and support. She has been a source of love and energy ever since I met her during my time at MIT. She is, by far, the best thing that has happened to me during this adventure.



# Contents

<b>List of Figures</b>	<b>xii</b>
<b>List of Tables</b>	<b>xv</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation . . . . .	5
1.1.1 Energy and Buildings . . . . .	6
1.1.2 Green Buildings and Energy Efficiency . . . . .	7
1.1.3 Smart Buildings . . . . .	10
1.2 Research Questions and Objectives . . . . .	16
1.2.1 Automatic HVAC Optimization . . . . .	17
1.2.2 Thermodynamic Modeling of Building Spaces . . . . .	20
1.3 Contributions . . . . .	24
1.4 Publications . . . . .	24
1.5 Thesis Outline . . . . .	25
1.6 Summary . . . . .	27
<b>2 Related Work</b>	<b>29</b>
2.1 Energy and Comfort Management . . . . .	30
2.1.1 Models for Building Systems and Comfort Evaluation . . . . .	31
2.1.2 Scheduling . . . . .	35
2.1.3 Building and Occupant Interfaces . . . . .	37
2.2 Context-awareness for Ambient Intelligence . . . . .	41
2.3 Context-based Framework . . . . .	43
2.4 Summary . . . . .	44
<b>3 Reinforcement Learning for HVAC Control</b>	<b>45</b>
3.1 Background . . . . .	46
3.1.1 Sequential Decision Problems . . . . .	47
3.1.2 Q-Learning . . . . .	50
3.1.3 Continuous State and Action Spaces . . . . .	52
3.2 Application Problems . . . . .	57
3.2.1 The <i>Bang-bang Heater</i> Problem . . . . .	59
3.2.2 The <i>Set Point Heater</i> Problem . . . . .	61
3.3 Discussion . . . . .	64

3.4	Summary . . . . .	67
<b>4</b>	<b>Context-based Thermodynamic Models</b>	<b>69</b>
4.1	Operational Semantics of Context-based Models . . . . .	69
4.2	Example (The Hysteresis Controller) . . . . .	71
4.2.1	Thermodynamic Model . . . . .	72
4.2.2	Heater Operation . . . . .	72
4.3	Hybrid Time Sets and Executions . . . . .	74
4.4	Model Formulation for a Thermal Zone . . . . .	78
4.4.1	Full-Scale Lumped RC Thermal Model . . . . .	80
4.4.2	Application Example . . . . .	85
4.5	Summary . . . . .	87
<b>5</b>	<b>Simulation Setup</b>	<b>89</b>
5.1	Reinforcement Learning Simulation Setup . . . . .	89
5.1.1	Occupant Schedule Simulation . . . . .	90
5.1.2	Modeling Comfort . . . . .	90
5.1.3	Performance Evaluation . . . . .	92
5.1.4	The <i>Bang-bang Heater</i> Problem . . . . .	94
5.1.5	Occupant <i>Uncomfortable</i> State Redefined . . . . .	97
5.1.6	The <i>Bang-bang Heater</i> Problem with Less States . . . . .	97
5.1.7	The <i>Set Point Heater</i> Problem . . . . .	99
5.2	Context-based Modeling Simulation Setup . . . . .	102
5.2.1	EnergyPlus . . . . .	103
5.2.2	Simulation Diagram . . . . .	104
5.2.3	Building Site and Construction . . . . .	107
5.2.4	Parameters of the RC Model . . . . .	108
5.2.5	Model Performance Evaluation . . . . .	115
5.2.6	Simulation Files . . . . .	117
5.3	Summary . . . . .	117
<b>6</b>	<b>Results</b>	<b>119</b>
6.1	Reinforcement Learning Simulation Results . . . . .	119
6.1.1	The Bang-Bang Heater . . . . .	120
6.1.2	Simulation Results With the Uncomfortable State Redefined . . . . .	122
6.1.3	<i>Bang-bang Heater</i> Results Using Less States . . . . .	123
6.1.4	The <i>Set Point Heater</i> Heater . . . . .	124
6.2	Context-based Thermodynamic Modeling . . . . .	130
6.3	Summary . . . . .	134
<b>7</b>	<b>Conclusions and Future Work</b>	<b>135</b>
7.1	Reinforcement Learning for HVAC Control . . . . .	136
7.2	Context-based Thermodynamic Models . . . . .	137
7.3	Future Work . . . . .	139

<b>Bibliography</b>	<b>142</b>
<b>A Heat Transfer Through Building Constructions</b>	<b>167</b>
A.1 Heat Transfer Through a Material Layer . . . . .	167
A.1.1 Fourier’s Law of Heat Transfer . . . . .	167
A.1.2 Conservation of Energy . . . . .	168
A.1.3 Laplace Transform . . . . .	169
A.1.4 Transmission Matrix . . . . .	169
A.1.5 Electrical Analogy . . . . .	170
A.1.6 Model Reduction . . . . .	170
A.2 Multi-Layer Walls . . . . .	171
A.2.1 Modeling of a Multi-Layer Wall using a 3R2C Model . . . . .	172
A.2.2 Modeling of a Multi-Layer Wall using a 3R4C Model . . . . .	174
A.3 Simulation Results . . . . .	174
A.3.1 3R2C Wall Model . . . . .	174
A.3.2 3R4C Wall Model . . . . .	177



# List of Figures

1.1	The smart building concept map . . . . .	3
1.2	Picture of a green building . . . . .	9
1.3	The intelligence appraisal of a smart building . . . . .	11
1.4	The challenge for building automation . . . . .	13
1.5	Intelligent behavior . . . . .	13
1.6	<i>Context</i> of a school building . . . . .	15
2.1	PPD – Predicted Percentage Dissatisfied. . . . .	33
2.2	Thermostat interfaces . . . . .	35
2.3	A behavioral feedback scheme . . . . .	38
3.1	The BAS as an intelligent <i>agent</i> . . . . .	47
3.2	The agent interacting with the environment . . . . .	48
3.3	Mapping the continuous state into continuous action-value pairs (Wires) . . . . .	53
3.4	Wire fitting interpolation . . . . .	55
3.5	Interpolation example . . . . .	56
3.6	Wire fitted neural network training. . . . .	57
3.7	Updating the set of wires . . . . .	58
3.8	Reinforcement learning for HVAC control . . . . .	59
3.9	The <i>Bang-bang Heater</i> heater problem . . . . .	60
3.10	Selecting actions with the <i>Set Point Heater</i> . . . . .	62
3.11	Comfort penalty function . . . . .	63
4.1	A hybrid system example with the <i>Bang-bang Heater</i> . . . . .	72
4.2	Hysteresis behavior of the <i>Bang-bang Heater</i> . . . . .	73
4.3	Graphical representation of the heating system. . . . .	75
4.4	Hybrid time set . . . . .	76
4.5	Execution of the heating system . . . . .	78
4.6	Mode for heat transfer . . . . .	79
4.7	RC building envelope model . . . . .	81
4.8	Full-scale RC model of the TZ . . . . .	83
4.9	Application Example . . . . .	85
4.10	Natural ventilation macro-state . . . . .	87
5.1	Simulation of the occupant behavior . . . . .	91
5.2	Representation of comfort using a fuzzy set . . . . .	91

5.3	The HVAC and occupant's state with no automation . . . . .	93
5.4	New simulation of the occupant behavior . . . . .	98
5.5	Simulating the thermal zone with a PID controller . . . . .	101
5.6	Temperature response with the PID controller . . . . .	102
5.7	Box-shaped building with a single thermal zone . . . . .	103
5.8	EnergyPlus simulation engine . . . . .	105
5.9	Buildings Controls Virtual Test Bed simulation model . . . . .	106
5.10	EnergyPlus 24-hour simulation example . . . . .	107
5.11	Evolution of outdoor and indoor temperatures depending on shades .	112
5.12	The unit heater . . . . .	113
5.13	48-h simulation of solar gains . . . . .	114
5.14	Airflow rate with two openings (door and window) . . . . .	115
5.15	Airflow and temperature, depending on the opening factor of a window	116
6.1	Convergence of Q-values and average reward . . . . .	121
6.2	Policy application results . . . . .	122
6.3	Policy application results with occupant's comfort interaction redefined	124
6.4	Convergence of Q-values and average reward using less states . . . . .	125
6.5	Policy application results using less states . . . . .	127
6.6	Average reward daily with the <i>Set Point Heater</i> . . . . .	128
6.7	Policy execution results for the <i>Set Point Heater</i> . . . . .	128
6.8	Execution and simulation results . . . . .	132
6.9	Histogram of the error distribution for the example in Figure 6.8. . .	133
A.1	Heat conduction through a single material layer . . . . .	168
A.2	Thermal two-port associated with a construction layer. . . . .	170
A.3	Heat transfers and construction of a multi-layer wall. . . . .	171
A.4	3R2C Wall model. . . . .	172
A.5	Two-port network configurations . . . . .	173
A.6	3R4C Wall model. . . . .	174
A.7	TZ model using a 3R2C RC model for the envelope. . . . .	175
A.8	Temperature evolution for a multi-layer wall (3R2C) . . . . .	176
A.9	Histogram of errors (3R2C) . . . . .	177
A.10	Histogram of errors (3R4C) . . . . .	178
A.11	Temperature evolution for a multi-layer wall (3R4C) . . . . .	179

# List of Tables

1.1	The energy usage in commercial and residential buildings . . . . .	7
2.1	The predicted mean vote (PMV) sensation scale. . . . .	33
3.1	Q-values used to solve the <i>Bang-bang Heater</i> problem. . . . .	61
4.1	The domain of each <i>context</i> . . . . .	87
4.2	Guards associated with each of the edges represented in Figure 4.9. .	88
5.1	Initial simulation parameters used for the <i>Bang-bang Heater</i> . . . . .	96
5.2	Properties for constructive elements . . . . .	108
5.3	Convective heat transfer coefficients. . . . .	110
6.1	Average performance metrics <i>Bang-bang Heater</i> (1) . . . . .	120
6.2	Average performance metrics <i>Bang-bang Heater</i> (2) . . . . .	123
6.3	Average performance metrics <i>Bang-bang Heater</i> (3) . . . . .	126
6.4	Comfort vs. Heating Cost . . . . .	129
6.5	RC values for the full-scale thermodynamic model . . . . .	131
6.6	The hybrid time set and discrete trajectory of the running example. .	131



# Nomenclature

## Acronyms and Abbreviations

<b>AI</b>	Artificial Intelligence
<b>AmI</b>	Ambient Intelligence
<b>BCVTB</b>	Building Controls Virtual Test Bed
<b>BAS</b>	Building Automation System
<b>CO<sub>2</sub></b>	Carbon Dioxide
<b>DR</b>	Demand Response
<b>EIA</b>	U.S. Energy Information Administration
<b>FSM(s)</b>	Finite State Machine(s)
<b>GB(s)</b>	Green Building(s)
<b>GHG</b>	Greenhouse Gases
<b>HVAC</b>	Heating, Ventilation and Air Conditioning
<b>HMI(s)</b>	Human-Machine Interface(s)
<b>ICTs</b>	Information and Communication Technologies
<b>IEA</b>	International Energy Agency
<b>IoT</b>	Internet of Things
<b>KR</b>	Knowledge Representation
<b>LCS(s)</b>	Large Complex System
<b>LTI(s)</b>	Linear Time-invariant
<b>MAE</b>	Mean Absolute Error
<b>MAS</b>	Multi-Agent Systems
<b>MAX</b>	Maximum Absolute Error
<b>MDP</b>	Markov Decision Process
<b>ML</b>	Machine Learning
<b>MPC</b>	Model Predictive Control
<b>NLP</b>	Natural Language Processing
<b>NN(s)</b>	Neural Network(s)
<b>ODE(s)</b>	Ordinary Differential Equation(s)
<b>PID</b>	Proportional–Integral–Derivative
<b>PMV</b>	Predictive Mean Vote
<b>PPD</b>	Predicted Percentage of Dissatisfied
<b>RC</b>	Resistance-Capacitance
<b>RL</b>	Reinforcement Learning
<b>SB(s)</b>	Smart Building(s)
<b>SG(s)</b>	Smart Grid(s)
<b>TD</b>	Temporal Difference
<b>TZ(s)</b>	Thermal Zone(s)

<b>NZEB</b>	Nearly Zero-Energy Buildings
w.r.t.	With respect to
aka	Also known as
vs.	Versus
i.e.	From the Latin <i>id est</i> meaning "that is"
e.g.	From the Latin <i>exempli gratia</i> meaning "for example"
etc	From the Latin <i>et cetera</i> meaning "and other things"
et al.	From the Latin <i>et alia</i> meaning "and the others"

## List of Symbols

### Reinforcement Learning

<b>Symbol</b>	<b>Description</b>
$S$	Set of states of the environment
$U$	Utility function that measures the utility of a state
$A$	Set of actions the agent can execute
$P$	Set of probabilities associated with action outcomes
$R$	Reward function
$s_t$	State at time instant $t$
$s_{t+1}$	State in the next iteration step $t+1$
$a$	Action selected and executed
$\pi$	Policy function that maps each state to an action
$\pi^*$	Optimal policy that maps each state to the optimal action
$r$	Reward received
$\gamma$	Discount factor associated with current and future rewards
$Q$	Function that gives the utility of doing an action in a state
$q$	Utility value (q-value)
$q_{\text{Max}}$	Highest q-value
$a_{\text{Max}}$	Action associated with the highest q-value
$\delta$	Learning rate
$\varepsilon_{\text{Rate}}$	Exploration rate
$t_{\text{Comf}}$	Total amount of time the occupant was comfortable
$\Delta_{\text{min}}$	Minimum time interval the occupant was uncomfortable
$\Delta_{\text{max}}$	Maximum time interval the occupant was uncomfortable
$\Delta_{\text{mean}}$	Average time interval the occupant was uncomfortable

### Context-based Modeling

<b>Symbol</b>	<b>Description</b>
---------------	--------------------

$\mathcal{M}$	Context-based model
$x$	Continuous state
$l$	Discrete control mode/ <i>context</i>
$\rho = (l, x)$	Full state of the hybrid system
$Init$	Set of initial states
$L$	Set of contexts
$D$	Domain
$u$	Model inputs
$E$	Set of edges
$y$	Model outputs
$t$	Time variable
$Rst$	Reset map
$G$	Guard condition
$\tau_i$	Time instant $i$
$\tau$	Hybrid time set
$\chi$	Execution of a hybrid automaton
$\chi_{EP}$	Simulation execution
$I$	Time interval
$\chi_{Model}$	Model execution

$\mathbf{1}_A(x)$	Indicator function: $\mathbf{1}_A(x) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{if } x \notin A \end{cases}$
$ x $	Absolute value of $x \in \mathbb{R}$
$\ x\ $	L <sub>2</sub> -Norm of $x \in \mathbb{R}^n$ : $\ x\  = \sqrt{\left(\sum_{i=1}^n  x_i ^2\right)}$
$\ x\ _\infty$	L <sub>∞</sub> -Norm of $x \in \mathbb{R}^n$ : $\ x\ _\infty = \max_{i=1, \dots, n}  x_i $

## Building Physics and RC models

Symbol	Units	Description
$A, B, C, D$		Model matrices
$z$	m	Thickness
$\rho$	kg m <sup>-3</sup>	Density
$\rho_{air}$	kg m <sup>-3</sup>	Density of air
$Q$	W	Energy flux (heat gain)
$Q_o$	W	Heat lost through the building envelope
$Q_s$	W m <sup>-2</sup>	Solar gains (outer face of the building envelope)
$Q_{sw}$	W	Solar gains transmitted through building windows
$Q_i$	W	Energy flux generated by each occupant
$Q_h$	W	Heat generated by heating equipment

$f_{sw}$	W	Heat transmission function ( $Q_{sw}$ to zone air)
$T_a$	$^{\circ}\text{C}$	Outdoor ambient temperature
$T_{in}$	$^{\circ}\text{C}$	Indoor zone air temperature
$T_d$	$^{\circ}\text{C}$	Desired temperature setpoint
$T_g$	$^{\circ}\text{C}$	Ground temperature
$T_h$	$^{\circ}\text{C}$	Temperature of the heater
$h_c$	$\text{W m}^{-2} \text{K}^{-1}$	Convective heat transfer coefficient
$\lambda$	$\text{W m}^{-1} \text{K}^{-1}$	Thermal conductivity
$R_{ext}$	$\text{K W}^{-1}$	Exterior surface convective resistance
$R_{int}$	$\text{K W}^{-1}$	Interior surface convective resistance
WS		Set of opening factor for window shades
$R_W$	$\text{K W}^{-1}$	Thermal resistance of windows
Heat		Set of heater states
$R_{WS}$	$\text{K W}^{-1}$	Thermal resistance of window shades
$\dot{M}_v$		Set of discretized ventilation levels
$R_{ih}$	$\text{K W}^{-1}$	Thermal resistance heater/interior air
$\dot{m}_v$	$\text{kg s}^{-1}$	Ventilation air flow rate
$R_{vent}$	$\text{K W}^{-1}$	Resistance for natural ventilation
$o$		Number of occupant in the thermal zone
$C_h$	$\text{J K}^{-1}$	Heat capacitance of the space heater
$C_p$	$\text{J kg}^{-1} \text{K}^{-1}$	Specific heat capacity
$C_z$	$\text{J K}^{-1}$	Thermal capacitance of zone air
$C_{air}$	$\text{kJ kg}^{-1} \text{K}^{-1}$	Specific heat capacity of air
W		Set of opening factor for windows
D		Set of opening factor for doors
$wof_i$		Opening factor state of window $i$
$d_i$		Opening factor state of door $i$
$ws_i$		Opening factor state of the shades in window $i$
$A_{wind}$	$\text{m}^2$	Window area off the TZ
$h$		The state of the heater
$A_e$	$\text{m}^2$	Area affected by $Q_s$
$A_{roof}$	$\text{m}^2$	Area of the roof

## Glossary

<b>Ambient Intelligence</b>	A term widely used to signify a vision in which environments support the people who inhabit them by incorporating data acquisition, computation, intelligence and behavior to everyday objects and spaces in an interconnected and unobtrusive way.
<b>Context</b>	Context is a structure or a frame of reference that defines which knowledge should be considered, what are the conditions of activation and limits of validity and when to use it at a given time.
<b>Effectors</b>	Also known as actuators in robotics, are components capable of moving a mechanical system such as a vent position, window or door.
<b>Green Building</b>	A combination of construction and a set of processes that are environmental friendly and resource-efficient throughout the entire building's life-cycle.
<b>Greenhouse Gases</b>	Gases in the atmosphere that cause the greenhouse effect, responsible for global warming and climate change.
<b>Hybrid System</b>	A system that combines <i>time-driven</i> and <i>event-driven</i> dynamics. The former is represented by differential equations, characterizing the behavior of continuous-time variables, while the latter, characterizing the switching conditions between different discrete modes, is described through various frameworks used for discrete event systems.
<b>Large complex system</b>	A system composed of many components which may interact with each other creating complex behaviors that are extremely hard to model and predict.
<b>Markov Decision Process</b>	Provides a mathematical framework for modeling sequential decision making in situations where decision outcomes are uncertain, and decisions are under the control of a decision maker with the objective of maximizing a utility function.

<b>Smart Building</b>	A building supported by technology that includes multiple systems and automation to control its operation.
<b>Smart City</b>	A urban development vision supported that uses information and communication technology to connect and exchange information between different city assets (e.g., buildings, cars, power, water and gas distribution networks) with the goal of improving overall efficiency of the city and the living conditions of its citizens.
<b>Smart Grid</b>	An electrical grid supported by information technology to interconnect the different network stakeholders and systems. Information is used in the smart grid to synchronize power production and demand in the network.
<b>Temperature set point</b>	Represents the desired or target temperature value, desired for a thermal zone.
<b>Ubiquitous computing</b>	A concept strongly associated with the Internet of Things (IoT), where many surrounding devices and systems can have computer processing capabilities, be interconnected, have advanced human-machine interfaces and weaver these capabilities into the fabric of everyday life until they are indistinguishable.

# 1

## Introduction

*“A house is a machine for living in.” – Le Corbusier (Architect).*

Over the past few decades, many countries have committed to adopting policies and initiatives explicitly designed to address climate change and define paths towards energetic sustainability. These initiatives have promoted efforts to find cost-effective ways to reduce or avoid the production of greenhouse gases (GHG) and dependence on fossil fuels [1]. To meet these objectives, some of the policy directives and guidelines are aimed at promoting energy efficiency in buildings, since indicators show that there is a high cost-effective potential for energy savings in this sector [2, 3]. Modern buildings are very *large complex systems* (LCSs) that are capable of supplying a variety of final services for occupants [4, 5]. It is through the process of delivering these services that a significant fraction of the energy used is wasted, among other reasons, due to inefficient building technologies [6]. To promote energy efficiency, buildings need to be built or retrofitted with well-applied sustainability solutions and incentives for energy savings.

Advances in technology are opening doors for entire new concepts and applications and energy can ultimately be saved by as much as 70% by the year 2030 in

new buildings with more efficient technologies [6]. *Smart buildings* (SBs)<sup>1</sup> [7–10] have been hailed as a solution to increase energy efficiency in the building sector. These buildings, represented in the concept map shown in Figure 1.1, are enabled by state of the art technologies and use automation and machine-to-machine communication to manage the complex concatenations of the internal building systems and equipment, and other external systems such as the electrical *smart grid* (SG) [11]. SBs are intended to deliver several building services over their lifecycle that promote a comfortable environment for their occupants, such as illumination, thermal comfort and air quality, while ensuring the efficient use of building resources [12, 13]. Therefore, there is a growing interest in developing and integrating technologies that make SBs more comfortable, economical, safe and efficient.

SB operation includes a *building automation system* (BAS) with advanced controls and algorithms to perform the real-time optimization of its energy-consuming systems. Building networks, supported by open standards such as BACnet, Modbus, LonMark and ZigBee [14–17], among other sensor network technologies, make building environments more observable for computation [18]. These networks provide access to fine-grained spatial-temporal resolution measurements of various variables from the building environment such as temperatures, CO<sub>2</sub> levels, humidity, movement, air velocity, occupancy, lighting conditions, energy usage, states of doors and windows, meteorological data and damper positions. Through online platforms, moreover, it is also possible to obtain energy pricing information, access weather forecasts, occupancy schedules and other information relevant for decision support and performance evaluation. Therefore, the opportunity to use all this information in various building applications becomes evident. However, trying to use all the available information to optimize the performance of the building makes optimization problems computationally intensive and extremely difficult to implement in real buildings.

In the beginning of the 90s, following Weiser’s publication on *ubiquitous computing* [19], a significant part of the research that has been done on smart environments falls within the *ambient intelligence* (AmI) research community [20]. AmI is a term widely used to convey a vision in which environments seamlessly adjust themselves to support the activities of occupants by incorporating data acquisition, computation and intelligent behavior into spaces and everyday objects. To create comfortable environments, SBs are expected to provide an environment with AmI that anticipates

---

<sup>1</sup>Throughout this thesis, we will use the term Smart Buildings, although part of the literature uses the term Intelligent Buildings.

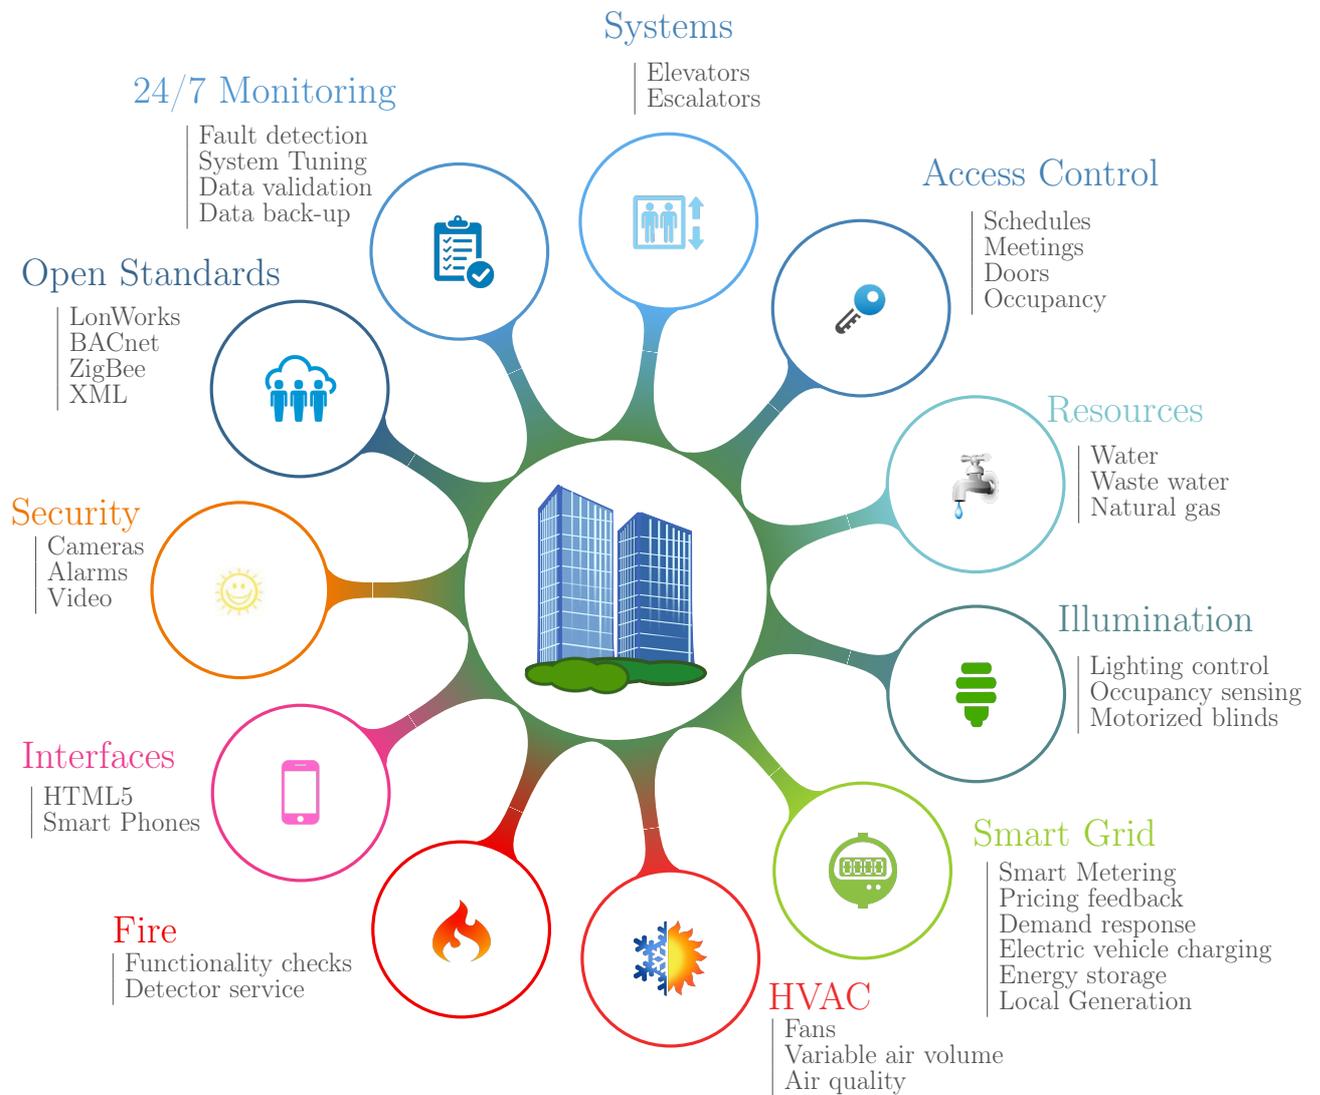


Figure 1.1: Concept map of a smart building as an integrated system of software hardware and services.

the needs of the occupants by learning their time-varying preferences and habits and respond in a timely and user friendly way [21–24]. However, most of these expectations are not trivial to implement and represent a massive challenge for system design and integration, knowledge representation (KR), machine learning (ML), con-

trol, and optimization. Moreover, given the current state of the art, it is not yet clear how to design and implement smart applications with the building's physical structure, services and processes. The integration of services and data from multiple sources, as well as interoperability between building systems, is a serious challenge for LCSs [25]. More research has to be carried out on these topics to support the expectations pertaining to SBs and AmI.

This dissertation has been motivated by the rich vein of research and development that exists for energy conservation in buildings, building systems and smart environments. Therefore, this study is another contribution in the domain of SBs. Research for this thesis was developed with a particular focus on ML techniques [26] that can be deployed in the BAS to perform automated tasks in the building. The focus of the study evolved in two distinct phases. In the initial research application, *reinforcement learning* (RL) [27–30] was used to optimize the operation of the building's heating, ventilation and air conditioning (HVAC) system, considering feedback information from the occupant and a cost function for heating. Simulation results with RL showed that the BAS was capable of minimizing the energy used for heating while maximizing comfort. However, it was also concluded that further optimization of the HVAC application would require adding additional information concerning the building structure and operation to the optimization problem. This requirement not only creates a problem for KR, but also makes the BAS learning process computationally expensive. Moreover, it was not clear how a RL HVAC controller could integrate with other SB systems and contribute to the AmI vision. Therefore, the research was guided towards finding an appropriate modeling paradigm for KR and organization using *context-based modeling*. The research proceeded with a clear focus on the thermodynamic behavior of building spaces due to its potential for *model predictive control* (MPC) and energy efficiency [31].

This chapter will proceed as follows. It will begin by presenting the main motivation for focusing on green and smart buildings. This includes a brief overview of important events and the global concern regarding environmental issues and energy efficiency. While addressing this concern, the chapter will discuss the contribution that buildings have on the global energy bill and the motivation for focusing on the HVAC system. Two important problems are presented in this domain, the contributions to solve each problem, a list of publications, and the outline of this thesis.

## 1.1 Motivation

As countries become more developed and populations increase, the demand for services such as comfort, education, leisure, and health also rises, and populations increase their energy usage in ways that are not always sustainable. As a consequence, the global energy demand is raising serious issues for energy supply and environmental impact mitigation. The world is far from achieving energy systems that are environmentally sustainable and the global energy demand continues to grow rapidly. Corroborating this fact, the U.S. Energy Information Administration (EIA) projects that world energy consumption will grow by 56% between 2010 and 2040 (from 524 quadrillion British thermal units (Btu) to 820 quadrillion Btu), mainly driven by emerging economies, such as India and China [32]. Fossil fuels such as coal, oil and gas will continue to supply nearly 80% of global energy use through 2040, led by natural gas, as the global supply of tight gas, shale gas, and coalbed methane increases. It is the consumption of these fossil fuels, the clearing of forests, agricultural practices, and other activities, that contribute to the amount of GHG released into the atmosphere [33].

Climate change is a real problem and human-induced GHG emissions are the primary cause of the global warming that has been observed over the past fifty years [33]. Therefore, this research has been intended, since the inception of this thesis, to become another contribution to help solve this very difficult and important problem. This motivation is further reinforced by the fact that Portugal imports 71.5%<sup>1</sup> of its energy [34], with these imports having a huge negative impact on economic growth and contributing to the nation's credit crisis.

Energy efficiency is the cornerstone of this thesis. The promotion of energy efficiency improves resource management and reduces energy demand, as well as its environmental impact. In this regard, this thesis addresses the buildings sector and, within that sector, the HVAC system. HVAC includes the equipment, distribution network and terminals used either collectively or individually to provide fresh filtered air, heating, cooling and humidity control in a building. Considering that the **HVAC systems are among the most energy-demanding systems integrated in buildings**, it is of the utmost importance that they are properly maintained and used in the most energy-efficient and cost-effective way.

---

<sup>1</sup>Values of 2012

### 1.1.1 Energy and Buildings

Over the past few years, the energy problem has driven different research areas and fields of interest, and a significant part of the research has focused on creating efficient systems for transforming energy. This widespread concern is being addressed in all energy systems at different levels of detail, extending from smart cities and buildings to specific systems, such as heating and lighting. People living in developed countries spend more than 90% of their time indoors and buildings are responsible for approximately 41% of primary energy usage, in many cases surpassing other large sectors, such as industry and transport [35–37]. According to the International Energy Agency (IEA) [38], in 2005, the final energy use in the building sector (Services and Domestic) was 108 EJ. Between 1990 and 2005, global GHG emissions increased to 21.2 Gt CO<sub>2</sub>, and 33% are associated with this sector.

The problem of energy conservation in buildings is a multidimensional one. It can be addressed in terms of building types, building sizes and building services. There are several types of buildings with distinct energy usage patterns, such as residential buildings, business stores, restaurants, hotels, hospitals, museums, shopping malls, heated indoor pools, large supermarkets, etc. As people spend more time in these buildings, the demand increases for more comfort and better building services, such as air quality, thermal comfort, lighting, domestic hot water, food preparation and communications. Tables 1.1a and 1.1b show the breakdown of the energy end-use in the commercial building sector for the U.S. and China, and in the residential sector for the U.S., China and Portugal, respectively. The residential sector in Portugal, with approximately 3.9 million households, uses 17.7% of the total national final energy, representing 30% of the electricity consumed [39]. In both residential and commercial buildings, space heating represents a substantial amount of energy in all countries (less in Portugal, due to its moderate weather), followed by water heating and other uses, primarily electric appliances in the U.S. (artificial illumination, elevators, office and kitchen equipment, etc.), which are also expected to increase in use over time in Portugal and China [40,41]. In service buildings, electricity is by far the most widely used energy commodity with a global share of 47% (in 2005). The use of electricity has increased by 73% since 1990, which has been the main factor driving the global increase in energy usage in this sector [38].

Among all building services, HVAC systems are the most significant energy-consuming services. In the U.S. alone, **one-sixth of the electricity consumed goes to cool buildings at an annual power cost of \$40 billion**, and forecasts

	China	U.S.		China	U.S.	Portugal
			appliances	21.0%		
space heating	45%	12%	space heating	32.0%	29.0%	10.7%
space cooling	14%	10%	space cooling		11.0%	0.7%
water heating	22%	7%	water heating	27.0%	11.0%	27.6%
cooking		2%	kitchen	7.0%	3.0%	40.0%
lighting		21%	lighting	9.0%	11.0%	6.1%
other uses	19%	48%	other uses	4.0%	35.0%	14.9%

(a) Commercial buildings. (b) Residential buildings.

Table 1.1: Breakdown of commercial and residential building sector energy usage in the U.S. (2005), China (2000), and Portugal (2010) (Source: IPCC [42]), INE/DGEG [39]).

show that this cost will continue to increase [37, 43]. Space heating and cooling can use up to 40% of the final energy in residential buildings, and 20% in commercial buildings [37, 42, 44]. Contributions in this domain can have huge economic implications at a global level and can contribute to mitigating energy usage and GHG emissions.

### 1.1.2 Green Buildings and Energy Efficiency

While traditional buildings have been viewed as a relatively static sector of the economy, experiencing relatively little change in technology or resource consumption patterns, *green buildings* (GBs) use key resources such as energy, water, materials and land more efficiently. In sustainable buildings, energy conservation plays a pivotal role. GBs provide cost and financial benefits, as compared to conventional buildings that are just built to code [45]. These benefits include energy and water savings, reduced waste, improved indoor environmental quality, greater employee comfort and productivity, reduced employee health costs and lower operations and maintenance costs. Savings associated with energy, water and waste can be predicted with reasonable precision after being measured and monitored over time. In contrast, productivity and health gains are much less precisely understood and far harder to accurately predict. Therefore, some benefits are more tangible than others.

To minimize the subjective nature of what exact combination of attributes defines GBs, sustainable building design tools and rating systems are currently being

used for building design and performance evaluation [46]. These rating systems and design tools promote the continuous process of systematically evaluating the performance and effectiveness of buildings in several aspects, including cost-effectiveness, functionality, productivity and sustainability. This benchmarking process promotes progress towards designing and operating the best buildings for their occupants, while still complying with national policies and regulations for the building sector.

The main goal of energy policies for the building sector is the promotion of energy efficiency. This goal improves the management of resources and reduces energy use, as well as its environmental impact. However, minimizing the negative impacts of buildings is a complex multidisciplinary problem that is currently being addressed by several different research fields. Since the building sector is very heterogeneous, any strategy for energy efficiency must be carefully adjusted to the type of building, climate characteristics, occupancy, and other parameters that influence the use of energy. There is no single solution to the energy efficiency problem. Taking a long-term perspective of sustainable development, different strategies need to be considered and deployed to address this problem.

An important step towards energy efficiency is to improve the buildings that already exist. For these buildings, strategies for energy frugality include retrofitting inefficient parts of the building and adding solutions adapted to the building's energy usage profile, such as using solar thermal energy for heating and cooling. However, more can be accomplished with new buildings. To establish new sustainable buildings, the approach usually taken includes following a set of rules and general methods that influence the siting, architecture, selection of building materials and systems incorporated in the building. As a general rule, to save energy, building constructions should promote an efficient interaction between the interior and exterior environments in such a way that promotes maximal environmental comfort levels in every season of the year, taking advantage of the best qualities of the surrounding environment. Environmental conditions that humans consider comfortable are mainly associated with air quality, humidity, lighting, sound and, above all, temperature [47]. These conditions are very important in terms of productivity and health, and the financial impacts on productivity due to poor indoor environmental conditions are well known [48–50]. Therefore, buildings should be built to operate in such a manner that they only bring into the indoor environment qualities of the outdoor climate that are desired by occupants, while minimizing the amount of energy required to guarantee comfort conditions that keep occupants productive and happy.

Energy efficiency can be achieved if proper considerations are taken regarding building design by using, for example, passive bio-climatic strategies, which are approaches to building design that use the architecture to minimize the use of energy [51]. To accomplish this requirement, *passive design* can be used as a design strategy to optimize the interaction of the building architecture with the local micro-climate. This strategy attempts to control comfort without using purchased energy by maximizing the use of free solar energy for heating and lighting, as well as *natural ventilation* for cooling. Figure 1.2 shows the architectural model for a new bio-climatic office building in Lyon, designed by Nicolas Laisné Associés and scheduled to be complete in 2018.



Figure 1.2: Plans for a new green building in Lyon using bio-climatic strategies, with green-layered outer facade (reprinted from <http://laisneroussel.com/fr>).

For energy savings, priority should be given to passive building design strategies [52]. However, the ultimate passive design vision is very hard to achieve [53]. Therefore, buildings must rely on active strategies that use purchased energy and mechanical systems to accomplish comfort and energy efficiency through the lifecycle operation of the building. In some applications, active-passive strategies can be combined for environmental control [54,55]. However, building environments are dynamic and the integration of both strategies is a complex multifaceted process. Further re-

search is needed to provide a detailed description of the building environment, which is required to deploy more intelligent and efficient strategies for automation and post-construction management [56]. Indicators show that a significant part of this research will be integrated in the development of AmI and SBs [57–60].

### **Nearly Zero-Energy Buildings**

Following technology developments in information and communication technologies (ICTs), GBs are also currently one of the most important foundation blocks of the *smart city* urban development vision. Integrated with SGs, there is a quantitative ongoing challenge towards creating *nearly zero-energy buildings* (NZEB), with zero net annual energy usage. The most ambitious part of this challenge is the idea that NZEB can meet all their energy requirements from low-cost, locally available energy sources (with a special focus on using renewable energy sources) such as micro-turbines, fuel cells, photo-voltaic panels and diesel generators. NZEB can also include energy storage devices (e.g. flywheels, batteries and thermal storage) and controllable loads such as electric vehicles, elevators and HVAC systems. These buildings are expected to be integrated with the SG with the capability for automated *demand response* (DR), which is a concept where the building is capable of automatically changing its electricity consumption patterns in response to time-varying electricity rates or incentive programs. To achieve NZEB, it is crucial to perform a correct energy management between supply and demand to minimize financial costs and maximize the use of available energy. Current buildings lack necessary information systems for energy analysis and DR control strategies. New active strategies are necessary for future GBs to act as a coordinated cluster of systems, taking into account DR signals, building-integrated energy storage, and availability of renewable energy, in order to manage in real-time controllable loads such as the HVAC system [61].

#### **1.1.3 Smart Buildings**

The term SB was first used for buildings in the United States at the beginning of the 1980s and, since then, the concept has gained traction as a result of technological developments in several different fields, including ICTs, computer science, artificial intelligence (AI), robotics, control, and building automation standards. Within the past few decades, a substantial amount of literature has been generated in this domain. Wong and Wang [10] have presented a comprehensive review of existing research on

intelligent buildings, with several references in which the broad concept of SBs is discussed. To formalize the concept, Wong et al. [62, 63] present a method to evaluate and benchmark the intelligent performance of a SB using a set of indicators for intelligence with an analytical model for computing a *building intelligence score*. The authors take into account the interrelationship, represented in Figure 1.3, between the intelligent attributes and their operational goals and benefits. For buildings, the

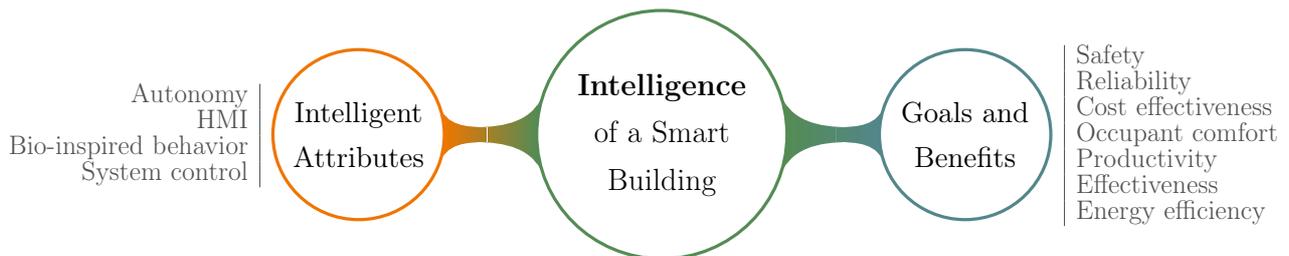


Figure 1.3: The intelligence appraisal of a smart building takes into account the interrelationship between intelligent attributes and operational goals.

term “smart”, synonymous with “intelligent”, has a functional definition: Smartness is typically associated with the automation and integration of different building systems in order to operate in ways that provide effective, responsive, and supportive building environments, within which organizations can meet their performance objectives [64]. Intelligent attributes include autonomy, controllability of complex system dynamics, human-machine interaction and bio-inspired behavior. In this sense, SBs extend the behavior of normal building systems (HVAC, lighting, local power generation, energy storage, etc.) beyond simple automation strategies to promote comfortable environments while managing building resources, such as water, power, and natural gas, with high efficiency and minimum waste. Goals and benefits include increased safety and reliability, lower costs, enhanced cost and operational effectiveness of building operations, and improved occupant comfort, productivity, and energy efficiency. These benefits can have significant impacts throughout the lifecycle of the building. For example, for operational effectiveness, SBs are expected to perform automatic evaluation and diagnosis of building systems using AI methods for continuous commissioning. Supporting this research, Louise Travé-Massuyès [65] exemplifies how different theories in the AI and Control literature can be integrated to provide better diagnostic solutions and to achieve improved fault management in different environments. Considering the fact that at least 10% of the energy wasted in buildings is due to excessive run time and problems in the HVAC equipment and controls, better

diagnostic solutions can have a significant impact in energy savings [66]. In another example, Ellis and Mathews [67] argued **that proper control of the heating and cooling demand in a building, using an integrated system design approach, can result in up to 70% energy savings for the HVAC system**. To accomplish these savings, the BAS is expected to “smartly” manage available building parameters and energy systems, e.g. by storing thermal energy locally and shifting energy demand to off-peak time periods when utility rates are lower; regulating natural lighting with shading devices while reducing glare and overheating; and managing temperature set points expressed in terms of space air temperatures in different conditioned spaces. The proposal of new efficient methods for building automation is therefore a key issue for research in energy and buildings.

### **Ambient Intelligence**

The vision for AmI envisages that building services should be managed in an unobtrusive and transparent way, contributing to the wellbeing of their occupants without intruding on daily activities. Therefore, minimizing energy waste and operating costs cannot be achieved at the expense of occupants’ comfort, productivity or health. Building environmental conditions should be regulated within an optimization space that is compliant with the comfort of its occupants i.e. **the BAS must avoid taking actions that will, with high probability, lead to counteractions by those occupants**. Implementing this vision requires addressing a vast set of challenges posed by building environments in a multidisciplinary effort to develop efficient ways for managing local energy generation, learning by interacting with occupants, modeling and controlling building systems, characterizing the thermal behavior of building spaces, and exchanging information with other systems outside the building, such as other SBs and SGs.

The problem of energy conservation in SBs is currently a popular research topic in energy management [68]. Although researchers have been working on intelligent control systems for energy and comfort management in SBs for over a decade, there is still a lot of work to be done in this domain [13]. The automated management and optimization of resources in a building environment faces many challenges. In most situations, the lack of observability, due to limited sensing capabilities, raises uncertainty about the building environment. There is a compromise, depicted in Figure 1.4, between the capability to manage and optimize building resources, with a set of constraints defined by e.g. occupant comfort requirements, and the number

of environmental variables can be observed and controlled.

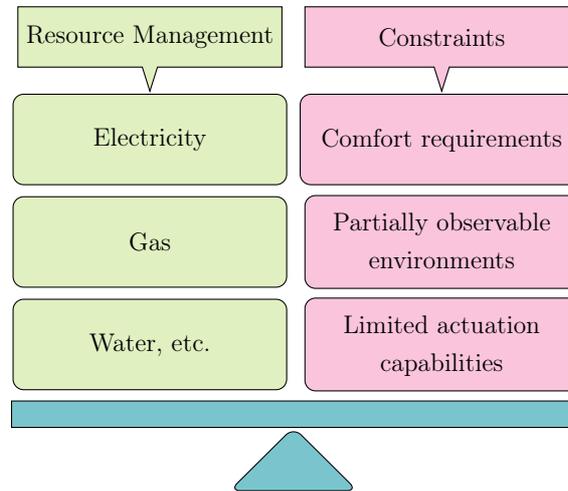


Figure 1.4: Compromise between the capability to manage and optimize building resources with a set of constraints for building automation.

This dissertation is guided by the vision of having a BAS with rational actions to proficiently manage the operation of the buildings through a thinking and learning process, mimicking human cognition. Therefore, the BAS must extract information from the environment, map that information into *actionable knowledge*, and execute intelligent behavior based on that knowledge, as represented in the sequential processing diagram illustrated in Figure 1.5 [68]. For energy efficient buildings, knowledge includes a set of rules and models that are useful to minimize energy waste. A SB should be able to use this knowledge in order to “understand” and predict its environment, including where and how energy is being used. **The smart BAS should acquire data, include feedback learning, and find ways to adapt in order to continuously improve its performance.** Each processing step represents a huge set of challenges that are not easy to address with current buildings and technology. Therefore, building intelligence is still an open problem for research.

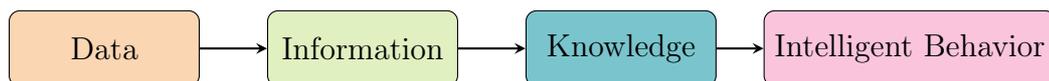


Figure 1.5: Sequential processing for intelligent behavior (adapted from [68]).

### Context-based Modeling

It is not easy to predict and quantify the effects that SBs may have on energy savings. SBs can have multiple spaces, occupants, human-machine interfaces (HMI), distributed systems, sensors, and a set of environmental variables (observed and unobserved) of a significant size that require controlling and/or monitoring. Despite all the available data, if no efforts are made on data management, data can produce little or no meaningful information in what is known as the *data-rich but information-poor syndrome* [69]. To meet the information expectations placed on building data, the process of formalizing empirical data into parsimonious theorems and principles, and the conceptual modeling of various building systems, must be part of the engineering process to describe the general knowledge of each part of the building. As stated by Weng and Agarwal, “Areas of SB research, such as modeling and prediction of building operations, can be used to augment and improve the control over a building” [7]. All building variables represent a huge amount of information and currently it is the facility managers responsibility to interpret high-level data and act upon the information available. This is not a trivial task in data-rich environments. Building managers need decision support systems that are ubiquitous advisors with automatic diagnostics capable of selecting relevant building information for decision support. To facilitate human intervention in building operations, information should be filtered and presented to the facility manager in a human-friendly way [70, 71].

Organizing information in LCSs requires using appropriate modeling paradigms. *Context-awareness* has been presented as one of these paradigms where the identification and adjustment of behavior according to specific conditions are primitive concepts [72–81]. Coarsely speaking, **context is a structure or a frame of reference which can be used as a mechanism to manage, organize or reason about knowledge** [82, 83]. It can define which knowledge (or model) should be considered at a given time, what the conditions of activation are and limits of validity under which it applies. Context-awareness provides the means to partition the operation of a complex systems such as smart buildings into “scenarios” (or *situations* [84]) where knowledge, strategies, parameters and objectives are organized. As an example of partitioning, depicted in Figure 1.6, consider how occupancy (and thus the use of energy) in a school building, depends on how the academic year is organized into vacations, exams, holidays and instructional days. Occupancy and energy usage, have different distinct profiles in each of these cases. Many other building variables, parameters, objectives and strategies, also depend on the season, weather conditions,

cost of energy, location, and other previously known facts such as meetings or other special events.

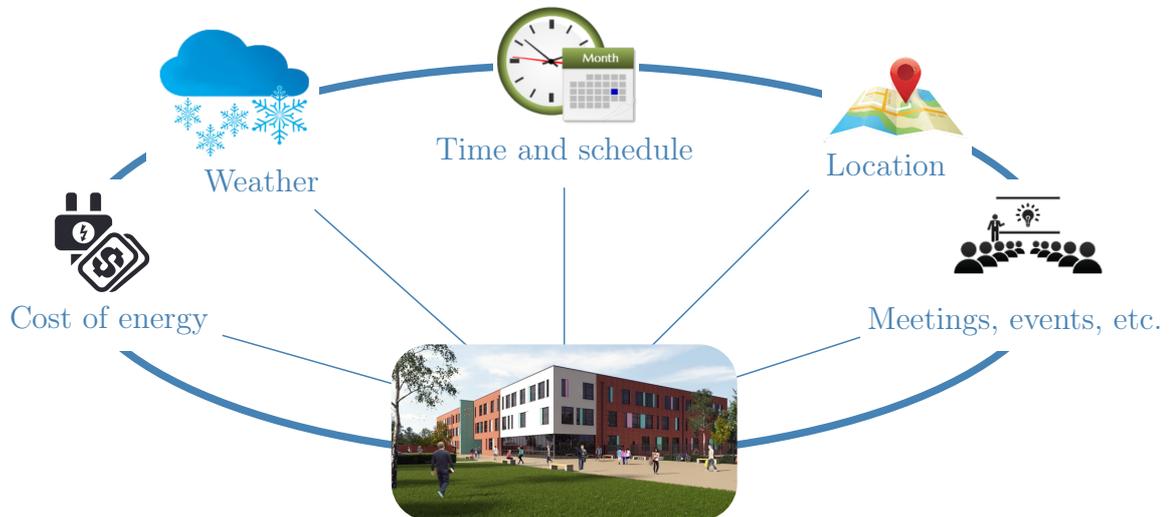


Figure 1.6: *Context* of a school building depends on location and changes depending on various factors such as weather conditions, special events and activities, cost of energy and holidays.

Context has always played an important role in human intelligence. The awareness of *context* about the environment, discussion, or problem in hand, allows many important aspects of human interaction to remain implicit [75]. Contexts act like adjustable filters creating a knowledge frame that is shared by all interlocutors thus minimizing the amount of information that needs to be exchanged for an effective communication. The interlocutors intuitively know, at each step, which knowledge pieces must be taken into account explicitly (contextualized knowledge) and which pieces are not directly necessary or already shared (contextual knowledge).

In the past few years context-awareness has been used in several different applications. A survey of the literature dealing directly and explicitly with context in different domains has been presented by Brézillon [72]. Context-awareness has been used for natural language processing (NLP), databases, ontologies, communication, electronic documentation, vision, AI and AmI [85, 86]. However, not much can be found on context-awareness and SBs, leaving a gap in the literature. There are several potential applications where context-awareness could be useful for building operations. For example, when identifying the thermodynamic model of a building *thermal zone* (TZ), considering a specific *context* associated with the TZ being oc-

cupied by a single occupant with all are windows *closed*, occupancy and ventilation rate are variables that can directly be inferred through *context*. These variables, as contextual knowledge, can be replaced by constants in that *context*, thus minimizing the number of explanatory variables in model equations. Therefore, when aiming at more simplified models, context-based modeling can be a natural way of including previous knowledge in the process of modeling systems and variables from the observed environment.

## 1.2 Research Questions and Objectives

In the building environment, the attitudes and preferences of occupants have a significant impact on the efficient use of energy resources. When taking into account the fact that occupants tend to forget to adjust the HVAC appropriately, and in many spaces, the conditioning requirements are not adjusted to the occupancy of those spaces, the result is energy being unnecessarily wasted. In this regard, developing a self-adaptive energy management system capable of minimizing the energy used for heating and cooling, by correctly scheduling the HVAC unit activity for occupancy and non-occupancy periods, becomes a key factor for promoting the efficient use of the HVAC system [87]. **Ideally, a multi-objective optimal supervisory controller for such a system would take into account activity schedules, occupancy patterns, the individual preference of each tenant, the cost of energy, weather predictions, and the thermodynamic behavior of building spaces, among other information that can be used for decision support. However, this is not a simple optimization task. HVAC is a multivariate nonlinear time-variant system, so optimizing the operation of this system is a complex optimization problem with numerous constraints [88,89].** Moreover, some of these constraints, such as the thermal sensation of an occupant, occupancy, and other environmental conditions, may not be directly observable and may need to be inferred, during execution time, from available data. Consequently, the optimization of energy and comfort management in buildings remains an open challenge for real-time computing, and more research is required in this domain [13].

Currently, there are still many aspects that need to be addressed in AI research in order to create intelligent systems [90]. Subsequently, most of this research will extend to SBs. This presents challenges and opportunities for research and an opportunity to foster synergy between different research areas. In the current state

of the art, there is no generally adopted approach for creating intelligent HVAC systems. Coordinated research efforts are needed to develop AI techniques for automatic HVAC control, and predictive systems capable of forecasting the demand for energy, comfort conditions, and occupancy time expectancy [13]. **To address these requirements, this dissertation starts by pursuing the goal of developing a supervisory controller capable of learning from the building environment, dealing with its uncertainty, and proficiently controlling the amount of energy used to operate the heating system, while still keeping occupants comfortable.** The research challenge includes scaling the BAS intelligent behavior to integrate additional information, such as the state of doors, windows and blinds. However, this integration task is too onerous for our HVAC controller. Integrating all of the available building information that could be useful for decision support makes the problem of energy conservation and comfort management computationally expensive and extremely difficult to implement in real buildings [68]. Therefore, **this research proceeds towards finding suitable paradigms for KR in SBs, with a particular focus on how to represent the thermodynamic behavior of building spaces.** Although there is some literature on ontology-based KR for AmI and buildings [91–93], not much can be found on modeling and reasoning about building dynamic systems. Moreover, the evolution of temperatures in building spaces has double continuous and discrete dynamics i.e. temperatures undergo abrupt changes of dynamics upon certain changes in the building environment. To this researcher’s knowledge, there are no thermodynamic models in the literature that are capable of describing this continuous and discrete dynamic behavior. Therefore, this dissertation explores a modeling paradigm for creating thermodynamic models with multiple modes of operation, capable of representing the thermodynamic behavior of building spaces in different contexts. In the following sections, each research problem is explained in more detail, including the methods used to solve them.

### 1.2.1 Automatic HVAC Optimization

In the complex building environment, building spaces with varying thermal temperatures can be subdivided based on occupancy-associated information and environmental conditions (thermal zoning) [94, 95]. These spaces, defined as TZs, are passive systems in the HVAC energy chain - exterior sources of energy are required for temperature regulation. The useful heat that must be added or extracted from each TZ is directly associated with the desired temperature set points and linked to the

chain of energy flowing through the HVAC network [95]. Demand for HVAC varies between TZs and the control algorithms need to be designed for each TZ based on predicted energy demands, directly influenced by the relationships between different configurations of floor area, office space layouts, occupancy, solar gains, state of doors, windows, electric equipment, lighting, etc. **To save energy, a smart HVAC controller must intelligently optimize temperature set points in different parts of the building according to the conditioning requirements of each TZ.** However, the information necessary to search for the optimal temperature profile, such as occupancy schedules and comfort requirements, is either not available or not sufficiently accurate to represent the best energy-efficient profiles for thermal demand. Consequently, building spaces are heated or cooled even if not needed by the occupants, and energy is wasted with incorrect thermostat settings.

Currently, in many building spaces, optimizing the operation of the HVAC includes various strategies, such as setting the operation of the system to a *low-power* state according to a certain pre-programmed schedule (setback control) [96]. This means that comfort requirements are not guaranteed in unscheduled hours when there is low occupancy, such as late classes or meetings. On the other hand, during normal operating hours, energy is unnecessarily wasted in many situations [97]. For instance, occupants tend to forget the HVAC *on* during non-occupancy periods, and heating and cooling loads are set to guarantee certain fixed set points instead of being intelligently exploited to save energy. This provides an opportunity to use ML to further optimize the HVAC. By learning the occupants' schedule, for example, the BAS can let the zone temperature "drift" without invoking heating or cooling a few minutes before lunch time - comfort conditions are maintained by thermal inertia. Energy can be saved during non-occupancy periods with no discernible change in thermal comfort, if the temperature set point trajectory is then set to guarantee a comfortable temperature in the TZ at predicted times of arrival. Therefore, to go beyond simple automation strategies, an intelligent HVAC controller must learn by observing the environment and find optimal strategies to schedule thermostat settings according to occupancy sensing and prediction.

Simple thermostats do not guarantee energy savings. In fact, in current buildings, at least 5-15 % of energy is wasted due to the course-grained, manual configuration of thermostats by occupants [98]. Therefore, learning the occupants' thermal sensation has a great potential for HVAC comfort control and energy savings [99]. For automatic HVAC control, calculations of human thermal comfort have historically

been based on predefined models that predict the mean thermal sensation of the occupants [100]. However, using fixed predefined comfort models may not always be the best solution, because comfort depends on several processes, such as physiological or even psychological factors, as well as on several circumstances, such as location, activity and season of the year. Different occupants can have different comfort preferences and the occupants' comfort state, associated with activity level and clothing, cannot be measured by sensors. Energy savings can be accomplished if the smart controller is capable of learning the occupants' comfort preferences by observing behavior and performing according to those preferences. Moreover, a smart BAS should be able to learn and explore the boundaries of comfort from the feedback received from the occupant. It can be assumed that the less an occupant needs to instruct the BAS (for example, by adjusting the thermostat) to change the environment (in this case, changing the temperature), the more the occupant is satisfied. Therefore, one of the goals of the BAS includes learning to set the temperatures to appropriate values throughout the day in a manner that minimizes the number of interactions with occupants and, at the same time, also minimizes the costs associated with heating and cooling.

To implement a smart BAS that is capable of dealing with the real-time uncertainties and constraints on occupancy and comfort requirements, this dissertation proposes a new closed-loop RL-based (Q-learning) supervisory control strategy for the HVAC system. The strategy is presented for two different applications: (1) to control the typical electric space heater, commonly used in many dwellings where the RL controller is only capable of switching the heater *on/off*; and (2) to control the HVAC system directly through an application programming interface, where the controller is capable of defining temperature set points. To validate each application, simulation results are presented for a single TZ and occupant. Results show that in both applications, the adaptive controllers are capable of learning the statistical regularities in the occupants' behavior, and appropriately scheduling the thermostat operation in order to guarantee comfort requirements and optimize energy costs. Scopes and limitations of the RL-based strategy are discussed on the basis of the vision described for AmI and SBs.

Using an RL-based BAS brings significant advantages for the automated learning of occupancy and comfort conditions. However, through this research, some important conclusions are drawn concerning the scalability of the computational representation for the state of the environment. In particular, taking into account the

thermodynamic state of the TZ and the laws that govern the evolution of temperatures in the TZ, there are different actions that could be explored by a smart BAS for additional energy efficiency, such as closing windows, blinds and doors. For example, to minimize the amount of energy used, the HVAC should be turned *off* if the zone is expected to be unoccupied, or if there is any other cost-effective means to guarantee the same comfort levels by using e.g. natural ventilation, solar gains, or by taking advantage of inter-zonal airflow with TZs that have higher thermal gains. The ideal goal for the BAS would be to obtain a control input profile that minimizes a given cost function, by using a dynamic system model (that establishes a direct input-output relationship from a temperature set point to the actual TZ temperature), as well as updated system information, including predictions such as weather, comfort requirements, occupancy, and the state of doors, windows and blinds. However, finding this control input profile is not a simple task and **using RL-optimized strategies to identify the most energy-efficient configurations may not be possible due to a dependence on the intractably large number of necessary states to represent the environment, as well as the size of the underlying policy set.** Therefore, as an important research step, appropriate models to represent the environment must be devised to support future optimization strategies.

Due to the importance that thermodynamic models have for MPC and energy demand forecasting, this dissertation focuses on context-based thermodynamic modeling and the identification of building TZs. MPC for energy and comfort management has proven to have clear advantages over other control strategies [101,102]. The drawbacks which currently hamper its widespread implementation include the proper identification of the TZ thermodynamics, the need for online estimation of the corresponding parameters which is robust in the presence of noise, and the fact that the adopted thermal comfort models do not reflect the complex, nonlinear features which characterize thermal comfort. Therefore, one of the goals of this dissertation is to contribute in the direction of obtaining a KR model capable of identifying the different dynamics of a TZ, with future applications for MPC, zero-net energy buildings, and ML algorithms.

### 1.2.2 Thermodynamic Modeling of Building Spaces

The development of adequate models to capture the dynamics of a building, especially the heat dynamics of building TZs for controlling indoor climate and improving energy efficiency, has fueled a great deal of research. Models are applied to simulation and

analysis problems, including passive design [103,104], energy use [105], and deriving predictive controllers for the building's thermal dynamics [101,102,106,107].

Environmental conditions inside a TZ depend on a plethora of factors, including zone architectural characteristics, construction materials, climate, occupancy and activities, and the state of electric equipment and temperatures in adjacent zones. Finding the appropriate thermal dynamics relating the control signals to average zone temperatures is a complex task, due to the complexity of the underlying physical processes [108,109]. Building environments are continuously changing with the occurrence of events, such as doors, windows and blinds being *opened* or *closed*. Changes in the configuration of the environment affect the underlying processes that govern the dynamics of temperature evolution of building spaces. For example, when a building is divided into environmental zones with occupancy-based HVAC control, and temperatures are adjusted to occupants' comfort preferences, an *open* door will increase the inter-zonal air-flow rate due to natural convection between two adjacent zones. This air flow has an impact on the thermal energy exchanged between both TZs and should be taken into account to obtain reliable analytical models. Luo and Ariyur showed, through simulation, that better modeling of the TZ environment, with more sensors to detect the state of doors and windows, can help to reduce the use of building energy by more than 20% [110]. Therefore, models should take these changes into account.

Using highly detailed physical models for prediction turns many approaches to solving energy management and control problems prohibitively large and complex, rendering them unusable for real-time applications. To circumvent this problem, several authors have used simplified and reduced models [111–115]. The purpose of model size reduction is to derive a low-order model of an intrinsically complex system to achieve a reduction in terms of computation effort, while preserving as much of the dominant dynamic description of the original system as possible. Methods for model reduction include, for example, selecting the appropriate time constants of the system [116], or selecting system modes according to their energy contribution [113]. For some modeling methods, a model should be detailed enough to provide a reliable representation of the TZ with a fast time-scale to control, for instance, the rapid flow of heat in a small room. In other situations, a slow time-scale model is enough to predict the mean temperature in the zone over each hour. Model reduction is always a compromise and the relative importance of various system characteristics is highly dependent upon the application. For this reason, Savo and Andrija state

that there can be no universal model reduction algorithm and state that “a reduced model is valid only over the range of conditions used to generate it” [117]. Therefore, notwithstanding the potential use for real-time applications, a reduced model fails to efficiently cover a broad range of conditions that would have to be described either by adding complexity to the model, or by using several different simplified models, with each model adapted to the range of conditions used to generate it.

**This dissertation shows how model reduction can be context-dependent, i.e. model parameters and structure depend on specific conditions that are relevant for a model during a certain time frame.** Even though a lot of research has been conducted within context-aware systems, the core term *context*, in many domains, is not yet a well defined concept [76, 118, 119]. This dissertation formalizes the concept of *context* to describe a particular thermodynamic behavior of the TZ. Different contexts are associated with different dynamics and, instead of using a single complex thermodynamic model for the TZ, a context-dependent model uses a set of simpler models and *context* as a concept to define the range of validity for each model. This range can depend on the state of discrete input variables that affect heat exchange, such as the position of window shades, and the opening of windows, or it can depend on the values of continuous variables, such as solar radiation, air-flow rate, and indoor temperature. For example, consider *context* being associated with the activation of an additional heater in the TZ if the outdoor temperature falls below a certain level, or with the state of a door connecting two adjacent TZs. This idea has been only superficially explored in the literature. However, there are references describing the importance of having different models to represent the dynamics of a TZ. Yashen Lin et al. state that convective heat transfer through an open door has a significant effect on the TZ’s thermal dynamics and showed that a door status sensor is required for temperature prediction [108]. For model-based control, the authors use two different models calibrated with data obtained in different door states (*opened/closed*), and use the door status signal to switch between these two models.

Most approaches in the literature address discrete changes in the building environment as disturbances, using statistical methods and stochastic frameworks to create models [120–124] and other closed-loop control strategies [125]. **However, this study conjectures that in many situations, models could be better adjusted to context.** This dissertation shows that if the boundary conditions that render some models more appropriate than others are observable and previously known, a context-based framework as a model selection strategy can be a very flexible so-

lution for immediate model commutation. This approach can complement, or even replace, multi-stage model selection strategies, such as those presented by Prívvara et al. [126] and Bacher and Madsen [127,128]. These strategies start with a set of initial candidate models for the TZ with different orders of complexity. The goal is to find the most suitable model for prediction. The maximum likelihood estimation is used to adjust each model’s parameters to optimal values, and the likelihood function value is used to compare performance between different model structures. Starting the selection from the simplest model, models with gradually higher levels of complexity are chosen iteratively up to a point at which the likelihood function value does not increase significantly. These strategies are adaptive and select the lowest-order model that best describes the model of the TZ. However, they need a certain processing time and, unlike our context-based framework, models do not immediately adjust to context changes as necessary for many real-time applications.

Context-based models are a contribution for multimodal thermodynamic modeling, whereby the environment is described by a set of distinct continuous time linear models, as opposed to just being described by a single time-invariant model. To the best of this researcher’s knowledge, a suitable framework to integrate these different models is not available in the scientific literature. **Therefore, this thesis advances the state of the art by providing a formal modeling framework based on hybrid automata that integrates the different models and describes the range of validity of each model** [129,130]. The term “hybrid” is used to characterize systems that combine *time-driven* and *event-driven* dynamics. The former are represented by ordinary differential equations (ODEs), characterizing the behavior of continuous-time variables such as temperature, while the latter, characterizing the switching conditions between contexts, is described through various frameworks used for discrete event systems, such as finite state automata or Petri nets [131]. In addition, the proposed framework also unifies hybrid automata semantics with the notion of context-based modeling from the AmI literature, where contexts represent discrete system configurations associated with the domain of each model, and context transition rules are associated with discrete transitions that govern model changes. An application example is provided to clarify this proposal, with a model for a single-zone building. Using previous knowledge about the physical characteristics of the building, lumped resistor-capacitor (RC) models are used to describe the building’s continuous dynamics in each *context*, due to their potential application for model MPC, where simplified and reduced analytical models are preferred. The provided example does

not address all of the details of HVAC modeling. However, the example is complex enough to convey the idea that context-based reasoning, hybrid systems, and thermodynamic models can be unified in a single framework. We compare the execution of the provided example with the simulation outputs of EnergyPlus - a simulator that uses models with higher complexity - to show that a context-based model can show comparable performance with more complex thermodynamic models concerning the prediction results, while keeping an acceptable low-level of complexity that is suitable for future use in a real-time feedback control loop.

### 1.3 Contributions

As a result of addressing the above research objectives and the research questions that they entail, this thesis contributes to advance the state of the art in the following topics:

- A new reinforcement learning strategy to control the heating ventilation and air conditioning system which, if implemented, may have a significant impact on reducing the global energy bill and greenhouse gas emissions over the lifecycle operation of the building.
- A description of the thermodynamics of a thermal zone using a set of distinct continuous linear time-invariant (LTI) models.
- A novel thermodynamic modeling framework based on hybrid automata that unifies the different LTI models that are used to describe the thermodynamic behavior of a thermal zone.
- Exploitation of the notion of context-based modeling from the AmI literature for simulation and knowledge organization of continuous LTI models.
- A description of a context-based modeling framework and the discrete event dynamics that govern *context* change.

### 1.4 Publications

This thesis is based on the following articles:

1. Pedro Fazenda, Paulo Carreira, Pedro Lima. Context-Based Reasoning in Smart Buildings. Proceedings of the First International Workshop on Information Technology for Energy Applications, 923:131-142, Lisboa 2012.
2. Pedro Fazenda, Kalyan Veeramachaneni, Pedro Lima, and Una-May O'Reilly. Using Reinforcement Learning to Optimize Occupant Comfort and Energy Usage in HVAC Systems. Journal of Ambient Intelligence and Smart Environments, 6(6):675-690, November 2014.
3. Pedro Fazenda, Pedro Lima, Paulo Carreira, Context-Based Thermodynamic Modeling of Buildings Spaces, Energy and Buildings, 124:164-177, 15 July 2016.

Part of Article 1 is summarized in Chapter 2, and Chapters 3 and 4 were based on Articles 2 and 3, respectively.

### **Additional publications, poster presentations and talks**

- (Talk) Edifícios Inteligentes, Jornadas de Engenharia Electrotécnica e de Computadores, Lisboa, Março, 2010.
- (Poster) Energy Efficiency Monitoring and Management to Promote Sustainable Behaviors, Second annual Conference for the MITPortugal, Porto, September, 2010.
- (Poster) Context-based Thermodynamic Modeling and Identification of Building Spaces, LARSyS annual meeting, Lisboa, June, 2013.
- (Poster) Context-based Thermodynamic Modeling of Building Spaces, LARSyS annual meeting, Lisboa, June, 2014.

## **1.5 Thesis Outline**

In this first chapter, we introduce the motivation, the research questions and objectives that define the scope of this thesis, highlight the contribution to knowledge, and list a number of support publications. The remainder of this thesis is organized as follows.

**Chapter 2, Related Work.** This presents a literature review and a discussion on related work. This line of work falls within the ambient intelligence domain.

This chapter gives an overview of the domain and discusses current and emerging technologies.

**Chapter 3, Reinforcement Learning for HVAC Control.** The optimization approach presented in this paper employs Q-Learning RL algorithms. Therefore, this chapter provides a brief but necessary technical introduction to Markov Decision Processes, reinforcement learning and sequential decision problems. Following this introduction, a reinforcement-learning-based controller for the heating ventilation and air conditioning is introduced. Two versions are presented. The first, called the *Bang-bang Heater*, presumes that the heating unit is controlled by simply turning it *on* or *off*. The second version, called the *Set Point Heater*, presumes that the unit has temperature set point control. The controller learns how to operate it in order to minimize a penalty function, which depends on the amount of energy used and the number of times an occupant adjusts the system. This chapter explains how Q-Learning can be used to solve the first problem, and how it can be extended to solve the second. It concludes with a discussion of scope and limitations.

**Chapter 4, Buildings and Context-Based Models.** This chapter describes the definition of *context* and the operational semantics of the proposed context-based framework and some important definitions. It then proceeds to describe resistor-capacitor model structure that was used to model the thermal behavior of a single-zone building. Illustrative application examples are presented.

**Chapter 5, Simulation Setup.** This chapter presents the simulation setup environment (Matlab, the Buildings Controls Virtual Test Bed, and EnergyPlus) for the experiments that was developed to evaluate the examples given in Chapters (3) and (4).

**Chapter 6, Results.** This chapter presents the results and conclusions of the experiments described in the previous chapter.

**Chapter 7, Conclusions and Future Work.** This chapter concludes the thesis and discusses directions for future research.

## 1.6 Summary

Energy can be saved in new buildings with more efficient technologies and smart buildings have been hailed as a solution to increase energy efficiency in the building sector. This chapter described the motivation for investing in energy efficiency, supported by smart building technologies. Important concepts were described including green and zero-net energy buildings, ambient intelligence, and context-based modeling. These concepts are important to understand the future of building research, and also why the initial research path, focused on optimizing the heating, ventilation and air conditioning system, was guided towards finding an appropriate modeling paradigm for knowledge representation and organization, using *context-based modeling*. Two research problems were presented for: automatic HVAC optimization, and thermodynamic modeling of building spaces. Methods to solve each problem were discussed, and contributions and publications were presented.



# 2

## Related Work

*“The cheapest energy is the energy you don’t use in the first place.” – Sheryl Crow (American singer and songwriter).*

Over the past three decades, a considerable amount of work has been done on using building automation strategies for energy savings, which naturally includes optimizing the operation of various systems, such as the HVAC. In this domain, a recent report from the U.S. Department of Energy has identified a set of high-priority initiatives for high-efficiency HVAC technologies that must be included in the HVAC research and development roadmap [132]. Some of these initiatives are highly focused on using emerging technologies to create new control schemes for HVAC operation, and developing open-source and open-architecture platforms to enable grid connectivity for demand response, and to exchange information with other building systems. Research supporting these technologies includes work in different research fields, including e.g., computational intelligence [133], distributed systems [134–136], context-aware systems [79], SGs [137], the Internet of Things (IoT) [138], sensor networks [139], control systems [13], thermodynamic modeling [140], information modeling [56], system identification [141], hybrid dynamical systems [142], multi-agent systems (MAS) [143] and

MPC [144]. Each research field is focused on its specific domain and there is no single field of research capable of addressing all the non-trivial problems that need to be solved for deploying AmI in SBs. There are a significant number of opportunities for interdisciplinary collaboration and more theories need to be synergistically integrated in order to provide better control solutions for building systems and contribute to the idea of having a rational BAS controlling the HVAC system. This chapter describes important work and ideas related to this thesis. It starts by describing related work on energy and comfort management for HVAC systems, and follows with a description of important concepts that should be taken into account in future research for AmI and HVAC control.

## 2.1 Energy and Comfort Management

Most of the past research for energy savings in HVAC systems has been focused on studying different mechanical designs and configurations to improve energy performance [97]. However, as stated by Vakiloroyaya et al., “the energy consumption of an HVAC system depends not only on performance and operational parameters, but also on the characteristics of the heating and cooling demand and the thermodynamic behavior of the building” [145]. This demand depends on meteorological conditions, occupancy, and many other factors, such as the state of doors, windows, lighting and equipment. Due to these factors, “the actual load of the HVAC systems is less than it is designed in most operation periods” [146]. Therefore, proper control of the heating and cooling demand becomes an essential requirement for HVAC energy reduction. In this domain, Dounis and Caraiscos [147] present a survey with the state-of-the-art on control systems for energy and comfort management in buildings. Ahmad et al. [133] also present a succinct review of computational intelligence techniques for HVAC systems. Work has spanned multiples strategies and control techniques, in general, are categorized into hard control and soft controls [148]. Classical control techniques, such as optimal, nonlinear, adaptive, and proportional–integral–derivative (PID) controllers are considered hard control. The focus of this thesis is on soft control techniques, where most AI algorithms are included. Most AI techniques have somehow been applied in building control strategies. This includes evolutionary algorithms [149–153], neural networks (NN) [99, 154–156], Bayesian networks [157] and other algorithms for control and optimization [158–161]. To express occupant’s preferences using linguistic labels that human operators can understand, a popular

approach for HVAC control has been to use fuzzy controllers [150, 151, 162]. Fuzzy logic is useful in capturing and representing imprecise notions, such as “hot” and “very hot”. However, it has several limitations [163, 164]. The fuzzy part is hard to program and prior knowledge is required to model the fuzzy system. The knowledge base is usually constructed based on the operators’ experience and requires fine-tuning and simulation before becoming operational. Their knowledge is often incomplete and episodic, rather than systematic.

The advantages of using a RL approach to machine learning, as an option over other paradigms, are well known [165, 166]. RL can be applied in situations where occupants do not know the correct answers required for supervised learning. A RL system is expected to learn how to achieve goals, by trial and error, in real-time, with continual feedback from its environment. This real-time, closed-loop, goal-seeking behavior, where the BAS automatically improves through experience, seems to be a crucial aspect of how humans operate, and is an interesting paradigm to be explored in SBs. Some of the unknown answers in the environment include occupancy and thermal preferences, as well as the thermodynamic behavior of building spaces. Moreover, there is also uncertainty concerning how occupants behave. For example, the BAS will not know if:

- Above a certain temperature, the occupant opens a window instead of using the HVAC system for cooling.
- The occupant will keep the lights *off* if the shades are *opened* before the room is unoccupied.

Most of these answers play an important role in energy and comfort control. Having the capability to automatically learn the uncertainties is one of the most important and desirable features for predicting heating and cooling loads and for smart HVAC control [99, 167, 168].

### 2.1.1 Models for Building Systems and Comfort Evaluation

For HVAC control, and decision support in algorithms, many authors use predefined models that represent expected behaviors for the occupants and building systems. Meyer and Emery [169], for example, propose an air conditioning system controller that generates an optimal plan to use thermal energy storage (the thermal capacitance of the building and a cold storage facility). The authors propose shifting part of

the daily on-peak cooling loads to the night off-peak hours when electricity prices are lower, also taking into account various factors, such as weather forecast and indoor heat gains. All parameters associated with the building's thermal response function and HVAC systems have to be previously characterized. However, most building characteristics are not time invariant and the state of windows, for example, affect the buildings' thermal capacitance. Although the authors suggest a procedure to identify system components, the resulting models will be highly prone to noise and modeling errors. Structures and systems degrade over time and depend on environmental conditions, such as temperature and humidity. In many situations, events in the building environment may change the context within these models are valid, thus potentially affecting their predictive performance. A SB must continuously keep learning in order to maintain building models that are adjusted to operating conditions.

### **Human thermal Comfort**

Another type of model, also used for HVAC control, predicts the thermal sensation of people. This type of model for thermal comfort can be used by an adaptive controller to guarantee comfortable temperatures in different buildings spaces that are subjected to different thermal loads. Dalamagkidis and Kolokotsa [170], for example, developed a RL environmental controller that sets the cooling/heating level and opens/closes a window by following a policy that maximizes a reward function based on a fixed weighted average of three factors: energy used, comfort and air quality. To estimate thermal comfort, the authors employ the predictive mean vote (PMV)/ predicted percentage of dissatisfied (PPD) (Fanger's comfort model) [171,172]. The PMV index is arguably the most widely used thermal comfort index today. It predicts the mean thermal sensation vote of a large group of people according to the seven-point thermal sensation scale, given by Table 2.1, proposed by the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ANSI/ASHRAE Standard 55-2010). The PMV index can be expressed as  $PMV = (0.303e^{-0.036 M} + 0.028)L$ , where  $M$  represents the metabolic rate quantifying human body heat production associated with human activity, and  $L$  is the thermal load defined as the difference between the rate of metabolic heat generation and the heat loss from the body to the surrounding environment, which depends on various factors, such as clothing, air velocity, air temperature and humidity.

The PPD index, which is a function of PMV, predicts the percentage of occupants that will be dissatisfied in a particular thermal environment. Figure 2.1 shows a plot

Value	Sensation
+3	Hot
+2	Warm
+1	Slightly warm
0	Neutral
-1	Slightly cool
-2	Cool
-3	Cold

Table 2.1: The predicted mean vote (PMV) sensation scale.

of PPD versus PMV, defined as  $PPD = 100 - 95e^{-(0.3353 PMV^4 + 0.2179 PMV^2)}$ . The ISO

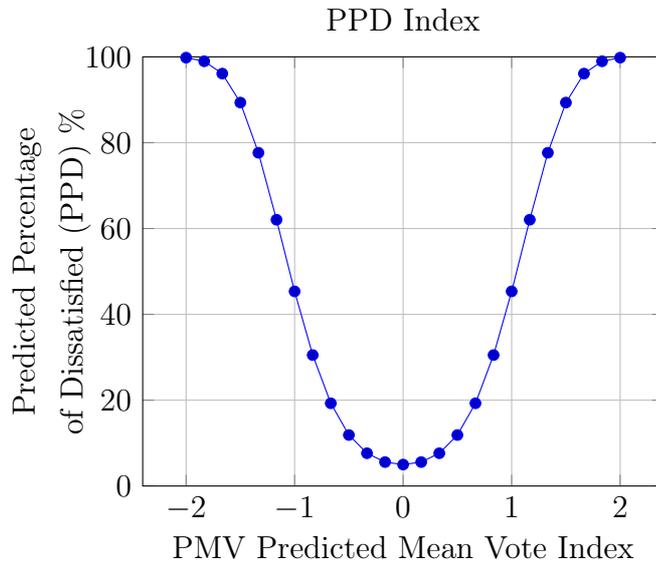


Figure 2.1: PPD – Predicted Percentage Dissatisfied.

7730 standard uses limits on PMV to define comfort. PMV values between  $[-1, 1]$  correspond to the range in which 75% of the occupants are satisfied, and between  $[-0.5, 0.5]$  is the range in which 90% of the occupants are satisfied. As a general rule, comfort conditions are considered acceptable when 85% of the occupants are satisfied [173]. Although this determination has been satisfactory for many HVAC applications, the PMV/PPD model is not optimal for every situation [174]. In the past few years, researchers have discussed its limitations extensively [175, 176]. Since the derived indexes are based on standardized assumptions for clothing, air velocity,

activities, and other factors, the algorithms that use these models may converge to temperature values that may not be optimal at all times. The PMV/PPD model does not take into account the fact that people usually react and adapt to the surrounding environment in order to avoid discomfort. Therefore, the model misses the fact the comfort can be affected by people’s actions [177]. In fact, occupants use numerous strategies to achieve thermal comfort e.g., activating ceiling fans, blinds, and activity level [178]. Moreover, people also adapt to the environment. Concerning this observation, adaptive comfort is based on the assumption that the comfort perception of people will depend on outdoor climate conditions. To quote from Mathews et al., “For low outside air temperatures, people will be comfortable if the indoor temperature is lower. The opposite is true for high outside air temperatures. This affords us the opportunity to potentially save even more energy. If the indoor temperature can be cooler when the outside air temperature is low, less heating would be required. During periods of high outside air temperatures, a higher temperature would imply that less cooling would be required” [179] (p. 153). Williamson and Riordan point out the fact that “...the reaction of people to a sense of being cold or hot is not necessarily to operate a heater or cooler, nor is such a reaction generally the sole response. Adjusting clothing, altering activity levels etc. are also common responses.” [180] (p.1). To save energy with the HVAC system, SBs should include adaptive environment controllers with the ability to learn and self-regulate according to the thermal preferences of occupants. **Therefore, SBs need to learn rules of behavior based on feedback they obtain from occupants, and continually adapt this knowledge [181].**

The RL-based controller developed in this dissertation explores the heating set points that satisfy human comfort conditions, while minimizing the needs for thermal energy. No modeling is needed neither for modeling comfort nor building structures and components. All the necessary information is extracted from the interaction between the occupant and the BAS, and the cost associated with the conditioning of the environment. The proposed method takes into account the interaction of the occupant and the energy used for heating and cooling in the reinforcement signal. The BAS is penalized in proportion to the amount of energy that it uses to guarantee the current temperature set point, including an additional penalty if an occupant acts on the interface at any instant of time.

### 2.1.2 Scheduling

The simplest strategy to save energy with an HVAC system includes using thermostats in an efficient manner. A thermostat acts as an interface between the occupants' thermal preferences and the operation of the heating and cooling systems by maintaining the temperature near a desired set point. Figure 2.2 shows an example of four different thermostats. Thermostats can range from simple mechanical control mechanisms (Figure 2.2a) to Internet-enabled programmable devices, offering programming options that allow tenants to define several set points according to a predefined schedule [182, 183]. They can also display important information, such as TZ temperature, ventilation rate, and in some cases, smart metering with pricing feedback [184]. Unfortunately, it has been shown that many occupants do not explore most functions of these interfaces, due to the fact that they do not understand them [182]. Even for people who understand these interfaces, programming setback schedules for every day of the week and time of year is a tedious task.



(a) Non programmable.



(b) Digital programmable with Wi-Fi.



(c) Smart-phone application.



(d) Smart with learning capabilities.

Figure 2.2: Thermostat interfaces from Honeywell (a, c), Emerson (b) and Nest (d).

Building administrators would be satisfied if a BAS could automatically and dynamically adjust the HVAC operation to optimal schedules. Therefore, to circumvent the task of manually programming these schedules, some authors have tried to estimate activities and occupancy by observing the environment. This included, for example, observing CO<sub>2</sub> levels [185, 186], monitoring the electrical load of the house and the hot water heating pattern over a certain period [187], or even using smartphones as tenant-location devices to predict the times of arrival and departure of occupants, and modifying temperature set points accordingly [188]. Undoubtedly, all this information can be useful for decision support. **However, this dissertation follows the plausible conjecture that the HVAC optimization problem can be reduced to learning how to operate within maximum energy efficiency, while trying to minimize the number of occupant “complaints”.** Adaptation to user inputs is allegedly used by Nest’s (Figure 2.2d) enhanced auto-scheduling thermostat [167]. However, the company does not disclose any details about its learning algorithm and, therefore, there are no means to evaluate the advantages and drawbacks of Nest’s learning algorithms. Nest’s smart thermostat proves, due to its popularity, that there is tremendous potential in the current market for smart thermostats. In fact, the smart thermostat market is estimated to be worth 5.9 Billion USD by 2020 [189]. Energy price and consumer environmental awareness will be a key driver, but widespread adoption of this technology is strongly dependent on network and platform interoperability [190].

Occupants should be able to express their preferences and satisfaction for building environmental settings in an easy and intuitive way. Following the same ideas behind Nest, in this dissertation, the research priority is to avoid demanding the occupant operate any temperature controller more complicated than the simple, inexpensive thermostat that is standard in most buildings – one that allows the occupant to command the temperature to be increased or decreased. Such a thermostat may optionally provide a temperature gauge, but it is assumed that such a feature is unnecessary. This is supported by the belief that occupants are perfectly satisfied with a simple interface if their inputs, using that interface, result in temperature set points that comply with their comfort requirements. The motive is to effectively circumvent sophisticated thermostats that require even small amounts of set-up or direct use, as these features are frequently ignored by and wasted on occupants.

### 2.1.3 Building and Occupant Interfaces

A classic and ambitious goal for AI is to have intelligent systems interact with humans up to the point where they can even engage in dialogue. With this vision, SBs are expected to include advanced HMIs, eventually with image-processing capabilities, NLP and speech recognition [191]. One of the drawbacks that hamper the widespread implementation of MPC is the lack of user-friendliness and occupant interactions [101,147]. SBs in contrast, aim to be adaptive, sensitive, and responsive to user needs, habits and emotions. Research in affective computing is enabling systems to recognize and respond to human emotions, and emotionally, SBs may have a clear advantage when it comes to human-computer interaction [192,193]. Human emotions can be used for decision support to help reduce frustration in a sense that if an occupant is angry or stressed, then the occupant is probably not very receptive to notifications about energy performance. Building technologies are expected to enable a wide variety of completely new user-friendly applications that facilitate the process of presenting information, enhance human decision-making, and increase passive control over the building environment.

#### Presenting Information

One of the desired requirements for AmI is to have occupants informed about important aspects of their environment. Therefore, building HMIs are expected to present values of power consumption and real-time information on how energy is being used with associated financial costs and GHG emissions. If this information is complemented with advice on efficiency improvements, which result in quantifiable savings, occupants can also assume active roles in promoting energy efficiency. Since buildings usually have a limited number of effectors that are capable of changing the building environment, the number of variables that can be automatically set or regulated by the BAS are limited. To change the environment without direct actuation, a SB can try to use a human-in-the-loop by giving him smart recommendations [194]. A diligent occupant can follow recommendations, and assume actions that change the state of the environment, if they are properly justified by showing, for example, that the same comfort level can be guaranteed using a different or more efficient environment configuration (e.g., opening windows and blinds to take advantage of natural lighting and ventilation). Therefore, SBs should be capable of generating these recommendations and working collaboratively with humans in order to improve living and working

conditions in an efficient manner.

There is a growing body of literature demonstrating that there is potential for energy savings through measures targeting occupant behavior [195]. HMIs can be used to notify occupants with information related to their individual behaviors such as, for example, unusual energy use patterns in areas of the building that they usually occupy, or the financial costs of leaving the HVAC and lights *on*, when the office is unoccupied. These notifications are important, because with information on individual behaviors occupants can be informed of the actual power consumption values and sensitized through feedback notification of their good behaviors. This behavioral feedback, as a form of self-regulation, has proven to impact the promotion of sustainable habits [196–198]. To inform occupants, efficient methods are needed for delivering messages that effectively motivate behavioral change. In a recent study, 500 dwellings in England participated in a social challenge to evaluate the impact that feedback had on behavior and energy efficiency [199]. The study results showed that by using a feedback scheme with one of the emotions shown in Figure 2.3, selected based on the evaluation of energy used by each dwelling ranked against its neighbors, an average daily electricity savings of 8.6% was accomplished. Feedback derives



A **“very happy” face** if the dwellings performance was in the lowest 25% of energy use among similar homes.



A **“smiley” face** if the dwellings performance was in the next 25% of users.



A **“neutral” face** for the next 25%.



A **“sad” face** if the dwellings performance was in the highest 25% of energy users.

Figure 2.3: A behavioral feedback scheme (using emotions) to demonstrate the amount of energy used by each dwelling, ranked against its neighbors (adapted from [199]).

from the straightforward comparison of their energy usage compared to similar-sized homes. This normalization is not always possible because houses, offices, and buildings have different sizes, configurations and systems, and are subjected to different environmental factors (e.g., solar radiation and wind), which affect the amount of en-

ergy used. Therefore, more accurate methods are necessary to generalize these types of feedback strategies.

In order to provide “smart” advice and send feedback notifications, the BAS needs to have the ability to reason over what’s happening in the building and have the ability to measure, predict, compare and find alternative building configurations. The BAS could execute, for example, some of the following verifications:

- **Compare power consumption with design predictions.** In the case of discrepancies, verify if it is because of weather conditions, thermostat settings, malfunctioning systems, etc.
- **How much will power consumption change with different thermostat settings or ventilation rates?.** This can give an estimate of the cost of comfort.
- **How much money can be saved by retrofits of the building shell or equipment?.** If retrofits are implemented, the BAS can verify if the results align with predictions.

Therefore, the problem of creating meaningful models of building systems that allow the BAS to characterize, for example, how energy is used in each room – and planning strategies based on those models that fulfil some expectations, like saving energy – are fundamental areas of research that need to be addressed to accomplish those requirements.

In light of the focus this dissertation, thermodynamic models enable a BAS to predict the thermal loads in different TZs which has, among other things, applications for MPC, DR, and planning strategies for zero-net energy buildings. Despite the variety of approaches that exist in the literature for inferring black-box thermodynamic models from time-series data, none are particularly well-suited for building apt descriptions of TZs, since they provide little or no interpretation of the underlying physical parameters [127, 128]. To create models with physical interpretation, other numeric approaches use prior physical knowledge (e.g., grey-box models) of the heat dynamics of the building to predict the thermodynamic behavior of TZs with sufficient accuracy [128]. However, as described in Chapter 1, models of arbitrary TZs often require a vast number of parameters, which obfuscates the interoperability of the inferred models. Therefore, the advantage of context-based thermodynamic models is that they are composed of a set of simplified models, each capable of fully describing

a distinct qualitative thermodynamic behavior of the represented TZ. The fact that these models are simple, and analytical, results in significant advantages for MPC and other real-time applications, such as planning [200] and model checking [201–203].

### **Explainable Actions**

When interacting with SBs, most users will not engage in trying to understand complex systems and reports or perform complex configurations or any other tasks that are time-consuming and require some amount of technical expertise. Moreover, when it comes to AI, humans are not very tolerant of accepting systems intended to perform like humans, but fail to do so. Humans will not delegate control to a BAS that becomes annoying by systematically contradicting their intentions when controlling the environment by, for example, inexplicably switching the lights or the heater *off*. Therefore, SBs should be able to explain automated decisions to facilitate engagement with occupants, by providing some insight on why algorithms are assuming certain behaviors. Having “explainable” actions is an important requirement for AmI because actionable information may be needed by a human operator to trace and understand the operation of the building, as well as detect malfunctions or undesirable behaviors [204]. An occupant or facility manager could ask the BAS, for example, the following questions:

- *Why is the HVAC off?*
- *What are the costs associated with lighting?*
- *How much energy is currently being wasted with the HVAC system?*
- *How much energy will be saved if the HVAC is off and a window is open?*
- *Will comfort conditions be guaranteed if the HVAC is off and a window is open?*
- *For how long was the HVAC on on a given date?*
- *How much energy and money will be spent on heating until the end of the day?*
- *If a window is closed, how much money will be saved in heating during the month of December?*

Answering these questions includes work on ML for intelligent environmental control systems, context-aware architectures, and KR. Knowledge should be organized in a

way that makes it easy for the BAS to reply to such queries. To answer these questions, the BAS must have a rich understanding of its environment, which includes the ability to create models from observed data, representing knowledge associated with these models, and reasoning over that knowledge. Although many parametric ML techniques, such as NN and support vector machines, are numerically accurate in predicting building systems, they shed little light on the internal structure of the system and its governing principles. Therefore, they are not suitable for supporting explainable actions. Conventional AI attempts to mimic human reasoning using expressive KR languages that are structured and formally well understood with reasoning algorithms capable of dealing with the expressiveness of such languages [205–209]. Although much of this work can be extended to SBs, it is not yet clear how to create models for building environments that integrate these expressive languages, taking into account that many building variables are continuous.

In order to have a BAS achieve human-level performance in answering questions posed by occupants, basic knowledge of the common-sense world will be necessary. However, in the current state of the art, it is not yet clear how this vision can be implemented. There is no easy, single solution for creating environments capable of accomplishing automated common-sense reasoning. Therefore, different approaches to AI need to be explored and combined in order to create intelligent systems [90]. In light of the current literature in context- and case-based reasoning [78, 210] and middleware architectures [80], planning in hybrid domains [200], and explicability and predictability for task planning [211], it is this researcher’s belief that context-based thermodynamic models can play an important role in supporting explainable actions. Not only do context-based models provide insight on the internal structure of the governing principles of the thermodynamics of the TZ, but they also produce accurate predictions of the TZ temperature, as shown in Chapter 6. Answering questions with context-based models, by extending the work being done in the current literature, is future work.

## 2.2 Context-awareness for Ambient Intelligence

There are many gaps in our understanding of large complex systems and our ability to engineer them [4]. To cope with such complexity, most software architectures for AmI are programmed in a modular way. When applied to SBs, this modularity deals with the complexity of the domain by dividing the operation of the building

into a number of interdependent modules, which are able to control independent building systems and services. Following this approach, many authors employ a MAS approach [212–216] as a decentralized solution to control SBs, with a number of advantages, including scalability and re-configurability. MAS consists of a collection of software agents that execute the behavior associated with each of the SB’s modules. Agents can also be used to represent occupants, in order to maintain their preferences concerning environmental conditions. The MAS approach is scalable, as new agents can dynamically enter the scenario and start participating in the operation of the SB, when new components and services are added. However, although modularity simplifies the development of AmI software architectures, most agents responsible for each control logic are largely deployed in isolation, and the interaction between multiple agents in the MAS may result in undesired emergent behaviors. The term emergent is frequently used to describe behaviors that arise from the interaction of subsystems and are not evident from the analysis of each subsystem. Consider the following example: *an agent, programmed to optimize the use of natural lighting in a room, will open the window blinds and turn off the lights. This action may inadvertently increase the temperature inside the space due to solar gains. The agent that manages the HVAC will notice this increase and will try to cool down the room, thus spending more energy.* Without the perception of this causal relation between lighting, temperature and energy, two agents designed to save energy by managing each of their isolated domains, may end up spending even more energy when working together in the MAS.

Recent publications have begun acknowledging the importance of integrating the control of several different systems. Luigi Martirano [217], for example, recognized the importance of integrating the lighting control system with HVAC control and with solar blinds. Therefore, to avoid emergent behaviors, this dissertation discusses another type of modularity: the operation of each TZ depends on a set of active *contexts*. [72–74]. With strategies organized according to *context*, emergent situations may be easier to detect because the definition of a *context*, as described in this dissertation, is explicitly associated with the state of the TZ and for each *context*, a specific thermodynamic behavior is expected.

## 2.3 Context-based Framework

While context-modeling deals with how contexts are represented, stored and presented, context-awareness is the capability to reason about the context in order to make decisions about the actions to be triggered [218–221]. Designing a conceptual context-aware framework includes identifying the set(s) of *context* and the transition rules that define how to transition from one context to another. The classical frame problem is closely related to this issue [222]. The design process has to include the experience of human experts to model the necessary knowledge associated with the operation of particular types of buildings, equipment, systems and services. Gonzales et al. [75] formalized the definition of context-based reasoning, with applications for modeling some human behavior that control an autonomous agent performing a tactical mission in some environment. In their formulation, *context* is a 3-tuple  $(Ak, Tk, Dk)$  composed of the following three basic elements:

- ***Ak*–Action knowledge.** Required for the agent to carry out the behavior encapsulated within the *context*. It represents the agent’s functional intelligence within its given environment for a specific situation.
- ***Tk*–Transitional knowledge.** Indicates when a transition to another context is warranted. It can be expressed as *IF*(conditions) *THEN* (activation) transition rules.
- ***Dk*–Declarative knowledge.** Describing tactical knowledge (represented as attribute-value pairs) required to successfully execute the action knowledge.

Exercising the context-based model is the process of activating the set of contexts that best suits the situation at hand. This activation allows the active contexts to take over and control the execution of a process, defining behaviors, constraints, and other context-dependent characteristics. The process may survey the environment, as well as its internal state (including transitional knowledge), to determine the conditions where the current context is deactivated and a new context is activated and execution resumes. This formulation, presented by Gonzales et al., is generic enough to cover a wide range of applications. Therefore, this dissertation draws inspiration from their work and explores the similarity that their context-based reasoning model has with the description of a *hybrid system*. Although we do not represent and use *action knowledge* (reserved for future work), *declarative* and *transitional knowledge*

can be directly extracted from the hybrid system formalism, as described in Chapter 4. For the particular context-based thermodynamic models, given in this dissertation, the *declarative knowledge* includes the different models that describe the TZ in each *context*, while the *transitional knowledge* includes the rules that govern the model changes in the time-variant TZ environment.

## 2.4 Summary

This chapter discussed some of the literature, related work and ideas that guided the research presented in this dissertation. It started by discussing some previous work for energy and comfort management and the advantages of using reinforcement learning. Models have been used for buildings and for evaluating the thermal sensation of people. This chapter described the most common model used in the literature, which is the predictive mean vote (PMV)/predictive percentage of dissatisfied people (PPD) model. The disadvantages of using the PMV/PPD model were discussed, pointing out the fact that most model parameters are hard to observe and that users can adapt to the surrounding environment. Therefore, many methods that use the PMV/PPD model do not explore the most efficient conditions for heating and cooling. To solve this problem, this dissertation presents a method for learning the heating set points that satisfy the occupant's comfort, while also saving energy. The method is also capable of learning the occupant's schedule in order to avoid using energy when the occupant is out of the office. The chapter proceeded with a discussion of important requirements for ambient intelligence, centered on building occupant interfaces that are capable of user-friendly interaction for information feedback presentation, and for explaining automated decisions. Knowledge representation should take these requirements into account. The chapter ended with a discussion and presentation of previous work on context-awareness for ambient intelligence. This dissertation draws inspiration from previous work and explores the similarity that exists between context-based reasoning models and hybrid systems.

# 3

## Reinforcement Learning for HVAC Control

*“Civilization advances by extending the number of important operations which we can perform without thinking about them.” – Alfred North Whitehead (Mathematician and philosopher).*

The BAS is an autonomous system capable of executing actions to change the building environment. It is autonomous to the extent that its behavior is determined by its own experiences. Therefore, the BAS should be able to learn and evolve by evaluating the quality of its actions with regard to certain performance metrics that are used to track progress when satisfying a set of goals. This chapter discusses the application of a discrete and a continuous RL approach for HVAC control. The purpose of this approach is to actively learn how to schedule the HVAC operation and set the thermostat temperature set points based on feedback obtained from the occupants and the amount of energy used for heating. The discrete approach, which we call *Bang-bang Heater*, presumes that a heating unit is controlled, at a low level, by a controller that guarantees the environment temperature will converge to a certain set point when the heater is *on*. Heating is supplied by an electric space heater with no thermostatic adjustment and the BAS can only act to turn the heater *on/off*. In the

continuous approach, which we call *Set Point Heater*, we presume the HVAC unit has temperature set point control and the BAS has to actively learn how to schedule the HVAC and set thermostat temperature set points. In both applications, performance metrics include minimizing the energy used by the HVAC system while minimizing the number of times an occupant interacts with the HVAC control interface to procure comfort.

We proceed by presenting a brief introduction to RL, with particular focus on the Q-Learning technique. We then describe how Q-Learning can be applied to solve both applications and finish the chapter with a discussion of the presented methodologies. Both the *Bang-bang Heater* and *Set Point Heater* problems are demonstrated in Chapters 5 and 6 by setting up a simulated environment with simulated occupant and low-level *Bang-bang* and *Set point Heater* control behaviors.

### 3.1 Background

In this thesis the smart BAS is portrayed as an **intelligent agent** coupled to the building environment. The BAS model follows the general architecture of a particular **utility-based** agent, as described by Russel and Norvig [223]<sup>1</sup> and represented in Figure 3.1. As with any other agent, the BAS perceives and acts upon the environment. It uses sensors to read environmental variables and acts using effectors to change building configurations. During the perception process the BAS has to “understand” the environment and map the often complex and noisy information received from its sensory inputs to a reduced and simplified internal representation of the environment, defined as *state*. As time passes, this representation transitions through a series of states taken from a set of finite possible states, denoted as  $S$ , repeatedly and in unpredictable ways. If the environment is **fully observable**, as we assume henceforth, the BAS is capable of detecting all possible states of the environment with its sensory inputs and can keep the internal model up to date without uncertainty. This is the case when percepts are digital inputs to read, for example, the activation state of electrical equipment or the opening states of doors and windows, and the internal state representation is a vector containing each input. In some applications, however, inputs are analogue and the straightforward approach to representing them uses quantization. In this case, the lower the quantization error, the higher the cardinality of  $S$  and in the limit, state representation can even include continuous components,

---

<sup>1</sup>Chapters 2, 17, and 21.

as we describe further in Section 3.1.3.

By acting on the environment, the BAS is capable of changing the state of the environment and consequently its internal state representation. Actions can change the binary state of outputs to connect/disconnect electrical equipment or change the value of a discrete variable such as the heating level of a heater, which can be set to a *Low*, *Medium*, or *High* state. Actions can also be continuous to set temperature set points or lighting levels. In either case, **the main challenge of a utility-based BAS is to sequentially choose the actions that produce optimal behavior, considering and balancing the risks and rewards of acting in an uncertain environment.**

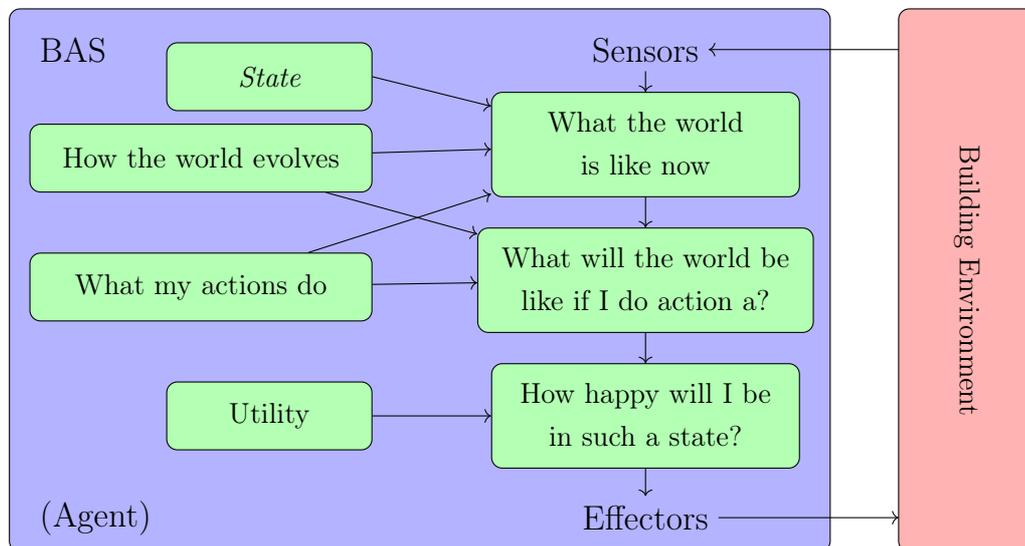


Figure 3.1: The BAS as an intelligent utility-based *agent* coupled to the building environment with goal-directed behavior (adapted from Russel and Norvig [223]).

### 3.1.1 Sequential Decision Problems

A BAS must find ways to adapt to its environment to continuously improve performance while manipulating state transitions with actions. As a utility-based agent, the BAS defines for each particular state a performance measure of how desirable that state is. This measure is given as a utility function  $U : S \rightarrow \mathbb{R}$ , which maps each state to a measure that quantifies the prediction of how “happy” the BAS will be in that state. This measure of “happiness” is directly associated with the **rewards** the BAS expects to receive in the states that will follow in the future.

During the interaction between the BAS and the environment, RL agents learn which actions to use through trial and error. Figure 3.2 graphically represents this interaction. At each time step and state  $s_t$ , the BAS chooses an action  $a_t$  from a finite set of possible actions, denoted as  $A$ , according to a **policy function**  $\pi : S \rightarrow A$  that maps each state of the environment to a suitable action. As a result, the environment’s state is updated to  $s_{t+1}$  and execution resumes. During this process the BAS has “knowledge” on how the world evolves, i.e., the BAS knows the set of possible states that are reachable from  $s_t$  through the outcome of each action  $a_t \in A$ . However, since most environments are non-deterministic, actions are unreliable. There is no guarantee that the environment will end up in a specific state. For example, if the BAS turns a heater *off*, there is no guarantee that the occupant will not turn it back *on*. Therefore, the **transition model** between states is stochastic and has to be described as a set of probabilities  $p(s_t, a_t, s_{t+1})$ , denoting the probability of reaching state  $s_{t+1}$  if action  $a_t$  is done in  $s_t$ . Transitions are assumed to be **Markovian** in the sense that probabilities depend only on the current state and not on the history of earlier states. Depending on this transition model, the environment will transition through a sequence of unpredictable states  $[s_0, s_1, s_2, \dots]$ , depending on the probabilistic outcome of each selected action.

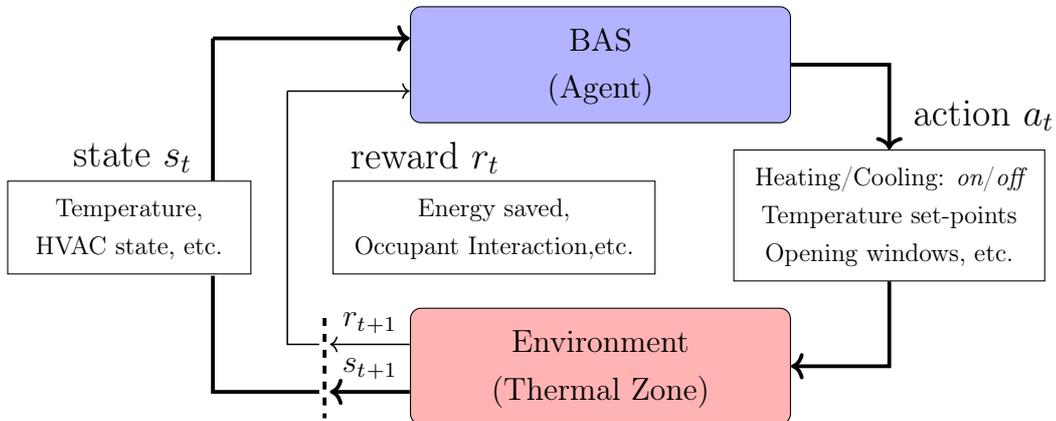


Figure 3.2: The interaction between an agent and its environment, represented as a series of state transitions that are associated with actions and rewards.

### Reinforcement Function

The RL BAS takes actions to maximize some notion of cumulative reward. It follows that in each state, the BAS receives a **reward** that may be positive or negative and is given by a bounded reward (or reinforcement) function  $R : S \rightarrow \mathbb{R}$ . The utility measure associated with a sequence of states depends on these rewards. For a specific state and time instant, the utility is associated with the **expected sum of discounted rewards** that the BAS will receive in the future while following a certain policy  $\pi$ , defined as:

$$U_t^\pi(s) = E \left[ \sum_{\tau=t}^{\infty} \gamma^\tau R(s_\tau) \mid \pi, s_t = s \right]$$

where parameter  $\gamma \in [0, 1]$  is the **discount factor** that describes the preference of the BAS for current rewards over future rewards. When  $\gamma$  is close to 0, rewards in the distant future are viewed as insignificant. When  $\gamma$  is 1, all rewards are considered with the same importance. The problem of sequentially choosing appropriate actions in order to maximize expected rewards is called a *Markov Decision Process* (MDP) and is formally defined by the 4-tuple  $(S, A, P, R)$  containing the sets of states, actions, transition model, and rewards.

### Learning the Optimal Policy

The task of the BAS is solving a MDP by learning the **optimal policy**, which is a function, denoted by  $\pi^*$ , that maps each state  $s_t$  to the optimal action to execute in that state. Considering the future set of states that are reachable from  $s_t$ , and the respective probabilities associated with each transition, the **optimal action** is the one that maximizes the expected utility of the subsequent state:

$$\pi^*(s_t) = \operatorname{argmax}_{a_t} \sum_{s_{t+1}} p(s_t, a_t, s_{t+1}) U(s_{t+1}) \quad (3.1)$$

To calculate the optimal policy, the BAS has to calculate the utility of each state, which depends on the utility of the neighboring states. Assuming that the BAS follows the optimal policy, the utility of a state is given by the sum of the immediate reward received and the **expected discounted utility** of the next state. This sum is obtained through the following the backward recursive **Bellman equation**:

$$U(s_t) = R(s_t) + \gamma \max_{a_t} \sum_{s_{t+1}} p(s_t, a_t, s_{t+1}) U(s_{t+1})$$

To solve the MDP problem, the Bellman equation has to be solved for each  $s \in S$  in order to calculate all the utilities. This can be done using an iterative approach, where utilities are updated in each iteration using the **Bellman update**:

$$U_{t+1}(s_t) \leftarrow R(s_t) + \gamma \max_{a_t} \sum_{s_{t+1}} p(s_t, a_t, s_{t+1}) U_t(s_{t+1})$$

Starting from arbitrary values for utilities, if the transition model is stationary and the Bellman update is applied repeatedly, all the utilities converge to equilibrium. The final utility values are unique solutions for the Bellman equations, and the corresponding policy obtained using (3.1) is optimal.

### 3.1.2 Q-Learning

If there are  $n$  possible states in the MDP then there are  $n$  Bellman equations – one for each state. The Bellman equations can be solved iteratively if the agent knows the transition model  $p(s_t, a_t, s_{t+1})$  and reward function  $R(s)$ . However, in most cases, the BAS has no prior knowledge of this information and needs to learn it as it is executing actions in the environment. Although the transition model can be learned by calculating transition probabilities (after observing the experienced transitions between states), this learning procedure can be expensive to execute in real time. Calculating the transition model and solving the MDP, by finding the solutions of the Bellman equations for all possible states, can become intractable if the number of states is too high. Hence, to circumvent this limitation, Q-Learning [224] can be used in lieu of policy search. A Q-Learning agent learns an action-value function (Q-function),  $Q : A \times S \rightarrow \mathbb{R}$ , instead of learning utility values. This Q-function ultimately gives the expected utility (*Q-value*) of doing an action  $a$  in state  $s$ , which is directly related to utility values as:

$$U(s) = \max_a Q(a, s)$$

The utility of a state is equated to the action-value of the most promising action to execute, which is defined in Q-Learning as the action with the highest *Q-value*. With Q-Learning the BAS can compare the *Q-values* of its available choices in each state without knowing the outcome of each action. The BAS can perform continuous updates to the Q-function during the learning cycle by executing the selected action in a particular state and immediately learning from that experience. This incremental model-free RL technique facilitates the implementation and deployment of the RL

agents in real-time applications, which is a fundamental requirement for building automation and HVAC control.

### Temporal Difference Learning

The implementation of the Q-Learning algorithm is accomplished by observing transitions and rewards and consequently adjusting action-values after each iterative step, using the temporal difference (TD) learning method [223, 225]. The BAS that learns a Q-function does not need a Markovian transition model for either learning or action selection. The Q-function is adjusted, after each state transition, using the following one-step TD update procedure:

$$Q_{t+1}(a_t, s_t) \leftarrow \overbrace{Q_t(a_t, s_t)}^{\text{Old Value}} + \overbrace{\Delta Q(a_t, s_t)}^{\text{Update}} \quad (3.2)$$

which is calculated whenever action  $a_t$  is executed in state  $s_t$ , leading to state  $s_{t+1}$ , with

$$\Delta Q(a_t, s_t) = \delta \left[ r_{t+1} + \gamma \max_{a_{t+1}} Q_t(a_{t+1}, s_{t+1}) - Q_t(a_t, s_t) \right] \quad (3.3)$$

where  $r_{t+1}$  is the reward observed after performing  $a_t$ ,  $\left[ \max_{a_{t+1}} Q_t(a_{t+1}, s_{t+1}) \right]$  gives an estimate of the optimal future utility of  $s_{t+1}$ , and parameter  $\delta$  is the **learning rate** ( $0 < \delta \leq 1$ ) that controls convergence to optimal action-values. If each action is executed in each state an infinite number of times, and  $\delta$  is decreased with an appropriate schedule, the algorithm converges to equilibrium values, denoted as  $Q^*$ , that verify the following constraint equation:

$$Q^*(a_t, s_t) = r_{t+1} + \gamma \sum_{s_{t+1}} p(s_{t+1}, a_t, s_t) \max_{a_{t+1}} Q^*(a_{t+1}, s_{t+1})$$

and the learned policy, which gives the action with the highest expected value, is provably optimal and given by:

$$\pi^*(s) = \max_a Q^*(a, s)$$

### Exploration versus Exploitation

When the BAS starts operating it has no previous knowledge of what actions to select. The strategy in this case is to explore available options while exploiting whatever feedback the BAS can get with respect to its objectives. During the learning process, the Q-Learning BAS must balance its selection of actions that currently appear to

be the most productive by following a known policy, with the need to further explore the state-space in order to find a better policy that will bring greater rewards.

Exploration is useful because the model of the environment can change and can become unadjusted to the learned policy. As an example, consider that the BAS can learn an optimum set point when a tenant is in the TZ and should exploit its utility by controlling it consistently. However, it must also expect that the zone may become unoccupied without any explicit signal. For this possibility, it must explore set points more appropriate to vacancy in order to learn a tenant occupancy habit. Therefore, the BAS must make a trade-off between exploitation to maximize its reward as reflected in its current utility estimates and exploration to maximize its long-term performance. There are many schemes on how the agent balances this decision. This thesis uses a simple **greedy in the limit of infinite exploration scheme** (GLIE), where the BAS chooses a random action in some selections depending on an **exploration rate** parameter, or follows a greedy policy by selecting the best action otherwise.

### 3.1.3 Continuous State and Action Spaces

When using the standard Q-Learning algorithm, the BAS can deal solely with a finite and discrete number of states and actions. The standard Q-function is usually implemented using a two-dimensional lookup table indexed by state-action pairs. However, for some applications, the BAS needs to have the ability to respond to “smoothly” varying states with smoothly varying actions. This requirement imposes a practical implementation limitation. Solving the smooth variation problem with the straightforward solution of scaling to large numbers of states and actions, as a means of covering a range of continuous values, becomes impractical, because learning and Q-function implementation becomes computationally expensive.

Considering that states and actions are vectors, possibly with a different number of continuous elements, generalization for the interval values between discrete states and discrete actions can be introduced to Q-Learning by using function approximation instead of table-based storage. Therefore, to deal with high-dimensional continuous states and actions in RL we used the **wire fitting** method proposed by Baird and Klopff [226] that has been used by Gaskett for robot control [227,228]. The wire fitting method uses a function approximation system, such as an artificial feedforward NN, to map the state vector to a set of  $n$  action-value pairs called **wires**. The set with all the wires is calculated in real time by propagating the state through the NN as

illustrated in Figure 3.3. The wire associated with the continuous optimal action with the highest associated  $Q$ -value, denoted by  $(a_{\text{Max}}, q_{\text{Max}})$ , is available at the output of the NN. The wire fitting method allows  $a_{\text{Max}}$  to be immediately calculated through the comparison of  $Q$ -value outputs, without any additional computation. This represents an advantage for real-time action selection and RL.

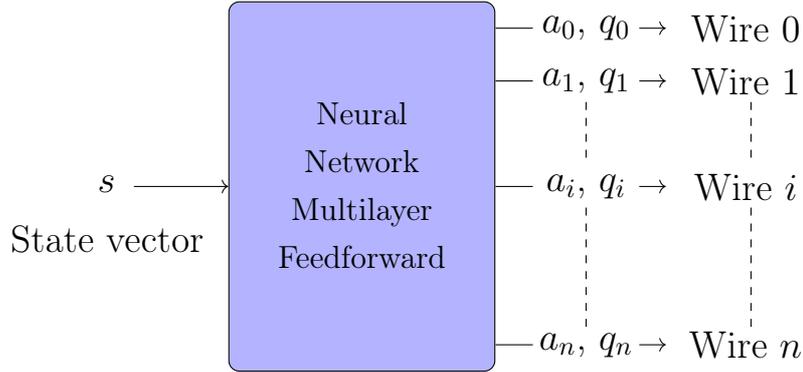


Figure 3.3: Using an artificial neural network to map the state vector into a set of action-value pairs (Wires) (adapted from [228]).

### Exploration and Action Selection

When following a greedy policy, the action selected for execution is the optimal action  $a_{\text{Max}} = \pi^*(s)$ . Since  $a_{\text{Max}}$  is a continuous action, exploration can be accomplished by adding random variable (noise) to the value of this action. The action selected for execution can be, for example, a random sample drawn from a normal distribution  $a \sim \mathcal{N}(\mu, \rho^2)$ , with a mean centered at  $\mu = a_{\text{Max}}$  and a standard deviation  $\rho$  proportional to an exploration rate parameter, denoted by  $\varepsilon_{\text{Rate}}$ . The additive noise provides the opportunity to explore the action space centered around the optimal action, depending on how this space is “stretched” by the  $\varepsilon_{\text{Rate}}$ .

During the learning phase, the set of wires at the output of the NN need to be adjusted after each execution step while considering the  $Q$ -value of the selected action. Since noise is added to  $a_{\text{Max}}$ , the selected action is most likely not part of the wire set. Therefore, generalizing and obtaining  $Q$ -values for action-values other than the ones included in the set can only be accomplished using a **wire fitting interpolation** function, as shown in Figure 3.4. The interpolation function takes as the input arguments the set of wires and the selected action. The output is the

locally weighted interpolation of  $Q$ -values given by:

$$Q(s, a) = \lim_{\varepsilon \rightarrow 0} \frac{wsum(s, a)}{norm(s, a)} \quad (3.4)$$

with,

$$wsum(s, a) = \sum_{i=0}^n \frac{q_i(s)}{dist_i(s, a)}$$

$$norm(s, a) = \sum_{i=0}^n \frac{1}{dist_i(s, a)}$$

and

$$dist_i(s, a) = \|a - a_i(s)\|^2 + c(q_{\text{Max}}(s) - q_i(s)) + \varepsilon$$

where  $i$  is the wire number corresponding to the action-value pair  $(a_i, q_i)$ ,  $c$  is a “smoothing” factor <sup>1</sup>, and  $\varepsilon = 0.001$  is a value that avoids division by zero.

For any particular action  $a$  the interpolator defines  $Q(s, a)$  as the weighted average of  $q_i$  values such that, if the action for evaluation is near a particular  $a_i$  (wire), then the corresponding  $q_i$  value is given more weight in the calculation of the average. Using the wire fitting interpolator, the RL algorithm has access to the  $Q$ -values of the entire continuous action-space range. This enables other actions to be used for execution aside from the ones included in the discrete set of wires. This generalization is necessary to implement the exploration task that the BAS must execute in order to adapt to changes that occur in the environment.

As an interpolation example, Figure 3.5 shows a graph with actions vs.  $Q$ -values, with the wire fitting interpolation using three wires placed at the following locations:

$$\{(0.2, 0.3), (0.5, 0.7), (0.8, 0.5)\}$$

The example shows the wire interpolation using two values for the smoothing factor:  $c = 0.5$  and  $c = 0.0$ . A property of the wire fitting interpolator is that in both cases  $q_{\text{Max}}$  always coincides with the highest interpolated  $Q$ -value. The smoothing factor defines how strictly the interpolation passes through the other wires. In the limit where  $c = 0.0$ , the interpolation coincides with all the wires.

### Learning with Wire Fitting

Learning with the wire fitting method includes adjustment of the wires after each execution step and training the NN to learn the new state-to-wire mapping [229]. This

---

<sup>1</sup>To simplify the interpretation of (3.4), consider  $c = 0$ .

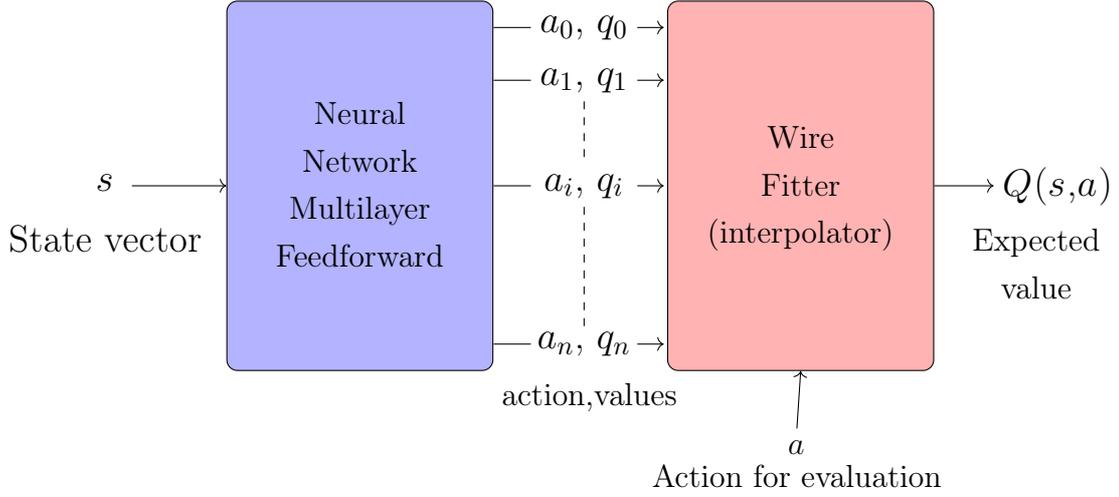


Figure 3.4: Obtaining the  $Q$ -value of a specific action  $a$  using the wire fitting interpolation function (adapted from [228]).

adjustment, represented in Figure 3.6, is accomplished by training the NN through backpropagation of errors. Since (3.4) is a continuous and smooth function of its inputs, it is possible to backpropagate errors through the wire fitting block to update the weights of the NN according to the chain rule. With this rule, wires are adjusted considering the partial derivatives of the interpolator and the Q-Learning TD update value given by (3.2), as follows:

$$q_i \leftarrow q_i + \frac{\partial Q}{\partial q_i} \Delta Q$$

$$a_{i,j} \leftarrow a_{i,j} + \frac{\partial Q}{\partial a_{i,j}} \Delta Q$$

where  $j$  selects a term of the selected action vector, and the partial derivatives are given by:

$$\frac{\partial Q}{\partial q_i} = \lim_{\epsilon \rightarrow 0} \frac{\text{norm}(s, a) \cdot (\text{dist}_i(s, a) + q_i \cdot c) - \text{wsum}(s, a) \cdot c}{[\text{norm}(s, a) \cdot \text{dist}_i(s, a)]^2}$$

$$\frac{\partial Q}{\partial a_{i,j}} = \lim_{\epsilon \rightarrow 0} \frac{[\text{wsum}(s, a) - \text{norm}(s, a) \cdot q_i] \cdot 2 \cdot (a_{i,j} - a_j)}{[\text{norm}(s, a) \cdot \text{dist}_i(s, a)]^2}$$

The wire fitting method allows smooth variations in actions with smooth changes in the input state vector. These changes allow fine adjustments to be made to the selected action as the input state progresses. Moreover, although the method enables

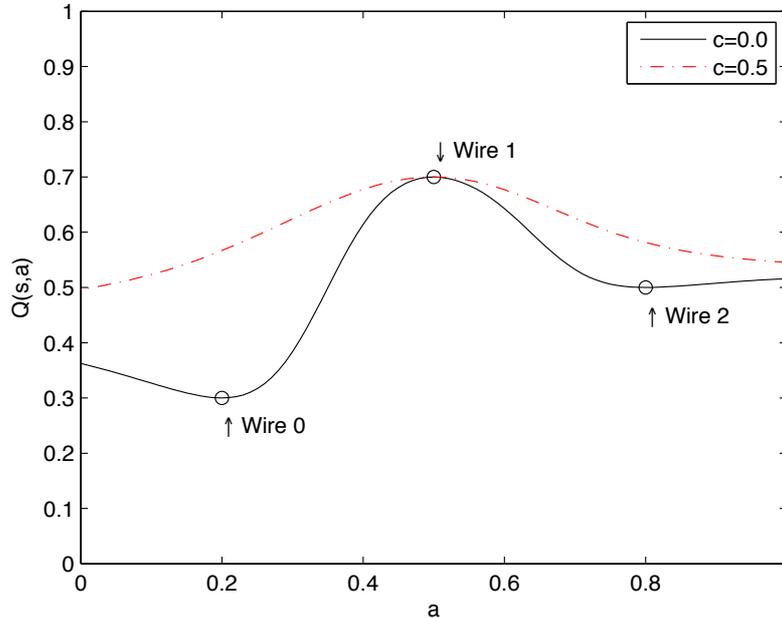


Figure 3.5: Weighted-nearest-neighbor interpolation with three wires (shown as  $\circ$ ) for  $c = 0.5$  and  $c = 0.0$  (smoothing factor).

action to converge smoothly, it also allows sudden changes in actions to happen in response to changes in state or policy updates. The method retains the possibility to “jump” to a totally different action by suddenly changing the selected wire. This flexibility in the selection of actions represents another advantage of the wire fitting method since it allows actions to converge smoothly to optimal values while retaining the possibility to execute immediate changes in action-values if needed. As an example, consider the three wires shown in Figure 3.5, where **Wire 1** is associated with the selected action  $a_{\text{Max}}$ . Updating this wire with  $\Delta Q(s, a_{\text{Max}}) = -0.5$  and  $c = 0.0$  results in the three wires being adjusted as shown in Figure 3.7. Arrows show the update made to **Wire 1**, with **Wire 0** and **Wire 2** also being slightly adjusted. After the update,  $q_{\text{Max}}$  will be associated with **Wire 2**, allowing an abrupt change in the selected action (from  $a_{\text{Max}} = 0.5$  to  $a_{\text{Max}} \approx 0.8$ ) as described. The practical value of the update feature is that it makes it possible to adjust actions smoothly during the learning phase while retaining the capability to explore different areas of the action-space if the reward penalizes the current selected action. This feature has applications for HVAC control, because although we expect a temperature set point to converge to an optimal stable value during most of the hours of the day, there

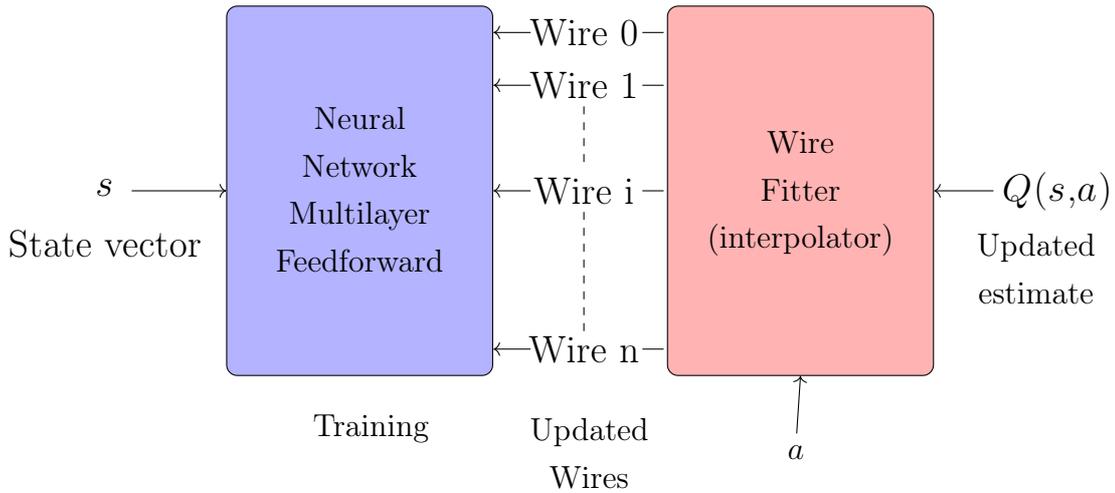


Figure 3.6: Wire fitted neural network training algorithm (adapted from [228]).

will be instances of time where the set point needs to change in order to avoid the occupant feeling uncomfortable or energy being wasted. Retaining the possibility to discretely select different temperatures for the set point range is an advantage for learning, because it reduces the convergence time needed to find the range of optimal temperature values.

## 3.2 Application Problems

Addressing the requirement that a smart BAS should adapt its operation according to the cost of energy and comfort level of its occupants, this section proposes two RL strategies for two different problems. The first problem, which is called *Bang-bang Heater*, presumes the heater can be turned *on* or *off* but does not have an interface for the BAS to set the temperature set point. The second problem, called the *Set Point Heater*, presumes the HVAC unit has a temperature set point control interface directly accessible by the BAS.

Both the *Bang-bang Heater* and *Set Point Heater* problems are applied to the TZ (environment) represented in Figure 3.8. In these settings we assume that an occupant uses a bi-modal interface to indicate to the BAS that s/he would like the temperature to either be increased (+) or decreased (-). We equate the interface to signaling “I am cold” or “I am hot”. An occupant can express dissatisfaction with the current temperature by pressing the “UP” or “DOWN” arrow on the thermostat. We expect that the occupant will sometimes be impatient and demand a bigger adjustment before

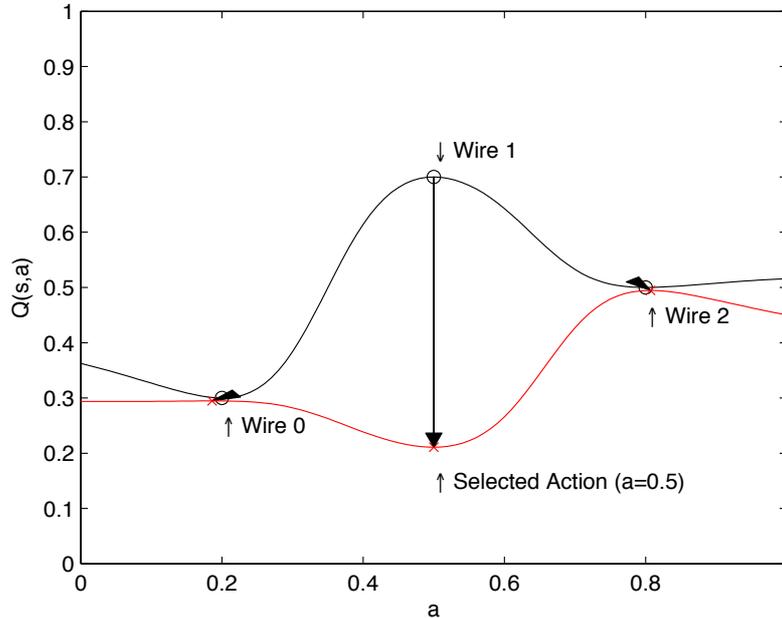


Figure 3.7: Updating the wires by  $\Delta Q(s, a) = -0.5$  (represented by the arrows) for the selected action  $a_{\text{Max}} = 0.5$ , using  $c = 0.0$ .

a control action has fully taken effect. We also expect that sometimes the occupant will be patient but will decide that the controller has not changed the temperature enough and consequently change it further. Tenancy may change, and the press of an arrow may also mean that a new standard of comfort is desired. With these two pieces of information and over repeated interactions, regardless of whether the TZ is vacant or not, the controller must adjust zonal temperature efficiently – to just the right temperature and with a maximum amount of energy savings achieved.

**Without loss of generality, we assume in this thesis that the outside temperature is lower than what is desirable if the TZ is occupied.** Therefore, in both problems, it is assumed that the occupant is either *comfortable* or *cold*. This implies that the controller’s operation is to turn *on* the heating only as frequently as it is necessary to make the occupant comfortable (*Bang-bang Heater*), or continuously set the temperature set points to appropriate values throughout the day (*Set Point Heater*). This control must be executed in a manner that minimizes the number of occupant signals and at the same time also minimizes the costs associated with heating. The BAS can neglect heating the TZ altogether when the TZ is unoccupied, but it may be required to “preheat” the zone to guarantee a comfortable environment

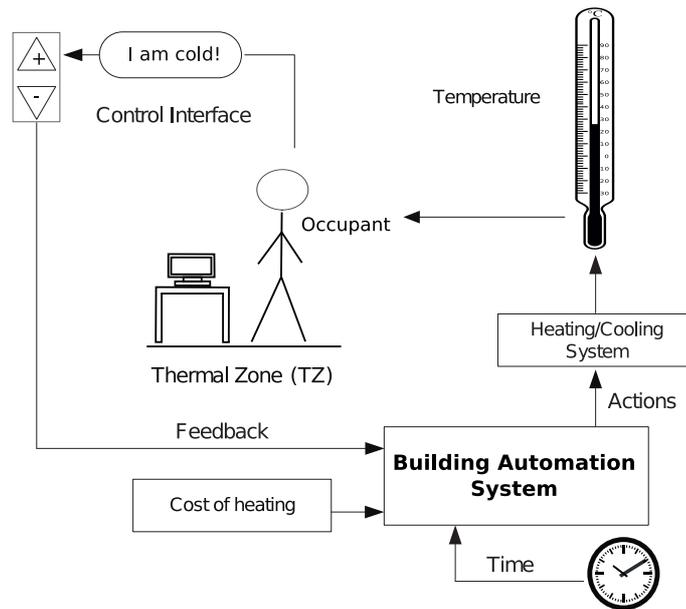


Figure 3.8: The occupant indicates when s/he is feeling cold. The smart Building Automation System must learn how to best respond to the occupant (by minimizing the required number of thermostat interactions) while minimizing the energy used for heating.

when it becomes occupied. Even when the occupant forgets to reset a temperature for zonal vacancy, the smart BAS must discover and exploit the opportunity made by the occupant’s absence or seek satisfaction with current environmental conditions – inferred through the lack of action over the control interface – for more efficient energy management.

### 3.2.1 The *Bang-bang Heater* Problem

The *Bang-bang Heater* presumes that a heating unit is controlled at a low level by a hysteresis controller that guarantees that the environment temperature will converge to a certain set point when the heater is *on*. That is, it can either be switched *on* or *off*, but it has no temperature range controllability. We assume that the occupant will adjust the temperature set point to the preferred comfort level and that this temperature is reached and maintained if the heater is left *on* for a certain amount of time. The *Bang-bang Heater* problem matches the common residential scenario. In most dwellings it is common to have electric oil heaters and other types of radiators for heating. These heaters, which are regulated to a certain heating level defined by

the desired temperature set point, represented by  $T_d$ , can be connected to a smart power switch that is capable of cutting off the power supply. This switch can be controlled by a home automation system with occupant feedback obtained through a smartphone application, as illustrated in Figure 3.9.

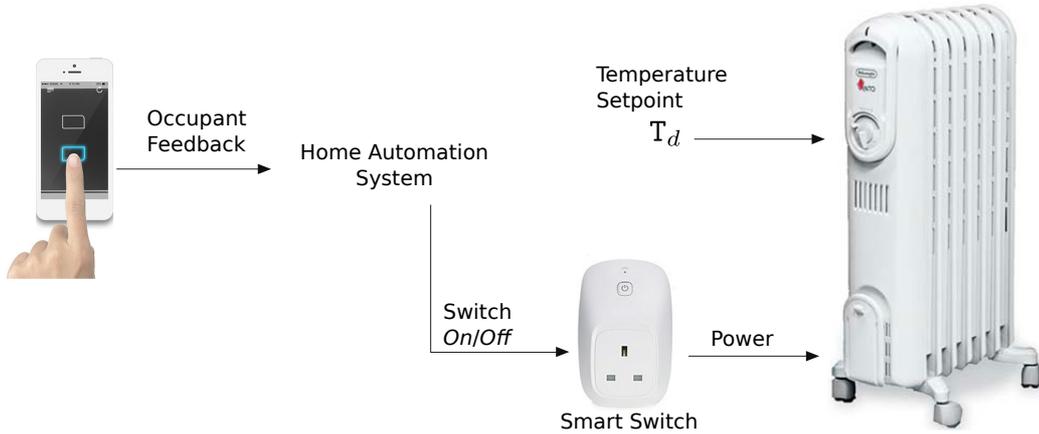


Figure 3.9: The *Bang-bang Heater* problem matches a common residential scenario where an electric heater can be controlled by the home automation system using a smart switch.

### States, Actions, and Q-values

The *Bang-bang Heater* heater problem is solved using standard Q-Learning with discrete states and actions. Therefore, the state is represented by a vector  $s = (t, h)$ , where  $t \in \{t_1, t_2, \dots, t_N\}$  represents a discrete instant of time counting from the beginning of the learning episode,  $t_N$  represents the corresponding time when the terminal state is reached (i.e., the amount of time the system has been in operation when the learning episode terminates), and  $h \in \{0, 1\}$  represents whether the heater is *on* ( $h = 1$ ) or *off* ( $h = 0$ ). To operate, we assume that the BAS observes and controls  $h$  and that the action set  $A = \{a \mid a \in \{0, 1\}\}$  includes the possibility to **Maintain** ( $a = 0$ ) or **Toggle** ( $a = 1$ ) the current state of the heater.

To evaluate the quality of each action in each state, the array of Q-values used by the Q-Learning algorithm are stored in the BAS's memory. For each instant of time, we have the four different state-action configurations shown in Table 3.1. This table is updated according to (3.2) and (3.3) in each learning iteration.

	HVAC State - $h$	Action - $a$
$Q_{off\_M}(t)$	<i>off</i>	<i>Maintain</i>
$Q_{on\_M}(t)$	<i>on</i>	<i>Maintain</i>
$Q_{off\_T}(t)$	<i>off</i>	<i>Toggle</i>
$Q_{on\_T}(t)$	<i>on</i>	<i>Toggle</i>

Table 3.1: Q-values associated with each HVAC state-action configuration for each instant of time.

### Reward Function

The RL BAS must learn to select its actions in order to minimize the energy used for heating and the number of times an occupant subsequently interacts with the thermostat. These accomplishments need to be reflected in the form of a reward. For the HVAC optimization problem, however, it is easier to invert the reward to a penalty (negative reward). **The BAS receives a higher or lower penalty if the user is uncomfortable or if excess energy is being used.** Whenever the heater is *off*, because we assume the outside temperature is too cold for the zone and we want to save energy, at every discrete point in time, no action by the occupant is considered a reward. When the occupant presses an arrow to increase the heat, the heater could already be *on* (without the zone being at the desired level of comfort) or the heater could be turned *off*. In both scenarios the pressing of the arrow is negative feedback. To include this feedback in the reward function, let  $interaction(t)$  represent a predicate on the occupant’s interaction during a certain time interval  $[t, t + 1]$ . The predicate evaluates to *true* (replaced by 1) if the occupant has acted upon the system and to *false* (replaced by 0) otherwise. Using this evaluation, the reward/penalty function is given by the following linear combination:

$$r_t = R(t, h) = -w_1 interaction(t) - w_2 h \quad (3.5)$$

where  $\{w_1, w_2 \in \mathbb{R} \mid w_1 + w_2 = 1\}$  are weights that regulate the trade-off between comfort and energy savings.

### 3.2.2 The *Set Point Heater* Problem

The *Set Point Heater* problem uses Q-Learning with continuous states and actions to determine the set point temperature at every instant of time. Just as it was

used in the *Bang-bang Heater* problem, the state of the *Set Point Heater* problem includes the current instant of time  $\{t \mid t \in \mathbb{R}, 0 \leq t < t_N\}$  from the finite learning period (corresponding to a 24-hour cycle). However, in the *Set Point Heater*, time is a continuous variable and it is the only variable necessary to represent the state of the environment. The temperature of the TZ, as opposed to the discrete state of the heater, was not considered as a state component. It is assumed, without loss of generality, that the low-level HVAC controller will heat the indoor temperature to the desired temperature set point, and that the HVAC controller will maintain the TZ at this temperature until further inputs are received.

### Actions and Wires

To obtain the selected action, the instant of time is mapped into wires by the NN block, as illustrated in Figure 3.10. After selecting  $a_{\text{Max}} \in [0, 1]$ , random noise is added for exploration. The thermostat set point temperature, defined by  $T_s$ , uses linear mapping to describe the relationship between the selected action and the following temperature interval:  $T_s \in [T_a, T_{\text{max}}]$ , where  $T_a$  is the outdoor ambient temperature and  $T_{\text{max}}$  is the upper bound temperature of the HVAC system. Temperature  $T_{\text{max}}$  should be set to a value such that the temperature operation interval includes a range of temperatures that guarantee comfort.

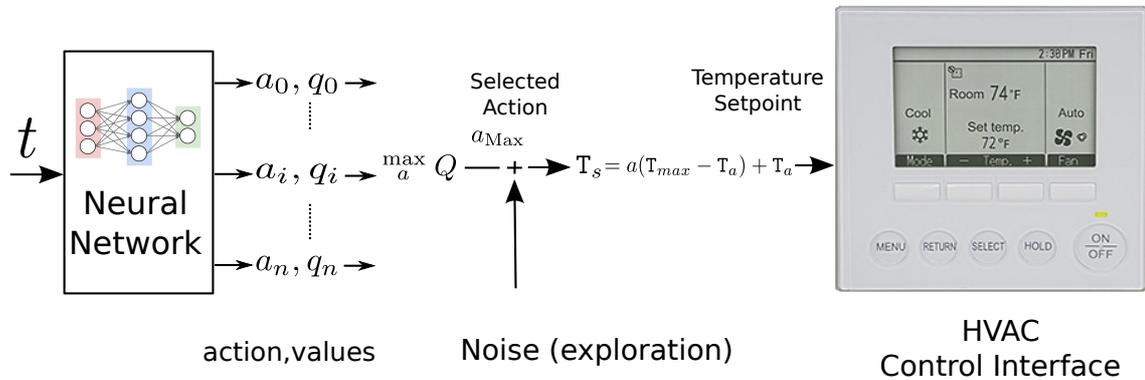


Figure 3.10: Selecting actions with the *Set Point Heater* using Q-Learning with continuous states and actions.

### Reward Function

Following the learning method described in Section 3.1.3, wires are updated at regular time learning intervals. For the *Set Point Heater* problem, there is a **heating cost**

proportional to the amount of heat supplied to the TZ. Heat transfer, governed by the laws of thermodynamics, is directly proportional to the temperature difference between the indoor and outdoor temperatures. The greater the difference, the greater the amount of thermal energy transferred across the boundary of the TZ. Therefore, after normalizing for the operating temperature range, the heating cost associated with the desired indoor set point temperature is given by the following expression:

$$hCost(T_s) = \frac{1}{T_{max} - T_a}(T_s - T_a) \quad (3.6)$$

For comfort performance, the reward function, as defined in the *Bang-bang Heater* example, only takes into consideration the feedback interaction from the occupant. The *Bang-bang Heater* has no direct control over the TZ temperature. On the other hand, the *Set Point Heater* has the direct capability to define the set point temperature. Therefore, when the occupant interacts with the thermostat, the reward function for this application must use a **comfort penalty** function as a heuristic to guide the search toward comfortable temperatures. This function, given by (3.7), is inversely proportional to  $T_s$  and has a maximum value when there is no heat supply to the TZ, which corresponds to the situation where the indoor temperature is equal to the outdoor temperature as shown in Figure 3.11.

$$cPenalty(T_s) = 1 - hCost(T_s) \quad (3.7)$$

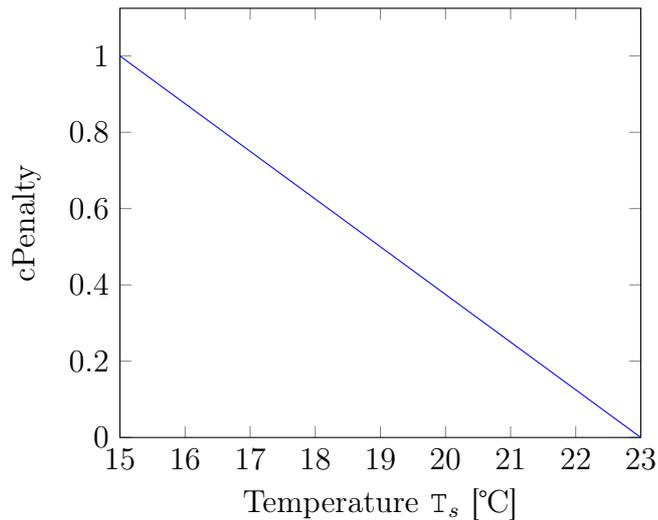


Figure 3.11: Comfort penalty function (using  $T_a = 15$  °C and  $T_{max} = 23$  °C).

The resulting reward function is given by the following linear combination:

$$r_t = R(t, T_s) = -w_1 \text{interaction}(t) c\text{Penalty}(T_s) - w_2 h\text{Cost}(T_s) = -w_1 \text{interaction}(t) - [w_2 - w_1 \text{interaction}(t)] h\text{Cost}(T_s) \quad (3.8)$$

### 3.3 Discussion

The methodologies presented in this chapter have included some of the following important facts that must be considered and discussed.

- **Interaction with the occupants.** Since the BAS learns by having the occupant interact with the thermostat, there will be learning periods when the occupant becomes uncomfortable with the selected temperatures. Depending on the learning period, it is this researcher's belief that this might not be a serious limitation. The BAS learns by observing habits, and the learning and optimization strategy is solely based on feedback obtained through observation of these habits. Since these observations take some time to acquire, a certain degree of system adaptation is expected from the occupants. Humans also need time to learn and adjust to the preferences and habits of their coworkers, family, and friends. Therefore, the BAS operates just as humans do and this adaptation period is natural and normal.

Another consideration about the interaction with occupants is that building environments are not stationary and learned policies may become irrelevant when occupants change their behaviors from time to time. Since humans are not always predictable, when behavioral patterns change, the BAS has to readjust by learning new policies. These changes impact performance in terms of finding optimal stable policies. The BAS learns to explore specific behavioral patterns, which is not different from humans adjusting to humans. It is known that the better we can predict someone's behavior, the better we can adjust to his/her habits. Therefore if the occupant has regular habits, the BAS can maximize the efficiency of its policy by fitting the policy to those habits. If the occupant changes his/her behavior frequently and becomes unpredictable, then for the sake of comfort, energy savings cannot be fully accomplished during the occupied periods (unless the environment becomes more observable by adding the capability to, e.g., predict the mood of the occupant). When the occupant is unpredictable, savings can only be achieved by learning how to

minimize the use of energy during time intervals when the occupant usually does not interact with system, or when s/he is normally out of the office. However there will always be situations that will disrupt the daily routine. The occupant can have holidays scheduled for unexpected days or may perhaps be running late. To increase system performance, some extra steps can be taken by considering more contextual information in the decision-making process (e.g., the occupant’s schedule or location), making some of these occupancy changes more “predictable.” If the BAS is able to predict that the occupant will not be following the usual schedule, then it can readjust the control policy according to a different strategy, defined for each specific situation.

- **Multiple occupants.** When considering more dynamic environments, such as rooms shared by several students, the BAS must optimize performance for multiple occupants by taking into account the fact that there may be different schedules and thermal preferences. Since our RL algorithm relies on feedback obtained through actuation on a button interface, the BAS merely learns how to minimize complaints. Therefore, the occupant that interacts more with the HVAC comes out as the winner and imposes his/her comfort preferences on the comfort preferences of other occupants. For energy savings, in the worst case, the system eventually learns how to minimize energy during time intervals where nobody is interacting with it. To optimize for multiple occupants, having access to more information can be beneficial for decision-making support. With more observability over the environment (by using, e.g., a smart video based identification system capable of identifying each occupant or by obtaining a user’s identification information from the HVAC smartphone application), the boundaries of comfort versus energy optimization can be further explored by considering the set of learned parameters for each specific tenant.
- **Explainability.** In order to deploy efficient HMIs, as described in Section 2.1.3, it is important to have the BAS explain the reason for its actions. With RL, policies are encoded in Q-Value tables or in NN weights. Therefore, presenting an explanation to a human operator is not a straightforward process. It is very difficult to extract semantic information from these policies. For a friendly user interface it would be desirable the possibility to have an hypothetical dialogue between a human and the BAS where the human asks the building, for example, why the HVAC has been turned *off*, and the BAS responds with something along

the lines of, “*I was expecting that you would be attending the group meeting that is currently taking place in room 11*”. To have such a dialogue, the BAS must be able to reason at a symbolic level, which is not possible with RL.

- **Scalability.** The solution of using RL with the current state representation does not scale well and cannot cope with the complexity of the building environment. In most cases, there is additional decision-making support information that spans out of the “room domain” that could be considered in the optimization problem. As an example attending to this fact, consider the situation when the occupant has an appointment on his/her schedule for a different physical location in the building. With the BAS predicting that there will be no interaction in the office during that time period, the optimal strategy for energy savings would be to leave the HVAC settings in a low-power state. With the current RL approach, adding new information to the optimization problem is not a simple matter, and this represents a serious limitation for AmI.

The HVAC control problem cannot be viewed as a simple state-space search strategy. To cope with the complexity of the building environment, the solution must be based on a systems thinking approach [230, 231]. The control problem cannot be solved by dealing with parts of the problem in isolation using closed and predefined state-space representations. It must be solved in concert with many other modules of the BAS that interact to produce behavior. Therefore, not only should the operation of the BAS be analytically partitioned into smaller components for simplification, it should also take into consideration that everything is systemic, i.e., everything interacts, affects, and is affected by the things around it. However, notwithstanding the fact that occupants in the future will probably expect this level of integration, addressing this requirement in an algorithm is a very complex problem. In many cases, the amount of information that is included in the decision-making process means that the optimization problem is computationally intractable. Therefore, only a subset of this information can be used in the optimization problem, and appropriate modeling paradigms are necessary to sort and organize what is relevant to use at each decision-making step. These modeling paradigms are a separate topic of research and still need to be developed for SBs.

## 3.4 Summary

This chapter presented a brief overview of reinforcement learning theory with a description of how it can be used to optimize the operation of the HVAC system. In this description the environment is the thermal zone, and the building automation system is a software agent capable of obtaining feedback on the control actions it executes to change the thermal zone temperature. These actions are expected to lower the energy cost associated with heating and cooling, while minimizing the number of occupant interaction signals in the form of “I am cold” or “I am hot” that are obtained through the HVAC user control interface. The chapter described the application of a discrete and a continuous Q-Learning based supervisory control approach for the building automation system, which actively learns how to schedule the operation of the HVAC system considering two different problems of low-level heating control: *Bang-bang Heater* and *Set Point Heater*. The *Bang-bang Heater* problem assumes that zonal heating is controlled by alternately switching the activation state of a heater. This problem is solved with straightforward Q-Learning using discrete states and actions. The *Set Point Heater* problem, on the other hand, assumes that the building automation system is capable of learning how to set temperature setpoints. The actions and states are continuous values: the zonal temperature set point is controlled in minutely small measures. Therefore, a continuous state and action Q-Learning algorithm was selected and customized to solve this problem.

The chapter finalizes with a discussion on the presented reinforcement learning methodologies. This discussion included the fact that reinforcement learning algorithms present limitations when it comes to explaining a system’s actions to humans. Moreover, the reinforcement learning problem becomes computationally expensive to solve as more states are added to represent the building environment and the interaction with its occupants with more accuracy. Therefore, the research efforts for this dissertation became focused on finding alternative solutions to represent the state of the building environment, taking into account that this environment may include an intractably large set of variables. Thereupon, the research problem changed towards the goal of finding accurate models for the building environment which can be used, among other things, for building simulation, model predictive control, and for synthesizing building automation plans that are explicable and predictable to humans.



# 4

## Context-based Thermodynamic Models

*“Minds inhabit environments which act on them and on which they in turn react.” – William James (Philosopher and psychologist).*

Many physical processes in buildings can be modeled by continuous dynamics. However, there are other building systems, such as TZs, that can exhibit both continuous and discrete dynamic behavior. Drawing a parallel with the context-based reasoning framework discussed in Section 2.3, this chapter describes a context-based framework to represent these systems.

### 4.1 Operational Semantics of Context-based Models

A context-based thermodynamic model includes all the different dynamic models that are available to describe the thermal behavior of a TZ. Therefore, a descriptive framework is required to capture the transitions between these models, with the different types of continuous and discrete dynamics. For this purpose, we have picked a suitable modeling formalism called *hybrid automata* [202, 232]. *Context* of a TZ, as defined in this dissertation, is equated to a **discrete configuration**  $l \in L$  (aka

*control state* or mode [233]) associated with and representing a particular thermal behavior of the TZ. A model for the operation of a TZ can include several contexts, between which the system evolves in an event-driven manner. In addition to this set of discrete contexts, there is a **continuous state**  $x \in X \subseteq \mathbb{R}^n$  containing the real-valued time-driven variables (temperatures) that “flow” in continuous time according to a set of different models. Each of these models, associated with one or more discrete states, are described using ODEs. The TZ’s **full state**  $\rho = (l, x) \in L \times X$  (also sometimes called *configuration* or just simply *state*) is defined, at a certain instant in time, by the discrete and continuous part of the model. An *open hybrid automaton/context-based model* is a dynamical system, with inputs and outputs, that describes how  $\rho$  evolves in  $L$  and  $X$ . It is given by the following definition (adapted from Lygeros [202]):

**Definition 1 (Context-based model)**

A context-based model  $\mathcal{M}$  is a collection  $\mathcal{M} = (L, X, U, Y, Init, m, m_Y, D, E, G, Rst)$ , where:

- $L = \{l_1, l_2, \dots, l_s\}$  with  $s \in \mathbb{N}$ , is a finite collection of **discrete system contexts**.
- $X$  is a finite collection of **continuous state variables**.
- $U$  is a finite collection of **input variables**. Assuming  $U = U_D \cup U_C$ , where  $U_D$  contains discrete variables and  $U_C$  contains continuous variables.
- $Y$  is a finite collection of **output variables**. Assuming  $Y = Y_D \cup Y_C$ , where  $Y_D$  contains discrete variables and  $Y_C$  contains continuous variables.
- $Init \subseteq L \times X$  is a set of **initial states**.
- $m : L \times X \times U \rightarrow \mathbb{R}^n$  is a **vector field** that characterizes the continuous dynamics in the domain of the corresponding context, which evolve in continuous time.
- $m_Y : L \times X \times U \rightarrow Y$  is a **vector field** associated with the output of the open hybrid automaton.
- $D : L \rightarrow 2^{X \times U}$  assigns to each context  $l \in L$  a **domain/invariant set**.
- $E \subseteq L \times L$  is a collection of **discrete transitions/edges** between contexts.
- $G : E \rightarrow 2^{X \times U}$  assigns to each edge  $e = (l, l') \in E$  a **guard**.

- $Rst : E \times X \times U \rightarrow 2^X$  is a **reset map** that assigns to each  $e = (l, l') \in E$  and  $x \in X$  a reset location.

with  $2^X$  denoting the power set (set of all subsets) of  $X$ .

For each *context*  $l \in L$ , there is an associated set of continuous states assigned by function  $D(l) \subseteq \mathbb{R}^n$ . Starting from an initial state  $\rho_0 = (l_1, x_0) \in Init$ , the continuous state  $x$  flows according to the model that characterizes the continuous dynamics in the domain of  $l_1$ :

$$\begin{aligned}\dot{x} &= m(l_1, x, u) \\ x(0) &= x_0\end{aligned}$$

while the *context*  $l = l_1$  remains constant. Continuous evolution can go on as long as  $x$  remains in the domain  $D(l_1)$ . Contexts are connected by discrete labeled transitions (edges) with guards and effects. If at some point the continuous state  $x$  reaches the guard  $G(l_1, l_i) \subseteq \mathbb{R}^n$  of some edge  $(l_1, l_i) \in E$ , **context may change** value to  $l_i$ . With this transition, the continuous state is reset to some value given by  $Rst(l_1, l_i, x, u)$  and continuous evolution resumes, according to a new model associated with the new *context*  $l_i$ . In this dissertation, *context* transitions are simulated assuming *must transitions*, i.e., transitions are immediately taken when guard expressions become true. This is a practical refinement of the hybrid system formalism, where guard expressions just act as enablers and do not force transitions to be made.

By comparing the hybrid system with the context-based reasoning framework discussed in Section 2.3, it can be assumed that the *declarative knowledge* equates to the vector field, and the *transitional knowledge* is represented by the domain, transition edges, and guards.

## 4.2 Example (The Hysteresis Controller)

As an application example of a context-based model, consider the operation of the heating unit used in the *Bang-bang Heater* problem described in Section 3.2. This section simulates the behavior of the hysteresis controller that controls the heating unit, and the evolution of the indoor temperature using a simple thermodynamic model to describe the TZ.

### 4.2.1 Thermodynamic Model

For this example, the TZ is assumed to be isolated from the outdoor ambient temperature through a single wall, as illustrated in Figure 4.1. When the BAS activates the

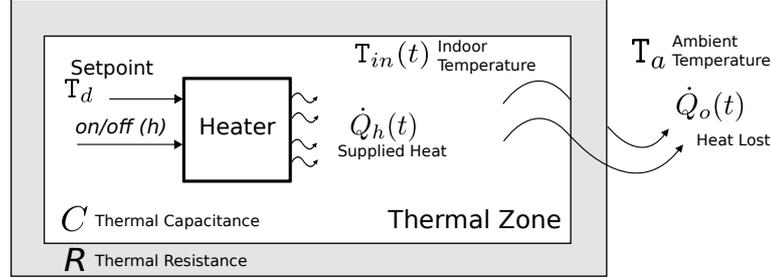


Figure 4.1: The *Bang-bang Heater* in a single thermal zone.

heating controller, by setting the input control signal  $h = 1$  (*on*), a certain amount of thermal energy, represented by an input  $Q_h \in \mathbb{R}^+$ , is supplied to the interior of the TZ to increase the indoor temperature, denoted by  $T_{in}$ . Depending on the temperature difference between the indoor temperature and outdoor ambient temperature ( $T_a$ ), a certain amount of heat is lost through the boundaries of the TZ, given by

$$Q_o(t) = \frac{1}{R} [T_a(t) - T_{in}(t)] \quad (4.1)$$

where  $R$  represents the *thermal resistance* associated with those boundaries. Depending on the amount of heat that remains in the TZ, the indoor temperature evolves as a function of time. To describe this evolution, we can use a simplified thermodynamic model given by

$$\dot{T}_{in}(t) = \frac{1}{C} [Q_h(t) - Q_o(t)] \quad (4.2)$$

where  $C$  represents the total *thermal capacitance* of the TZ (including the walls and the interior air). By replacing (4.1) in (4.2), we get the corresponding first order ODE:

$$\dot{T}_{in}(t) + \frac{1}{RC} T_{in}(t) = \frac{1}{C} Q_h(t) + \frac{1}{RC} T_a(t) \quad (4.3)$$

that represents the continuous-time flow of the indoor temperature.

### 4.2.2 Heater Operation

In order to regulate the TZ temperature to the desired set point, defined by  $T_d$ , the heater includes a hysteresis controller that is activated/deactivated by the external

## 4.2. EXAMPLE (THE HYSTERESIS CONTROLLER)

control input  $h$ . When activated, this controller operates by repeatedly interchanging the *on/off* power state of the heating element (electric coil) in order to maintain the temperature difference between  $T_{in}$  and  $T_d$ , within certain boundaries. The controller switches the heat supply *on* if the indoor temperature goes below a certain temperature threshold, defined by  $T_{min}$ , and switches it back *off* when the temperature reaches above another defined threshold, defined by  $T_{max}$ . Figure 4.2 illustrates this hysteresis behavior with arrows representing the direction of the temperature evolution and the state commutation.

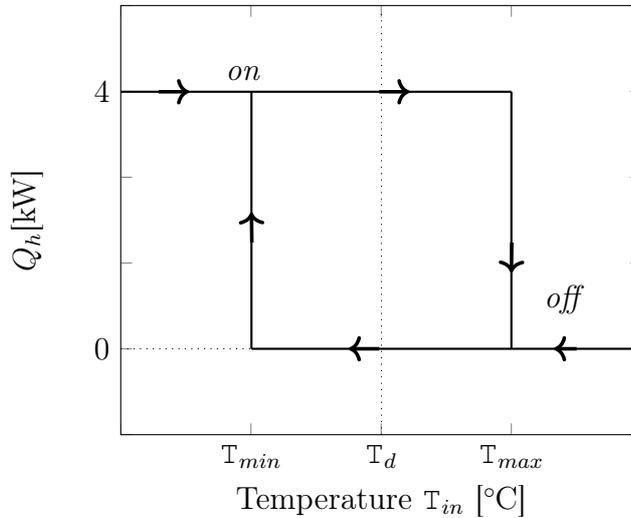


Figure 4.2: Hysteresis behavior of the *Bang-bang Heater*.

### Context-Based Model

The thermal behavior of this heater and TZ can be described formally by a context-based model with:

- $L = \{l_1, l_2\}$ . A collection of contexts with the following names: *Heater off* ( $l_1$ ) and *Heater on* ( $l_2$ ).
- $X = T_{in} \in \mathbb{R}$ . The continuous state, defined by the temperature in the TZ.
- $U = U_D \cup U_C = \{h\} \cup \{T_a, Q_h, T_{min}, T_{max}\}$ . The input variables that define the activation state of the heater, ambient temperature, heating energy, and the hysteresis controller parameters.

- Vector field given by (4.3), depending on the state of the heater:

$$\dot{T}_{in}(t) = m(l, T_{in}, u) = \begin{cases} -\frac{1}{RC}T_{in}(t) + \frac{1}{RC}T_a(t), & \text{for } l = l_1 \\ -\frac{1}{RC}T_{in}(t) + \frac{1}{C}Q_h(t) + \frac{1}{RC}T_a(t) & \text{for } l = l_2 \end{cases}$$

with  $u \in U$ .

- $Y=X$ . The output variable is given by the state variable.
- $Init = \{l_1\} \times \{T_{in} = 15\}$  to start with the heater *off*, and  $T_{in} = 15$  °C.
- The domain associated with each *context*:

$$D(l_1) = \{T_{in} \geq T_{min}\} \times \{h = 0\}$$

$$D(l_2) = \{T_{in} \leq T_{max}\} \times \{h = 1\}$$

- $E = \{(l_1, l_2), (l_2, l_1)\}$  edges between turning the heater *on/off* respectively.
- A guard to switch the heater *on* if the heater is enabled, and  $T_{in}$  reaches below  $T_{min}$ , and a guard to switch the heater *off*, if  $T_{in}$  reaches above  $T_{max}$  or if the heater is disabled, given respectively by:

$$G(l_1, l_2) = \{T_{in} < T_{min}\} \times \{h = 1\}$$

$$G(l_2, l_1) = \{T_{in} > T_{max}\} \cup \{h = 0\}$$

- $Rst(l_1, l_2, T_{in}, u) = Rst(l_2, l_1, T_{in}, u) = \{T_{in}\}$ . Because the TZ temperature does not change instantaneously.

The resulting hybrid automaton is represented in Figure 4.3 as a directed graph, where the vertices of the graph are locations from  $L$  and the (labeled) edges are the transitions from  $E$ . The initial state is marked by an incoming edge without a source.

### 4.3 Hybrid Time Sets and Executions

In order to characterize the evolution of the state of a context-based model, consider a set that contains the **continuous intervals**, over which continuous evolution takes place, and the distinguished **discrete points in time when discrete transitions happen**. Such a set is called a **hybrid time set** (adapted from Lygeros [202]).

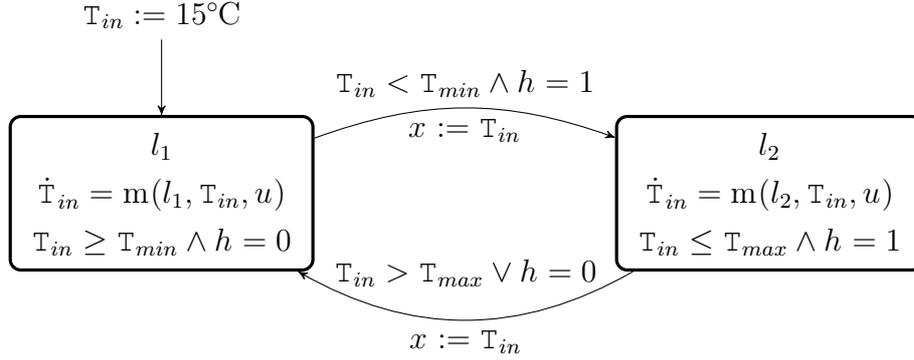


Figure 4.3: Graphical representation of the heating system.

**Definition 2 (Hybrid time set [202])** A hybrid time set is a sequence of intervals  $\tau = \{I_1, I_2, \dots, I_N\} = \{I_i\}_{i=1}^N$ , finite or infinite (i.e.,  $N = \infty$  is allowed), such that:

- $I_i = [\tau_i, \tau'_i]$  for all  $i < N$ ;
- if  $N < \infty$  then either  $I_N = [\tau_N, \tau'_N]$  or  $I_N = [\tau_N, \tau'_N[$ ;
- $\tau_i \leq \tau'_i = \tau_{i+1}$  for all  $i$ .

Discrete transitions are assumed to be instantaneous with  $\tau'_i = \tau_{i+1}$ . Figure 4.4 illustrates an example of a hybrid time set, labeled with the names of the corresponding contexts, with the automaton jumping from  $l_1$  to  $l_2$ , and then immediately to  $l_3$  and so on.

Hybrid time sets are used to define the time horizon over which  $\mathcal{M}$  evolves, i.e., the time horizon associated with the execution/evolution of the full state  $\rho(t) = (l(t), x(t))$  with  $t \in \tau$ . Considering *context* and the continuous state, we have the following definitions:

**Definition 3 (Hybrid trajectory [202])** A hybrid trajectory is a triple  $(\tau, l, x)$  consisting of a hybrid time set  $\tau = \{I_i\}_{i=1}^N$  and two sequences of functions,  $l = \{l_i(\cdot)\}_1^N$  and  $x = \{x_i(\cdot)\}_1^N$  with  $l_i(\cdot) : I_i \rightarrow L$  and  $x(\cdot) : I_i \rightarrow X$ .

**Definition 4 (Execution of a context-based model (adapted from [202]))** An execution  $\chi$  of a context-based model  $\mathcal{M}$  is a collection  $\chi = (\tau, l, x, u, y)$  with  $\tau \in I$ ,  $l : \tau \rightarrow L$ ,  $x : \tau \rightarrow X$ ,  $u : \tau \rightarrow U$ , and  $y \in Y$  satisfying the following conditions:

- Initial condition  $(l(\tau_1), x(\tau_1)) \in \text{Init}$ ;
- Continuous evolution: for all  $i$ , with  $\tau_i < \tau'_i$ ,

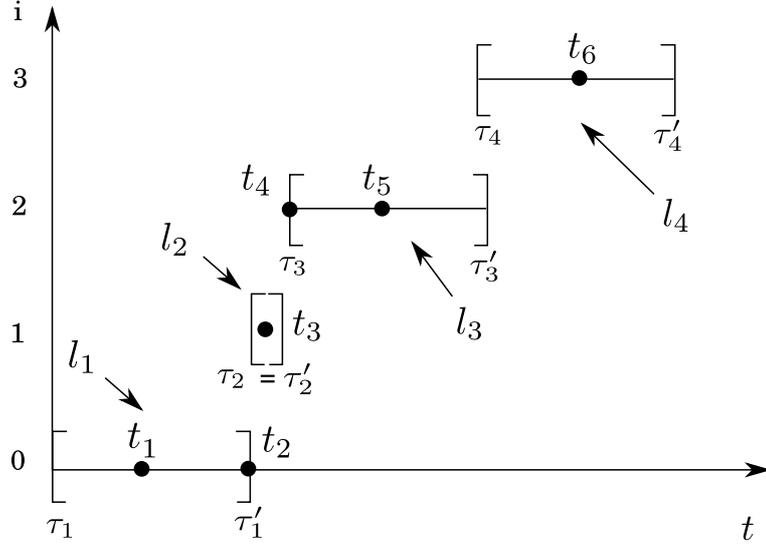


Figure 4.4: Graph showing a hybrid time set  $\tau = \{[\tau_i, \tau'_i]\}_{i=1}^4$ , as a sequence of  $i$  time intervals (adapted from [202]).

1.  $l_i(\cdot) : I_i \rightarrow L$  is constant over  $t \in I_i$ , i.e.  $l_i(t) = l_i(\tau_i)$  for all  $t \in I_i$ ;
2.  $x, u$  and  $y$  are continuous over  $[\tau_i, \tau'_i]$ ;
3.  $x_i(\cdot) : I_i \rightarrow X$  is the solution to the differential equation,

$$\dot{x}_i = m(l_i(t), x_i(t), u(t))$$

over  $I_i$ , starting at  $x_i(\tau_i)$ ; and,

4. for all  $t \in [\tau_i, \tau'_i[$ ,  $(x_i(t), u(t)) \in D(l_i(t))$ .

- *Discrete evolution:* for all  $i$ ,

either  $(l(\tau'_i), x(\tau'_i)) = (l(\tau_{i+1}), x(\tau_{i+1}))$ , or

$e_i = (l(\tau'_i), l(\tau_{i+1})) \in E$ ,

$(x(\tau'_i), u(\tau'_i)) \in G(e_i)$

$x(\tau_{i+1}) \in Rst(e_i, x(\tau'_i), u(\tau'_i))$

- *Output evolution:* for all  $t \in \tau$ ,  $y(t) = m_Y(l(t), x(t), u(t))$ .

### Example (Execution of the Hysteresis Controller)

Consider the example of the hysteresis controller given in Section 4.2. Algorithm 1 shows the simulation program that executes the context-based model using a specific

set of model-parameters. In this example, the heater is enabled ( $h = 1$ ) and inputs  $\{T_a, Q_h, T_{min}, T_{max}\}$  are constant values. Figure 4.5 shows the resulting execution with the continuous evolution of  $T_{in}$ , for a total duration of 7200 s divided between the hybrid time set  $\{I_0, \dots, I_{13}\}$ , as the *Bang-bang Heater* switches between  $l_1$  and  $l_2$ . The *Bang-bang Heater* maintains  $T_{in}$  within the pre-set boundaries of the desired temperature set point  $T_d = 22^\circ\text{C}$ .

---

**Algorithm 1:** Simulation program for the hysteresis controller.

---

```

N ← 7200 // 2-hour simulation
∂t ← 1 // time-step (s)
Ta ← 15 // Ambient temperature (°C)
Qh ← 4000 // Supplied heat (W)
R ← 0.0023 // Thermal resistance (K W-1)
C ← 300000 // Thermal capacitance (J K-1)
h ← 1 // Heater input enable
l(1) ← l1 // Initial context
Tin(1) ← 15 // Initial temperature (°C)
Tmax = 23.0 // Maximum Guard temperature (°C)
Tmin = 21.0 // Minimum Guard temperature (°C)
begin
  for each t ← 2 to N do
    if l(t - 1) == 1 then
      | ∂Tin = - $\frac{1}{RC}$ Tin(t - 1) +  $\frac{1}{RC}$ Ta
    else
      | ∂Tin = - $\frac{1}{RC}$ Tin(t - 1) +  $\frac{1}{C}$ Qh +  $\frac{1}{RC}$ Ta

      // Update Tin
      Tin(t) = Tin(t - 1) + ∂Tin * ∂t

      // Check Guards and update context
      if Tin(t) > Tmax ∨ h = 0 then l(t) = 1;
      else if Tin < Tmin ∧ h = 1 then l(t) = 2;
    end
  end
end

```

---

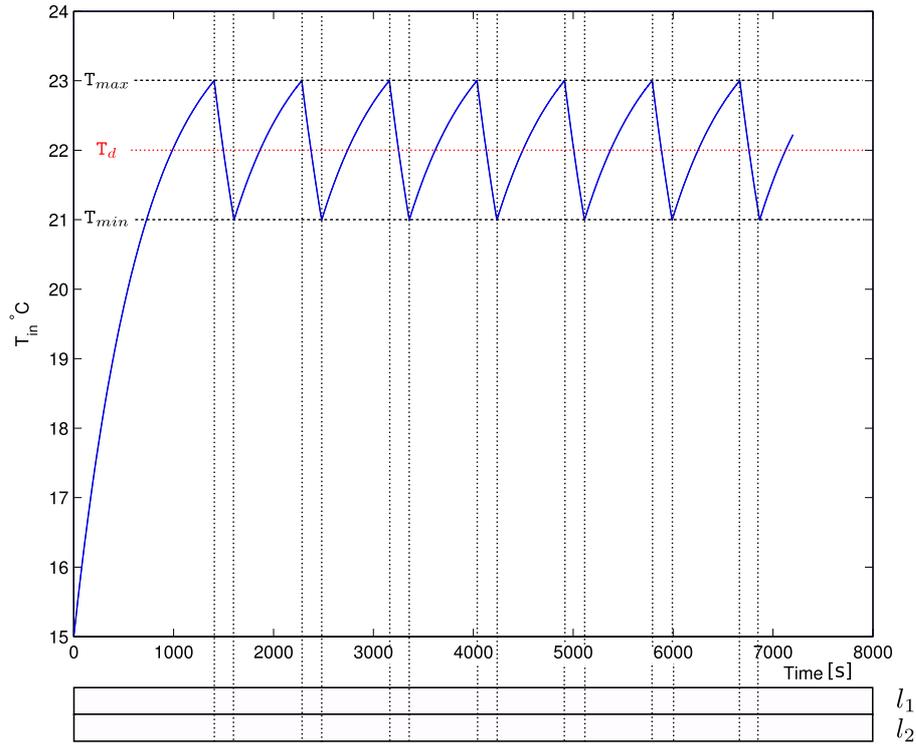


Figure 4.5: Execution of the heating system: continuous evolution of  $T_{in}$ , as the *Bang-bang Heater* switches between discrete states  $l_1$  and  $l_2$ .

## 4.4 Model Formulation for a Thermal Zone

The thermodynamic state of a multi-zone building consists of temperature variables associated with many subsystems that are geographically distributed. Each subsystem corresponds to a TZ and the thermodynamic interactions between adjacent zones occur mainly due to conduction (heat transfer through a medium) or convection (heat transfer between two different media) [234, 235]. The average temperature of a TZ ( $T_{in}$ ) evolves according to both heat and mass transfer laws with different modes for heat transfer, as illustrated in Figure 4.6. These modes include the following: (1) heat exchanges through zone surfaces that are in contact with the ambient temperature ( $T_a$ ), such as walls, roofs, doors, floors, windows and shades; (2) air exchanges by the HVAC system ventilation supply, natural ventilation, inter-zonal air-flows, infiltration and exfiltration; (3) solar gains represented by the rate of solar radiation acting on the building exterior surface, and transmitted through windows denoted, respectively, by  $Q_s$  and  $Q_{sw}$ ; (4) internal gains, denoted by  $Q_i$ , representing the rate of heat generated

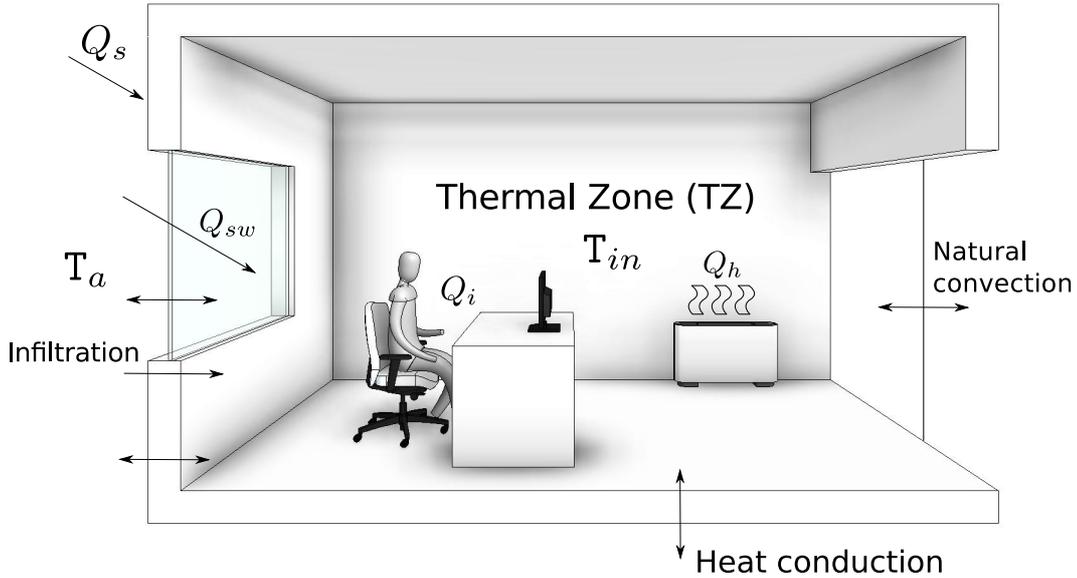


Figure 4.6: An illustration of different modes for heat transfer of a thermal zone. Convection (natural ventilation); solar gains through walls  $Q_s$  and windows  $Q_{sw}$ ; conduction through zone surfaces, infiltration, interior gains  $Q_i$ , and hydronic/electric heating  $Q_h$ . The variables  $T_a$  and  $T_{in}$  denote the ambient and zone air temperatures.

by occupants, electrical appliances, lights, computers, etc.; and (5) heat generated by heating equipment ( $Q_h$ ). All the combined modes for heat transfer present a complex interaction and TZ thermodynamic models can have different levels of complexity. Complex simulation strategies of heat transfer include computational fluid dynamic models that take into account the conservation equation of mass and energy to calculate, with high spatial resolution, the heat-energy transfer through building materials and ventilation [94, 107]. These simulations are computationally intensive and therefore inappropriate for online optimization-based control schemes such as MPC. To get faster responses, simpler models can be employed. These models, based on a thermal network assumption, are generally obtained after the reduction of more detailed and complex models. In the thermal network model, nodes represent temperatures, edges represent thermal paths, and nodes influence each other according to energy balance equations. In building simulation, nodes represent the average temperatures of the TZs assuming that the air volume in each zone is well mixed with a uniform temperature [236].

### 4.4.1 Full-Scale Lumped RC Thermal Model

Lumped RC networks are commonly used for constructing reduced-order thermal network models to describe the mono-dimensional (1D) heat transfers between network nodes [115,123,124,127,128,237,238]. Heat transfer by conduction between two media is proportional to the temperature difference between the two media involved, and inversely proportional to the resistance of the material layer(s) that separates them. Heat transfer by convection can also be approximated as a direct proportionality between the surface temperature of a media, and the temperature of the surrounding air [239]. The ability that materials (including interior air and materials for building walls, windows, furniture, etc.) have to accumulate heat is modeled with capacitors. Most building walls consist of several homogeneous layers of different materials. Depending on the complexity of the thermodynamic model, building walls can be described with more or less detail depending on the number of RC components used in the model. Fraisse et al. [240] compare different RC models for multi-layer walls and demonstrate that the thermophysical characteristics of multi-layer walls can be modeled using 3 resistors and 4 capacitors (3R4C). This model is sufficient to capture the conductive transfers between two TZs separated by a single wall, if the temperature distribution within the walls is not necessary (see more in Appendix A). Heat transfers through the TZ envelope can be modeled by different paths, depending on the number of surface elements. Figure 4.7 illustrates a RC building envelope model with heat transfer through walls, windows, roofs and floors. For the sake of simplicity, it is assumed that the walls and roof are composed of the same construction materials, and different surface elements are combined to construct the full-scale model of the TZ. The entire surface area of these elements is represented by a single 3R4C network with exterior and interior convective resistances  $R_{ext}$  and  $R_{int}$ , respectively.

The building's floor is composed by a single-layer material and is modeled with a 2R1C network and convective resistance  $R_{gint}$ , exposed to ground temperature  $T_g$ . Windows are modeled using a resistance  $R_W$  and, in some cases, shading with operable coverings like drapes, blinds, screens or pull-down shades, represented by variable resistance  $R_{WS}$ , can be used to reduce the amount of solar radiation entering the windows to reduce daylight glare. They can also be used to reduce heat loss through windows (movable insulation) using, for example, aluminium roller shutters/shades with thermal insulation.

The full-scale RC model is influenced by several discrete conditions. For this purpose, consider the discrete sets Heat, WS, W, D as containing the respective heat-

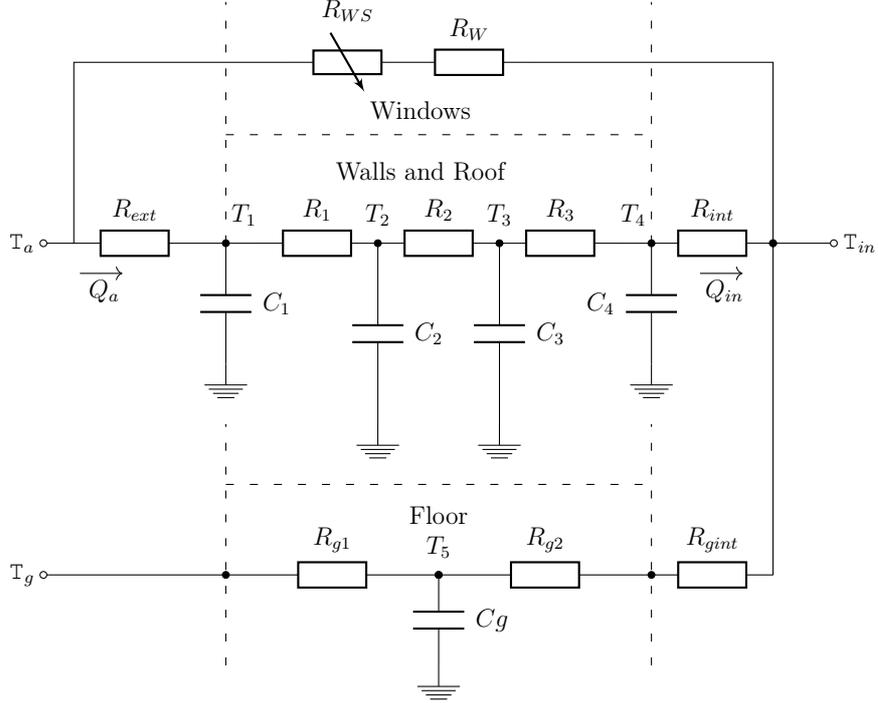


Figure 4.7: RC building envelope model composed by a 3R4C network (wall and roof), 2R1C network (floor), and a thermal path through windows and shutters/shades with thermal insulation.

ing levels, activation state for window shades, and opening factors (fraction of the opened area) for windows and doors. Consider a finite discrete set of ventilation levels  $\dot{M}_v \subseteq \mathbb{R}$  associated with air that passes through openings (windows and doors). Figure 4.8 shows a **full-scale RC model** of the entire TZ with several different context-dependent inputs. In particular:

- **Solar radiation.** Solar heat gains are an important part of free heat in a building. Heat inputs from solar radiation enter the building envelope through an effective area defined by  $A_e$  and, if shades are not active, through the effective window area defined by  $A_{wind}$ . The total transmitted solar radiation rate through windows acts directly on the building mass, and heat is transferred to the air inside the TZ, which has an associated capacitance represented by  $C_Z$ . Heat transmission to zone air is modeled using a transmission function  $f_{sw}(Q_{sw})$ .
- **Space heaters.** Controllable space heaters feed heat into the TZ according to a heat level  $h \in \text{Heat}$  and a maximum heating power defined by constant

$Q_h$ . This can be modeled by a current source  $hQ_h$ , a heat capacity  $C_h$  associated with the heater, and thermal resistance  $R_{ih}$  which is connected to the interior temperature node through a switch that is *on* if  $h \geq 0$ .

- **Internal gains.** The presence of people and the use of computers and other office equipment generate internal heat gains. The number of occupants is given by  $o$ , and heat gains are assumed to exist if there is at least one tenant occupying the TZ ( $o \geq 0$ ). It is assumed that each person has the same heat gain defined by constant  $Q_i$ .
- **Ventilation.** Airflow is caused by pressure differences and thermal buoyancy. They are influenced by the distribution of openings in the building shell, openings between rooms, and actions of occupants. Heat exchange by natural ventilation is assumed to be linear with the temperature difference between indoor and outdoor air. This approximation holds for low wind speed, but it is well known that for high wind speed, the natural ventilation of buildings becomes non-linear with both wind speed and direction. The rate of energy transferred between zones by means of airflow, on a per unit of temperature difference between them, is given by the expression  $U_{air} = \dot{m}_v C_{air}$ , where  $\dot{m}_v \in \dot{M}_v$  is the ventilation airflow rate and  $C_{air}$  is the specific heat capacity of air. A single air exchange between the TZ and the ambient environment is modeled using an alternative circuit through resistance  $R_{vent} = 1/(\dot{m}_v C_{air})$ , which is selected depending on the opening factor of windows and doors. This selection is based on the fact that when some windows and doors are *open*, the energy transferred by airflow is typically much more significant than the energy transferred by conduction through windows that are *closed*.

### Model Equations

Consider the **indicator function** defined on a generic set  $X$ , that indicates membership of an element in a subset  $A$  of  $X$ :

$$\mathbf{1}_A : X \rightarrow \{0, 1\}$$

defined as

$$\mathbf{1}_A(x) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{if } x \notin A \end{cases}$$

#### 4.4. MODEL FORMULATION FOR A THERMAL ZONE

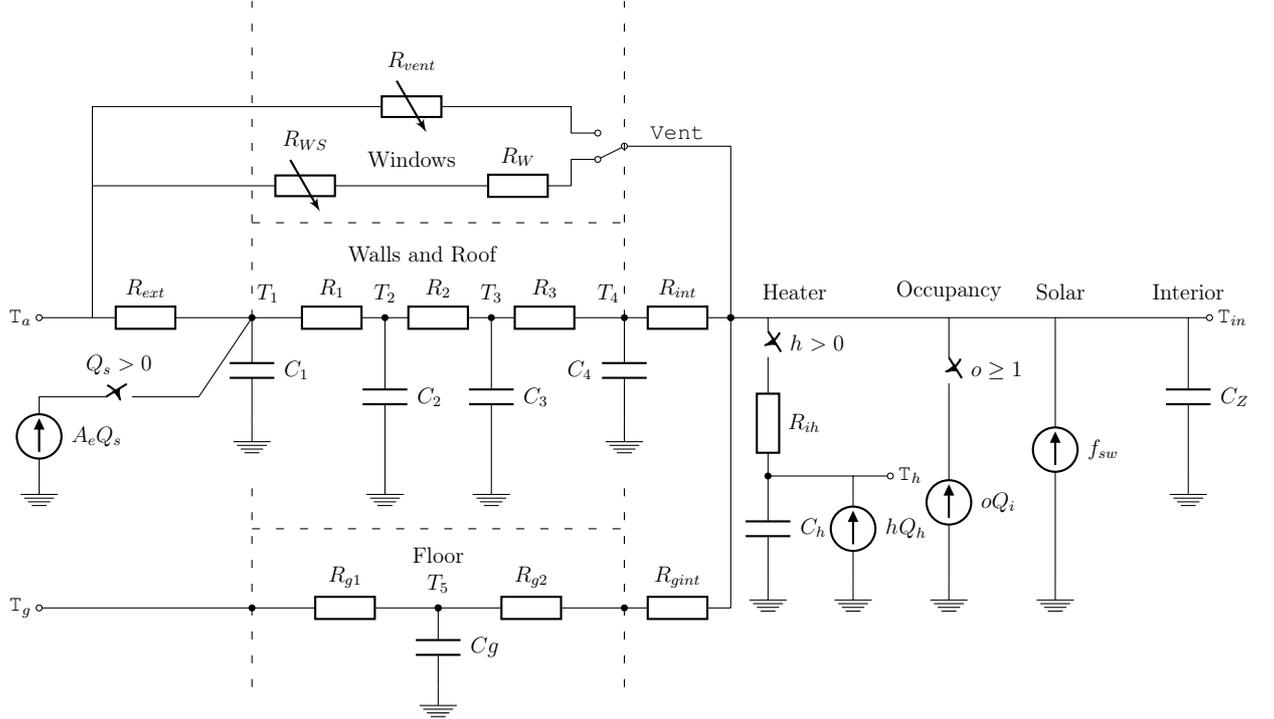


Figure 4.8: Full-scale RC model of a TZ with different heat inputs. Vent is a ventilation activation condition associated with the state of windows, and  $o$  represents the number of occupants.

Using this function to represent each switch in the full-scale RC thermal model, the ODEs that describe how temperatures evolve, given by Kirchhoff's nodal rule, are given by (4.4).

$$\begin{aligned}
 C_1 \dot{T}_1 &= - \left( \frac{1}{R_{ext}} + \frac{1}{R_1} \right) T_1 + \frac{1}{R_1} T_2 + \frac{1}{R_{ext}} T_a + A_e Q_s \mathbf{1}_{\mathbb{R}>0}(Q_s) \\
 C_2 \dot{T}_2 &= \frac{1}{R_1} T_1 - \left( \frac{1}{R_1} + \frac{1}{R_2} \right) T_2 + \frac{1}{R_2} T_3 \\
 C_3 \dot{T}_3 &= \frac{1}{R_2} T_2 - \left( \frac{1}{R_2} + \frac{1}{R_3} \right) T_3 + \frac{1}{R_3} T_4 \\
 C_4 \dot{T}_4 &= \frac{1}{R_3} T_3 - \left( \frac{1}{R_3} + \frac{1}{R_{int}} \right) T_4 + \frac{T_{in}}{R_{int}} \\
 C_g \dot{T}_5 &= - \left( \frac{1}{R_{g1}} + \frac{1}{R_{g2} + R_{gint}} \right) T_5 + \frac{1}{R_{g1}} T_g + \frac{T_{in}}{R_{g2} + R_{gint}} \\
 C_h \dot{T}_h &= - \frac{\mathbf{1}_{\mathbb{R}>0}(h)}{R_{ih}} T_h + \frac{\mathbf{1}_{\mathbb{R}>0}(h)}{R_{ih}} T_{in} + h Q_h \\
 C_Z \dot{T}_{in} &= \frac{T_4}{R_{int}} + \frac{T_5}{R_{g2} + R_{gint}} + \frac{T_a}{R_{RW}} + \frac{\mathbf{1}_{\mathbb{R}>0}(h)}{R_{ih}} T_h - T_{in} \left( \frac{1}{R_{int}} + \frac{1}{R_{g2} + R_{gint}} + \frac{1}{R_{RW}} + \frac{\mathbf{1}_{\mathbb{R}>0}(h)}{R_{ih}} \right) + o Q_i + f_{sw}(Q_{sw})
 \end{aligned} \tag{4.4}$$

where  $RW$  is the **thermal resistance associated with windows and natural ventilation**. This resistance depends on the activation of the natural ventilation airflow circuit which can depend, for example, on the opening factors of an *open* window and door, with  $Vent$  being the set:

$$Vent = \{vState = (wof_1, d_1) \mid wof_1 > 0 \wedge d_1 > 0, wof_1 \in W, d_1 \in D\}$$

or on the ventilation airflow rate with,

$$Vent = \{vState = \dot{m}_v \mid \dot{m}_v > 0, \dot{m}_v \in \dot{M}_v\}.$$

Using the indicator function to model the switch that connects either the resistance associated with natural ventilation circuit, or the circuit containing with windows and shades, we obtain:

$$RW = \mathbf{1}_{Vent}(vState)R_{vent} + \mathbf{1}_{\overline{Vent}}(vState)(R_{WS} + R_w).$$

### State-space Representation

The RC thermodynamic model represented in Figure 4.8, with all inputs active, can be represented by state-space equations [241], given by:

$$\dot{x} = \mathcal{A}x + \mathcal{B}u \quad (4.5)$$

$$y = \mathcal{C}x + \mathcal{D}u \quad (4.6)$$

where

$$x = [T_1, T_2, T_3, T_4, T_5, T_h, T_{in}]^T \in \mathbb{R}^7 \quad (4.7)$$

and  $y \in \mathbb{R}^7$  are the state and model observed outputs. The exogenous inputs to the system are given by:

$$u = [T_a, T_g, o, h, Q_s, f_{sw}]^T \in \mathbb{R}^6 \quad (4.8)$$

and matrices  $\mathcal{A}_{7 \times 7}$  and  $\mathcal{B}_{7 \times 6}$  are model parameters, extracted directly from (4.4), which define how the current state and input affects future states. Matrices  $\mathcal{C}_{7 \times 7}$  and  $\mathcal{D}_{7 \times 6}$  are the output and feed-through matrices.

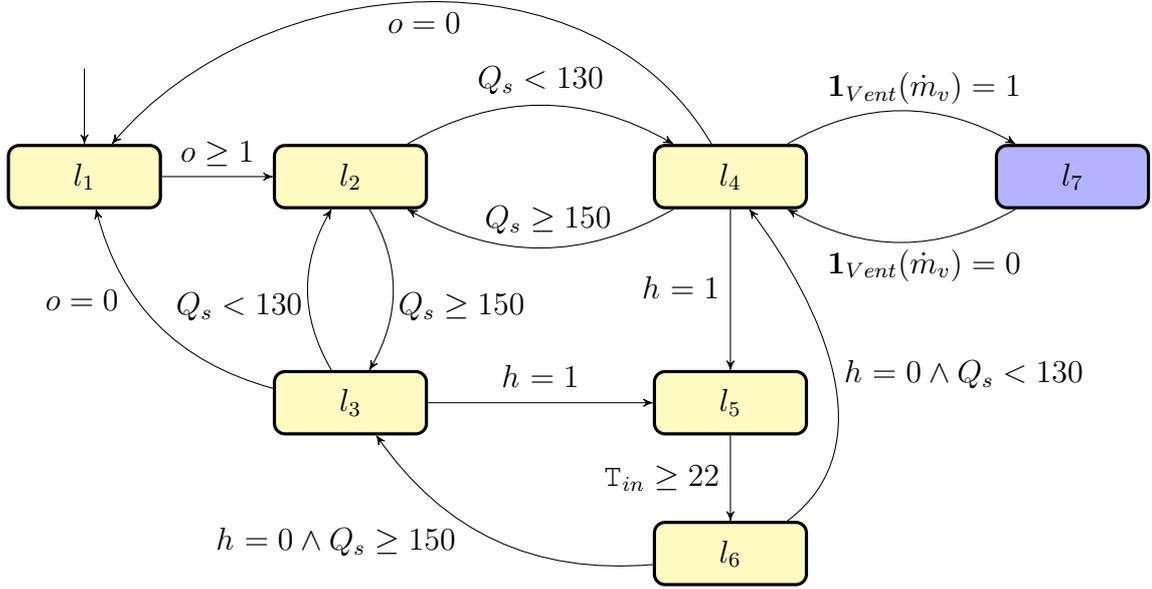


Figure 4.9: Graphical representation of the *context* transitions of the context-based model.

#### 4.4.2 Application Example

Consider the operation of a TZ where occupancy, ambient temperature, solar gains, the opening factor of window shades, airflow rate and heater state are observable input variables. A heater with two heating levels is switched (*on/off*) manually. When the heater is *on*, a temperature controller automatically regulates its operation to guarantee a temperature set point of 22 °C. The heater only operates when the windows are closed, but it can remain *on*, even when tenants go out for lunch. The TZ has a hysteresis controller for the solar shading device. To maximize solar heat gains (and thereby reduce heating loads), while also reducing daylight glare, this controller pulls the shades all the way up ( $ws = 0$ ) when  $Q_s < 130 \text{ W m}^{-2}$  and lowers the shades ( $ws = 1$ ) when  $Q_s \geq 150 \text{ W m}^{-2}$ . When shades, doors and windows are *open*, the TZ is naturally ventilated.

The thermal behavior of the TZ can be described by the context-based model, represented in the directed graph depicted in Figure 4.9. Formally, the automaton is described by the following model:

- $L = \{l_1, l_2, l_3, l_4, l_5, l_6, l_7\}$ . A **collection of contexts** that include the initial *context* ( $l_1$ ), when the TZ is unoccupied ( $o = 0$ ). This *context* changes when

an occupant arrives to  $l_2$  ( $o \geq 1$ ), and then immediately transitions between shades fully closed ( $l_3$ ) and shades fully opened ( $l_4$ ), depending on solar gains ( $Q_s$ ). When the heater is switched *on* ( $h = 1$ ), *context* changes to  $l_5$ , while preheating the TZ, and then to  $l_6$  (where the temperature remains at 22 °C), while the heater remains *on*. TZ *context*  $l_7$  is activated when the TZ is naturally ventilated.

- $X = \mathbb{R}^7$ . The **continuous state**, defined by the temperatures in (4.7).
- $U = U_C \cup U_D$ . The **input variables**, with

$$U_C = \{T_a, T_g, Q_s, Q_{sw}\}, T_a, T_g \in \mathbb{R}, Q_s, Q_{sw} \in \mathbb{R}^+$$

$$U_D = \{o, h, ws, \dot{m}_v\}, o \in \mathbb{N}_0^+, h \in \text{Heat}, ws \in \text{WS}, \dot{m}_v \in \dot{M}_v$$

- $Init = \{l_1\} \times X$ . The initial *context* is marked in the graph of Figure 4.9 by an incoming edge to  $l_1$ , without a source.
- A **vector field** assuming that temperatures follow according to the RC model given by (4.4), depending on input variables of the current *context*. The vector field, represented in the state-space form, is given by:

$$\dot{x} = m(l, x, u) = \mathcal{A}(l)x + \mathcal{B}(l)u$$

with matrices  $\mathcal{A}$  and  $\mathcal{B}$  parameterized by  $h, o, ws$  and  $\text{Vent}$ . The exception is *context*  $l_6$ , where it is assumed that temperatures do not change with  $\dot{x} = 0$ . *Context*  $l_7$  represents a macrostate where parameter  $R_{vent}$  in  $\mathcal{A}(l_7)$  is progressively adjusted according to the natural ventilation airflow  $\dot{m}_v \in \dot{M}_v$ , as shown in Figure 4.10.

- An **output expression** to access all temperatures in all contexts:

$$y = m_Y(x, u) = \mathcal{C}x$$

where  $\mathcal{C}$  is the identity matrix.

- The **domain** of each *context*, listed in Table 4.1.
- A set with all the **edges** represented in Figure 4.9:

$$E = \{(l_1, l_2), (l_2, l_3), (l_2, l_4), \dots, (l_6, l_4), (l_7, l_4)\}$$

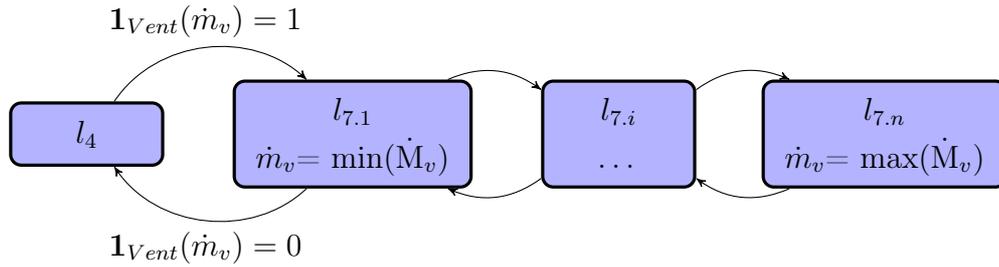


Figure 4.10: The RC model is adjusted according to the natural ventilation airflow rate. *Context*  $l_7$  represents a macrostate where resistance  $R_{vent}$  is progressively adjusted according to  $\dot{m}_v \in \dot{M}_v$ .

$l$	$D(l)$
$l_1$	$X \times \{o = 0, h = 0, \mathbf{1}_{Vent}(\dot{m}_v) = 0\}$
$l_2$	$X \times \{130 \leq Q_s < 150, o \geq 1, h = 0, \mathbf{1}_{Vent}(\dot{m}_v) = 0\}$
$l_3$	$X \times \{Q_s \geq 130, o \geq 1, h = 0, ws = 1, \mathbf{1}_{Vent}(\dot{m}_v) = 0\}$
$l_4$	$X \times \{Q_s < 150, o \geq 1, h = 0, ws = 0, \mathbf{1}_{Vent}(\dot{m}_v) = 0\}$
$l_5$	$\{T_{in} < 22\} \times \{h = 1, \mathbf{1}_{Vent}(\dot{m}_v) = 0\}$
$l_6$	$\{T_{in} = 22\} \times \{h = 1, \mathbf{1}_{Vent}(\dot{m}_v) = 0\}$
$l_7$	$X \times \{h = 0, ws = 0, \mathbf{1}_{Vent}(\dot{m}_v) = 1\}$

Table 4.1: The domain of each *context* of the automaton represented by the graph in Figure 4.9.

- The **guard conditions** associated with each edge  $e \in E$ , listed in Table 4.2.
- A **reset map** given by:  $Rst(e, x, u) = \{x\}$ , assuming that temperatures do not change instantaneously when there is a change in *context*.

Section 5.2 contains the simulation setup of the TZ described in this example, and section 6.2 shows the execution of the corresponding context-based model.

## 4.5 Summary

This chapter shows how the thermodynamic behavior of a single thermal zone can be described by an hybrid automata. The *context* of a thermal zone, as defined in this dissertation, is equated to a discrete configuration associated with and representing a particular thermal behavior of the thermal zone. Based on the hybrid automata

e	$G(e)$
$(l_1, l_2)$	$X \times \{o \geq 1\}$
$(l_2, l_3)$	$X \times \{Q_s \geq 150\}$
$(l_2, l_4)$	$X \times \{Q_s < 130\}$
$(l_3, l_2)$	$X \times \{Q_s \geq 150\}$
$(l_3, l_1)$	$X \times \{o = 0\}$
$(l_3, l_5)$	$X \times \{h = 1\}$
$(l_4, l_2)$	$X \times \{Q_s \geq 150\}$
$(l_4, l_1)$	$X \times \{o = 0\}$
$(l_4, l_5)$	$X \times \{h = 1\}$
$(l_4, l_7)$	$X \times \{\mathbf{1}_{Vent}(\dot{m}_v) = 1\}$
$(l_5, l_6)$	$\{T_{in} \geq 22\} \times U$
$(l_6, l_3)$	$X \times \{h = 0\}$
$(l_6, l_4)$	$X \times \{h = 0\}$
$(l_7, l_4)$	$X \times \{\mathbf{1}_{Vent}(\dot{m}_v) = 0\}$

Table 4.2: Guards associated with each of the edges represented in Figure 4.9.

formalism, the definition of a context-based model is given with some important definitions. The chapter uses a resistor-capacitor model to describe the full-scale thermodynamic model of a thermal zone and uses this model to describe an application example of a context-based model.

# 5

## Simulation Setup

*“You’re never going to get the amount of CO<sub>2</sub> emitted to go down unless you deal with the one magic metric, which is CO<sub>2</sub> per kilowatt-hour.” – Bill Gates (American entrepreneur, philanthropist and programmer).*

This chapter describes the simulation setup for the set of experiments that were devised to validate the examples given in Chapters 3 and 4. The chapter is divided into two main sections, with each describing, in detail, the simulation environment associated with each chapter, respectively. The associated simulation results for each section are presented and discussed separately in Chapter 6.

### 5.1 Reinforcement Learning Simulation Setup

This section presents a set of experiments to test and validate the RL algorithms described in Section 3.2 for the BAS. These RL experiments, developed using Matlab, are discrete time event simulations where the thermodynamics of the TZ and the behavior of a single occupant are both simulated. The duration of each simulation includes the learning phase, which takes several episodes to execute, where each

episode, corresponding to a 24-hour cycle, is divided into the  $N$  discrete time steps:

$$t \in I_{24} = \{t_1, t_2, \dots, t_N\}$$

For each simulation it is assumed that the occupant has a regular working schedule, and becomes uncomfortable if the indoor temperature is not within a predefined comfort range. Since heat is lost through the boundaries of the TZ, additional heat needs to be supplied to guarantee comfort conditions. To control the HVAC system for comfort and energy savings, policies for the BAS are obtained for the *Bang-bang Heater* and *Set Point Heater* problems. The execution and performance of each learned policy is evaluated at the end of each learning period, and results are presented and discussed in Section 6.1.

### 5.1.1 Occupant Schedule Simulation

For both *Bang-bang Heater* and *Set Point Heater* problems the behavior of the occupant was simulated using a finite state machine (FSM) with the following states: *Out*(0), *Working*(1), and *Uncomfortable* (2), as shown in Figure 5.1. State transitions between *Out* and *Working*, and between *Uncomfortable* and *Out*, depend on the **arrive** and **depart** events that occur according to a stochastic schedule that is generated prior to each simulation period. The time instances of arrival and departure from the TZ are generated from *normal distributions*, with mean and variance variables set according to the occupant's usual regular schedule.

### 5.1.2 Modeling Comfort

To model comfort, a fuzzy-set is used to control state transitions between states *Working* and *Uncomfortable*. Uncertainty in occupant's thermal preferences is expressed using a  $\alpha$ -level fuzzy set of temperature values, as described by Dounis and Cariscos [242]. To model comfort, the desired temperature value or set point is chosen as a trapezoidal type-1 fuzzy set. The degree of membership, in this type of set, can take any value in the interval  $[0, 1]$ . The membership function, denoted as  $\mu$ , assigns a degree of membership to each temperature  $T_{in} \in T$ , where  $T$  represents the entire range of environmental temperatures.

The fuzzy set used for modeling comfort is described by the membership function illustrated in Figure 5.2. The set is characterized the by upper and lower bounds

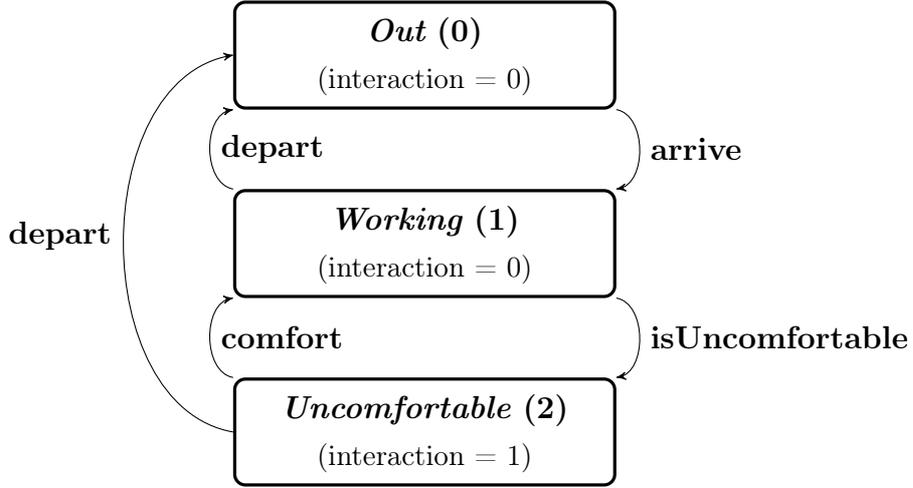


Figure 5.1: Simulation of the occupant behavior. The occupant is simulated using a finite state machine.

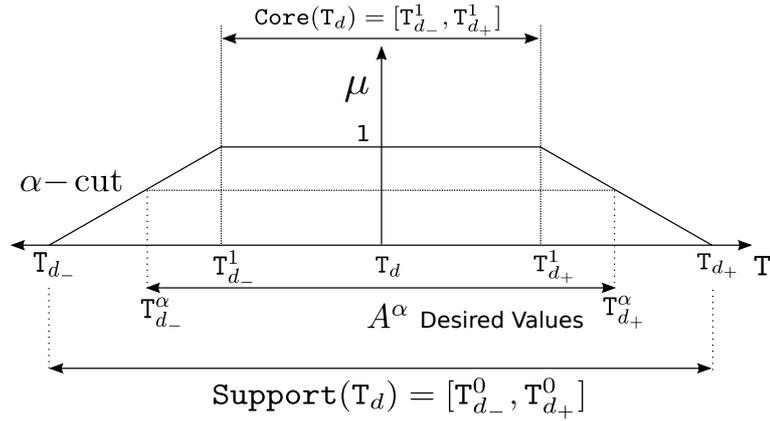


Figure 5.2: Trapezoidal fuzzy set for desired temperature values. The occupant becomes uncomfortable if the indoor temperature  $T_{in} \notin A^\alpha$  (adapted from [242]).

for comfort defined by the  $\text{Support}(T_d) = [T_{d-}, T_{d+}]$  temperature interval, and it is assumed that the occupant is always uncomfortable, if  $T_{in} \notin \text{Support}$ .

The interval associated with the most acceptable temperatures, centered around the most desirable set point  $T_d$ , is given by  $\text{Core}(T_d) = [T_{d-}^1, T_{d+}^1]$ . This interval represents the range of temperatures where the occupant is always comfortable. To evaluate if the occupant is comfortable for temperatures between the limits of  $\text{Core}$  and  $\text{Support}$ , a random evaluation is made to model uncertainty. By using a random variable  $\alpha \in [0, 1]$  taken from a uniform distribution, and the  $\alpha$ -cut set  $A^\alpha$ , which

includes the set of temperatures whose degree of membership in  $A$  is no less than  $\alpha$ :

$$A^\alpha = \{\mathbb{T}_{in} \in \mathbb{T} : \mu(\mathbb{T}_{in}) \geq \alpha\}$$

the occupant's state transitions from *Working* to *Uncomfortable*, if the indoor temperature  $\mathbb{T}_{in} \notin A^\alpha$ .

Figure 5.3 shows the tenant and HVAC state, **with no BAS actuation**, assuming full certainty in the occupant's preferences by setting  $\text{Core}(\mathbb{T}_d) = \text{Support}(\mathbb{T}_d) = [20, 24]^\circ\text{C}$ , with  $\mathbb{T}_d = 22^\circ\text{C}$ . The model assumes that the occupant shows his careless behavior towards energy use, by always leaving the heating *on*, even when the TZ is vacant. The occupant arrives at  $t = 80$  and starts working at instant  $t = 81$ . He becomes *Uncomfortable* at instant  $t = 82$  and acts on the HVAC system by switching it *on* at that instant. The tenant leaves for lunch during the interval  $I = [135, 175]$  and returns home at  $t = 220$ . The temperature graph shows the outdoor temperature  $\mathbb{T}_a = 15^\circ\text{C}$  and two lines representing the comfortable set point temperature ( $\mathbb{T}_d$ ), and the lower temperature limit for comfort ( $\mathbb{T}_{d-}$ ).

### 5.1.3 Performance Evaluation

As the BAS learns how to perform the appropriate actions throughout each learning episode, it should, on average, receive lower penalty values. To track the progress of the learning curve, the **average reward** is calculated for each episode as:

$$\bar{r} = \frac{1}{N} \sum_{t=t_1}^{t_N} r_t \quad (5.1)$$

where  $r_t$  is the reward calculated using (3.5) or (3.8). After learning, performance metrics associated with comfort and heating costs are calculated for each episode to evaluate the execution of the learned policy.

#### Duration of Comfort

Comfort is evaluated by taking into account the amount of time the occupant is comfortable in the *Working* state, measured in discrete time intervals  $\Delta_t$ , as:

$$tComf = \sum_{t=t_1}^{t_N} \text{comfortable}(t) \quad (5.2)$$

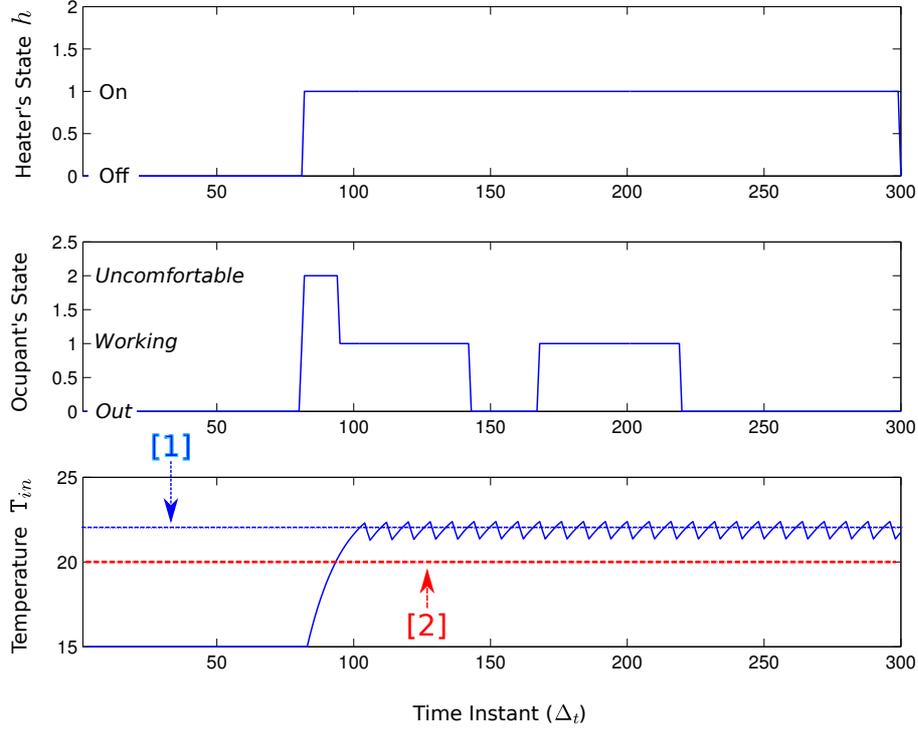


Figure 5.3: The HVAC and occupant’s states over a simulation episode with no BAS. The episode is divided into  $N = 300$  discrete time intervals. The occupant arrives at time instants  $\{80, 175\}$  and leaves at  $\{135, 220\}$ . The occupant is *Uncomfortable* during the interval  $I = [81, 90]$ , and the HVAC is kept *on* even when the tenant leaves the thermal zone. Line [1] represents the desired comfortable set point temperature ( $T_d = 22$  °C), and line [2] represents the lower temperature limit for comfort.

with,

$$\text{comfortable}(t) = \begin{cases} 1 & \text{if } \text{occupantState}(t) = \text{Working} \\ 0 & \text{otherwise} \end{cases}$$

For maximum comfort performance,  $tComf$  should equate to the total amount of time that the occupant is in the TZ. However, considering that there will be several time intervals when the occupant is in the *Uncomfortable* state, additional performance metrics are also calculated to evaluate the **minimum**, **maximum**, and **mean** duration of these intervals denoted, respectively, by  $\Delta_{\min}$ ,  $\Delta_{\max}$  and  $\Delta_{\text{mean}}$ .

## Heating Cost

The evaluation of the **total heating cost** associated with the duration of each episode is calculated differently for each RL problem. For the *Bang-bang Heater* problem, the heating cost is given by the average amount of time the heater is in the *on* state, given by:

$$thCost_1 = \frac{1}{N} \sum_{t=t_1}^{t_N} h(t) \quad (5.3)$$

which is linearly proportional to the amount of energy and monetary costs used for heating. Energy and monetary costs can be estimated for a specific heater, considering the heater’s power consumption and price of energy. This estimation is harder to obtain for the *Set Point Heater* problem since it uses the HVAC network to supply heat. Calculating the efficiency of the network and associated costs is not straightforward. These costs are, however, directly related to the amount of heat supplied to the TZ. Therefore, the total heating cost for the *Set Point Heater* is calculated as:

$$thCost_2 = \frac{1}{N} \sum_{t=t_1}^{t_N} hCost(T_{in}(t)) \quad (5.4)$$

where  $hCost$  is the heating cost given by (3.6).

## Averaged Metrics

To obtain accurate results, each simulation (containing several 24-hour episodes) is executed 30 times independently to obtain the average values of each performance metric. The averaged metric is denoted by appending the notation “*avg\_*” to the associated metric.

### 5.1.4 The *Bang-bang Heater* Problem

To test the *Bang-bang Heater* problem, a single TZ is simulated using the simplified thermal dynamic model and heater described in Section 4.2.1. The simulation program, described by Algorithm 2, executes **250 episodes**, with each episode divided into  $N = 300$  discrete-time intervals. Several simulation parameters are used for controlling the thermodynamic behavior of the TZ, the occupant, and the learning algorithm. At the beginning of the simulation, these parameters are set to the initial values listed in Table 5.1.

---

**Algorithm 2:** Simulation algorithm for the *Bang-bang Heater*.

---

```

Initialize Simulation Parameters;
Execute the Learning Phase;
begin
  for each episode  $\leftarrow 1$  to  $nbrEpisodes$  do
    Initialize Episode Parameters;
    Generate the Occupant's Schedule;
    Update (at a certain learning stage) Parameters  $\varepsilon_{Rate}$  and  $\delta$ ;
    // Start a New 24-hour Cycle
    for each  $t \leftarrow t_1$  to  $t_N$  do
      Simulate the TZ and Heater (Bang-bang Heater);
      Select  $a$ ;
      Execute  $a$ ;
      Simulate the Occupant;
      // Execute the Q-learning update
      Calculate the Reward value;
      Update Q-Values;
    end
    Calculate  $|\nabla Q_{off\_M}|, |\nabla Q_{on\_M}|, |\nabla Q_{off\_T}|, |\nabla Q_{on\_T}|$ 
  end
end

```

---

### Simulation Schedule

The simulated occupant follows the regular schedule shown in Figure 5.3 with a variance of  $3 \Delta_t$  on the arrival and departure time instances. Q-learning parameters  $\delta$  and  $\varepsilon_{Rate}$  change value as the simulation progresses though each episode. After *episode* 50, the BAS follows a greedy policy by setting  $\varepsilon_{Rate} = 0$ . The learning rate that controls convergence to the optimal policy is also decreased, at *episode* 80, to  $lrate = 0.70$ .

TZ	$T_{min}$	21.9 °C
	$T_{max}$	22.1 °C
	$T_a$	15.0 °C
	$R$	$2.30 \times 10^{-3} \text{ K W}^{-1}$
	$C$	$3 \times 10^3 \text{ J K}^{-1}$
	$Q_h$	4 kW
Occupant	$T_d$	22.0 °C
	$\text{Core}(T_d)$	[21.0, 23.0] °C
	$\text{Support}(T_d)$	[20.7, 23.3] °C
Q-Learning	$\delta$	0.90
	$\gamma$	0.80
	$\varepsilon_{Rate}$	0.20
	$w_1$	0.99
	$w_2$	0.01

Table 5.1: Initial simulation parameters used for the *Bang-bang Heater* simulation program.

### Simulation Outputs

The simulation outputs of the *Bang-bang Heater*, for the last executed episode, include: the state of the heater; the state of the occupant; and the TZ temperature, for the entire  $I_{24}$  time interval. To obtain an estimate of the amount of learning days it take for the algorithm to converge, outputs also include the average reinforcement received, given by (5.1).

The magnitude of Q-value updates is expected to be convergent towards zero as the policy converges to the optimal policy. The convergence of Q-values is observed as episodes progress through the simulation. The Q-values for the *Bang-bang Heater*, shown in Table 3.1, are represented as  $N$ -tuples denoted as  $Q_{off\_M}$ ,  $Q_{on\_M}$ ,  $Q_{off\_T}$ ,  $Q_{on\_T}$ , containing the corresponding Q-value for each  $t \in I_{24}$ . The convergence of these values, between learning episodes, is recorded by calculating the magnitude of each  $N$ -tuple update, denoted respectively by  $|\nabla Q_{off\_M}|$ ,  $|\nabla Q_{on\_M}|$ ,  $|\nabla Q_{off\_T}|$  and  $|\nabla Q_{on\_T}|$ . All the output results, including performance metrics, are shown and discussed in Section 6.1.1.

### 5.1.5 Occupant *Uncomfortable* State Redefined

In the previous section, the occupant FSM executed an extreme non-realistic simulation of behavior by persistently interacting with the BAS when the occupant is in the *Uncomfortable* state. In reality, humans are not that pertinacious and, in most cases, occupants expect a certain delay until comfort conditions are re-established. Occupants may even avoid interacting with the HVAC if they are busy, leaving the TZ, or if they are dressed appropriately. Therefore, simulation results using the previous behavior are based on the premises that comfort feedback is always available.

In order to evaluate the effects that more realistic behavior has on performance metrics, the occupant simulator can be modified to incorporate an additional level of random behavior for state *Uncomfortable*, described by the flow chart in Figure 5.4. Variable  $I_L$  represents the time interval that the occupant has been waiting to interact,  $I_L$  represents a *random waiting time* drawn from a standard uniform distribution, and  $act$  is a random variable drawn from a Bernoulli distribution with a given success probability.

New simulations were executed to validate the *Bang-bang Heater* problem with the new behavior. These simulations consider that the occupant, when in state *Uncomfortable*, decides to  $act$  with a probability of 70% and waits, on average, for an interval of  $10 \Delta_t$  before deciding to act again. Performance metrics and execution results are presented and discussed in Section 6.1.2.

### 5.1.6 The *Bang-bang Heater* Problem with Less States

Depending on the thermodynamics of the TZ the effects of actions on the state of the environment present a time delay that must be taken into account in order to increase learning performance. This conclusion, discussed in Section 6.1.2, shows that the RL algorithm for the *Bang-bang Heater*, in its present form, is inefficient in finding the optimal policy from the information obtained with immediate rewards. In the current RL methodology, system states are represented for each  $t \in I_{24}$ , and the BAS executes an action in every time instant. To take into account the heating delay, states and actions can be set further apart in time. In this section, results are evaluated with state variables represented for time instants  $t$  distant  $10 \Delta_t$  apart, according to the following set:

$$s = \{ (t, h) \mid t \in I_{24}, t \equiv 0 \pmod{10}, h \in \{0, 1\} \}$$

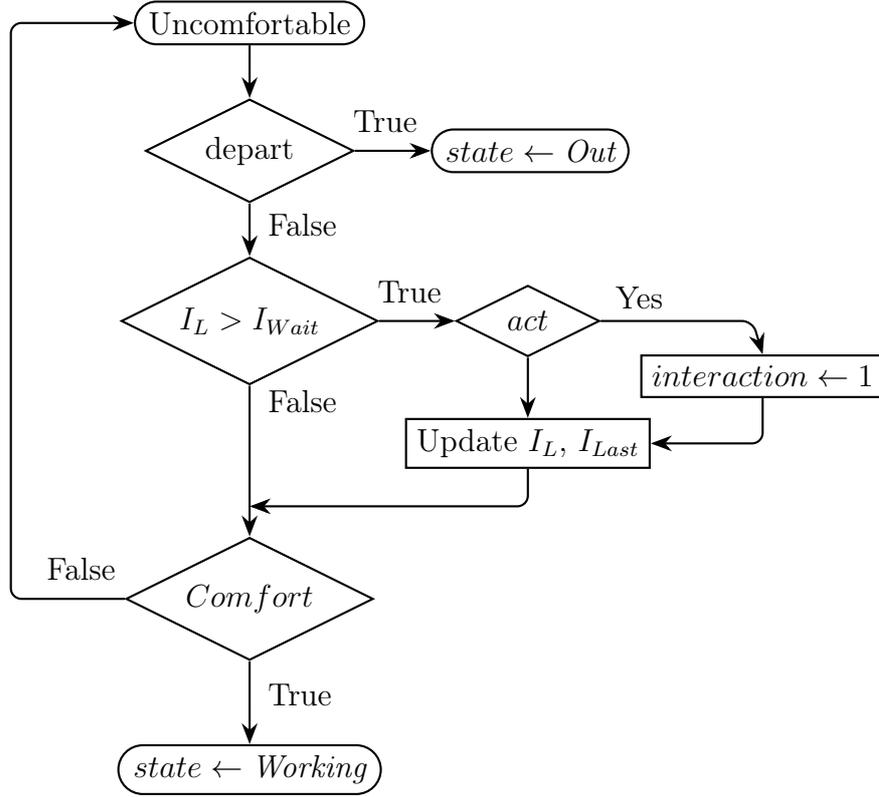


Figure 5.4: New simulation of the occupant behavior for the *Uncomfortable* state. Variable  $I_L$  represents the time interval the occupant has been waiting,  $I_{Wait}$  represents a random waiting time, and  $act$  is a random variable that controls the willingness of the occupant to interact with the HVAC system.

The reward function is modified to take into account the average number of interactions received in the time horizon between the selected action and the following state. The Q-value associated with an action selected at time instant  $t_i$  is updated at time instant  $t_{i+10}$ , considering the following reward:

$$r_t = R(t, h) = -w_1 avgInt - w_2 h$$

where  $avgInt$  is the average number of occupant interactions that occurred during the time interval between  $t_{i+10}$  and  $t_i$ .

The *Bang-bang Heater* problem with less states is simulated using the simulation parameters listed in Table 5.1, and the simulation results for a duration of 90 episodes are shown in Section 6.1.3.

### 5.1.7 The *Set Point Heater* Problem

This section describes the simulation program for the *Set Point Heater*. The structure of the program is described by Algorithm 3.

---

**Algorithm 3:** Matlab simulation program for the *Set Point Heater*.

---

```
Initialize the Wire Set;
Create the Neural Network;
Initialize Simulation Parameters;
Execute the Learning Phase;
begin
  for each episode  $\leftarrow$  1 to nbrEpisodes do
    Initialize Episode Parameters;
    Generate the Occupant's Schedule;
    Train the Neural Network for the Set of Wires;
    Update (at a certain learning stage) Parameters  $\varepsilon_{Rate}$  and  $\delta$ ;
    // Start a New 24-hour Cycle
    for each  $t \leftarrow t_1$  to  $t_N$  do
      Calculate the Output (Wires) of the Neural Network;
      Select the Wire/Action with the Highest Q-value;
      Add exploration noise to the selected action;
      Calculate and Set the HVAC Temperature ( $T_s$ );
      Simulate the TZ and the HVAC PID Controller;
      Simulate the Occupant;
      // Execute the Q-learning One-step Update
      Calculate the Average Number of Interactions (in the last  $4\Delta_t$  int.);
      Calculate Heating Cost and Reward;
      Calculate the New Set of Wires;
    end
    Save the Average Reward;
  end
end
```

---

### The Neural Network and Output Wires

The definition of state for the *Set Point Heater* is a continuous time variable and the algorithm uses a NN to map each state to its respective set of wires. However, due to the discrete nature of the simulation with Matlab, the program iterates only through the discrete time instants included in  $I_{24}$ , and all the wire sets are stored in an N-dimensional vector containing  $N = 300$  sets. Without a loss of generality, there is a set with 6 wires for each  $t \in I_{24}$  and these sets are all initialized to the following values:

$$\begin{aligned}\text{Wire 1} &= (a_1, q_1) = (0.0, 0.1) \\ \text{Wire 2} &= (a_1, q_1) = (0.2, 0.1) \\ \text{Wire 3} &= (a_1, q_1) = (0.4, 0.1) \\ \text{Wire 4} &= (a_1, q_1) = (0.6, 0.1) \\ \text{Wire 5} &= (a_1, q_1) = (0.8, 0.1) \\ \text{Wire 6} &= (a_1, q_1) = (1.0, 0.1)\end{aligned}$$

Since all the wires are stored in the memory vector, Algorithm 3 can be implemented without the NN. Wires can be immediately accessed and updated without having to train the NN. This gives the advantage of avoiding the NN training period and fitting errors. However, some applications may need to control temperature with higher time resolution. For those applications, the NN option is available to provide access to the set of wires for any continuous value of  $t$ .

To evaluate the effects of fitting errors on the learned policy, a feed-forward back-propagation NN was included in the Matlab simulation using the NN Toolbox [243]. **This NN was created empirically** with two hidden layers and one output layer. The first hidden layer was composed of 10 neurons with a Tan-Sigmoid (*tansig*) transfer function, the second hidden layer had 10 neurons with a Log-Sigmoid (*logsig*) transfer function, and the output layer included 12 neurons with a linear (*purelin*) transfer function. To train the NN the Levenberg-Marquardt (*trainlm*) algorithm was used with the *epochs* training parameter set to 10.

### HVAC and Thermal Zone Simulation

To simulate the HVAC for the *Set Point Heater*, a PID controller was implemented to regulate the indoor temperature of the TZ [244]. Both the TZ and controller are simulated. The TZ is described by the thermodynamic model given by (4.3), using the same RC thermodynamic parameters as the *Bang-bang Heater*. Figure 5.5 shows the block diagram of the feedback loop with the model of the TZ and the controller,

where  $K_p$ ,  $K_i$ , and  $K_d$  are non-negative coefficients for the proportional, integral and derivative terms, respectively.

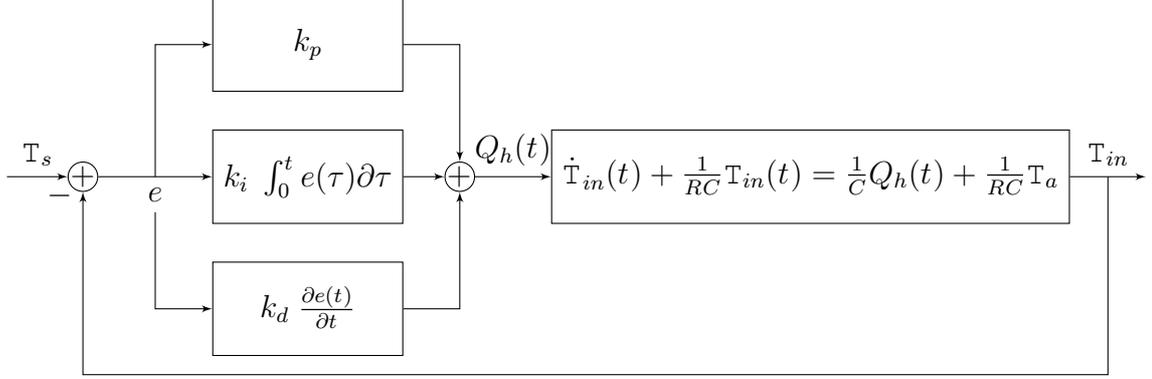


Figure 5.5: Simulating the thermal zone with a proportional–integral–derivative controller.

Using the gains  $(k_p, k_i, k_d) = (1500, 200, 10)$  for the PID controller, Figure 5.6 shows the step response of the system to the following temperature set points, set at different time instants:

$$\{T_s(50) = 22, T_s(100) = 16, T_s(200) = 20, T_s(250) = 22\}$$

The resulting controller regulates the temperature with no overshoot, and no offset error in the final steady-state condition. It is suitable to simulate how the HVAC will regulate the TZ to a given temperature set point.

### Simulation Outputs

The *Set Point Heater* simulation simulates the occupant with the *Uncomfortable* state redefined, and the BAS learns with less states, following the discussions given in Sections 5.1.5 and 5.1.6. Therefore, the program described by Algorithm 3 updates each set of wires with a delay of  $4 \Delta_t$  and the reward function, given by (3.8), now becomes:

$$r_t = -w_1 avgInt - (w_2 - w_1 avgInt) hCost(T_s)$$

where  $avgInt$  represents the average number of interactions that occurred during the time interval  $[t, t + 4 \Delta_t]$ .

The simulation program sets parameters  $\gamma = 0.4$  and  $(w_1, w_2) = (0.90, 0.10)$  at the beginning of the simulation. All the other simulation parameters used for the

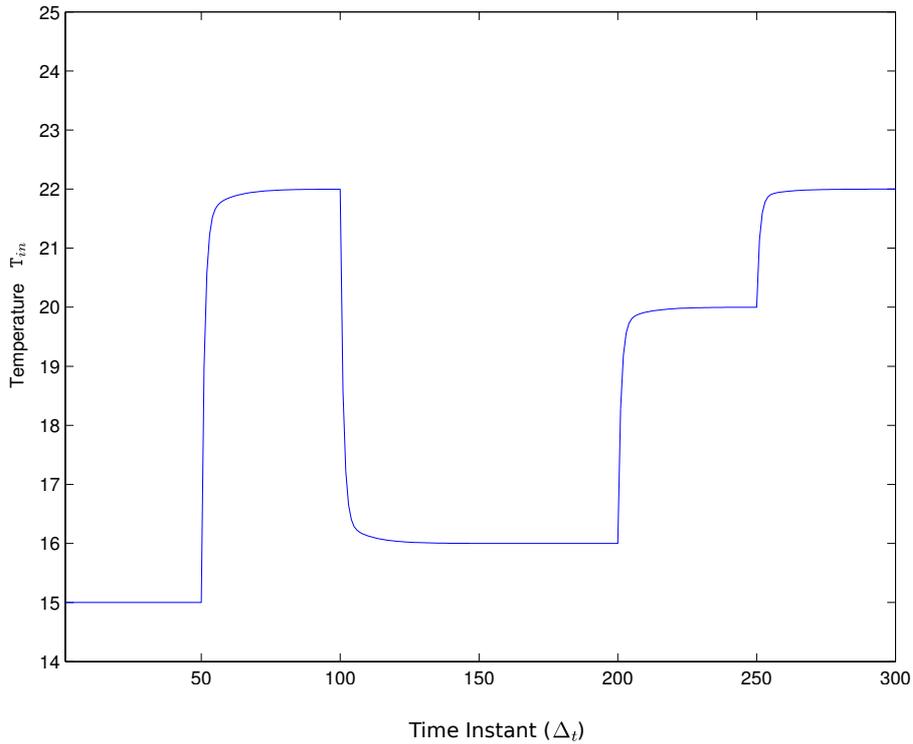


Figure 5.6: Temperature response of the controller and the thermal zone, with different temperature set points.

*Bang-bang Heater* remain constant. With a duration of 60 episodes, the BAS follows a greedy policy after *episode* 10, and the learning rate is set to  $lrate = 0.60$  at *episode* 25. The optimal policy, obtained with the last episode, is executed and the results are presented and discussed in Section 6.1.4.

## 5.2 Context-based Modeling Simulation Setup

This section contains a description of the experimental setup used to execute and evaluate the context-based model described for the example given in Section 4.4.2. The operational semantics of this model were implemented using the MATLAB simulation environment, considering the hypothetical TZ enclosed by the box-shaped building illustrated in Figure 5.7.

The execution of the context-based model was simulated by solving the ODEs given by (4.4), to obtain the model's continuous-state signals, while simulating the discrete *context* transitions according to the model's *guards* and *edges*. The execution

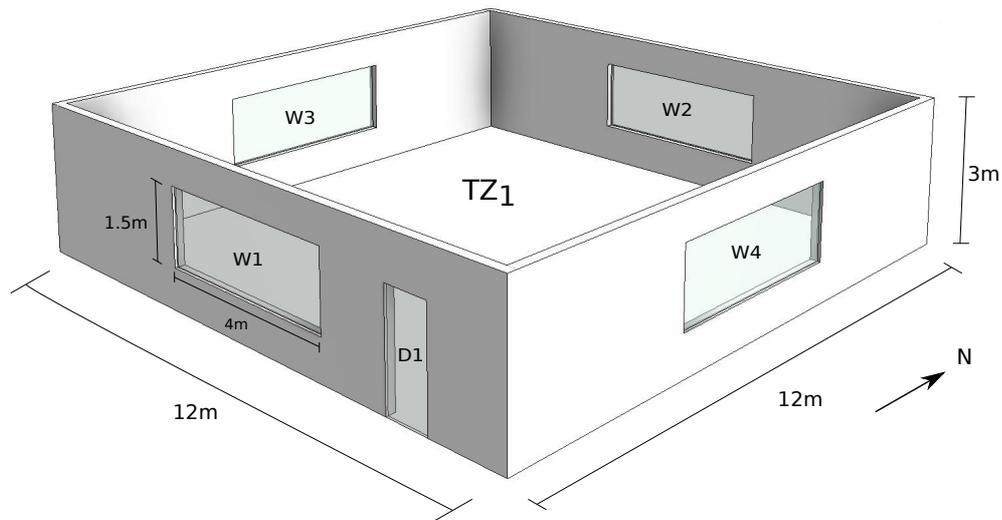


Figure 5.7: Detail of the box-shaped building used for the simulation, with multi-layer walls, windows, and a single thermal zone.

is then compared with the simulation outputs of EnergyPlus [245,246], applied to the same building. Conclusions are drawn from this comparison about the effectiveness of the context-based model in obtaining the same results as EnergyPlus.

### 5.2.1 EnergyPlus

EnergyPlus is a popular simulation tool developed for building performance and energy analysis that integrates, as shown in Figure 5.8, a set of information about the simulated building. This information includes the following descriptions:

- **Architecture and construction materials.** The description of the building's architecture and construction details, including all the interior and exterior walls, floors, roofs, doors, windows, blinds, and construction materials.
- **Information about weather and climate.** Collected from statistically assembled weather data, to be used as the reference for the typical weather at the building's location. This simulation input strongly influences the external loads from outdoor temperature, humidity, wind and insulation.
- **HVAC systems and components.** A description that can be detailed enough to include air and water loops with splitters and mixers, zones, heating and cooling coils, pumps, set point managers, ventilators, furnaces and chillers. Control

strategies can also be included to describe the behavior of HVAC systems during the simulation time.

- **Internal loads.** Associated with occupancy, activity levels, equipment and lighting.
- **Operating strategies.** Description of how blinds, temperature set points and other parameters are set during the simulation time. This item also includes various schedules that define, during this period, the value of many simulation parameters, such as occupancy, opening factors of doors, windows and blinds.
- **Simulation parameters.** Describing, among other simulation parameters, the simulation time interval, the duration of the simulation time step, numeric convergence tolerances, and simulation outputs (files and variables).

The input descriptions for EnergyPlus are composed of several input objects included in a simulation *input definition file*<sup>1</sup> [247]. EnergyPlus loads this file prior to simulation and, depending on these objects, executes several algorithms to calculate, among other things, the heating and cooling loads in different zones, temperatures, air flow through doors and windows, energy requirements, lighting conditions, and financial costs. Most of the simulation algorithms are complex and predict these variables with acceptable accuracy [248]. In order to calculate temperatures through walls, for instance, the simulation engine divides each material layer within a construction into a number of nodes between 6 and 18 [248]. With this level of detail, we can assert that the simulation of the associated heat-balance equations, including the simulation of the energy transmitted through windows and the airflow network, uses models with much higher complexity than the context-based model used in the application example. Therefore, **the execution of an EnergyPlus simulation is a suitable validating reference for model performance evaluation.**

### 5.2.2 Simulation Diagram

Modeling, simulation and analysis of buildings and building control systems is becoming increasingly complex. Many systems involve multiple domains, such as thermodynamics, fluid dynamics, heat and mass transfer, electrical systems, control systems and communication systems [249]. Modeling and simulating such systems not only

---

<sup>1</sup>File with extension .idf

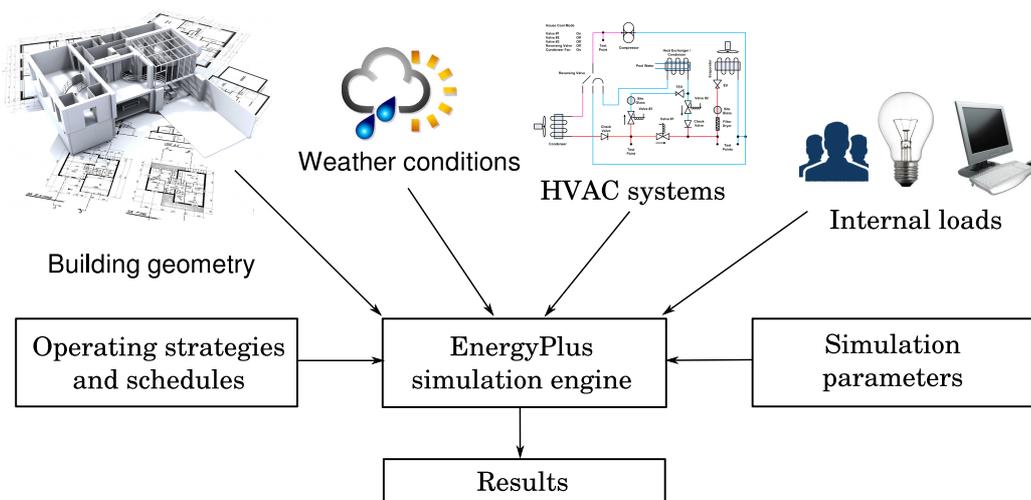


Figure 5.8: The EnergyPlus simulation engine integrates all the detailed information about the building structure and geometry, such as weather conditions, HVAC systems, internal loads, schedules, etc. Simulation results are used as validating references for performance evaluation.

requires a higher level of abstraction and modularisation, but also a proper platform to integrate each simulated system. For this particular experiment, in order to simulate occupancy, the operation of windows, shades, and heating equipment, some variables needed to be continuously set during the simulation time. Therefore, all the inputs for the building simulation, with EnergyPlus, were implemented using the Buildings Controls Virtual Test Bed (BCVTB) [250–253], which is an open-source modeling and distributed co-simulation environment, based on UC Berkeley’s Ptolemy II program [233, 254]. Ptolemy allows different simulation programs and tools<sup>1</sup> to be coupled, in the same coherent simulation work-flow, by synchronizing and exchanging data between them as the simulation time progresses. Programs and tools are abstracted as actors and connected (using TCP/IP sockets) in a diagram using the Ptolemy II graphical environment [233, 254, 255]. This diagram has data-flow semantics: actors are invoked/triggered when data is available at their input.

In the experimental setup, the BCVTB environment is used to control an actor that wraps around EnergyPlus for a building simulation, as illustrated in the example of Figure 5.9. The EnergyPlus actor is parameterized with the *input definition file* describing the simulated building and receives, during the simulation time,

<sup>1</sup>The BCVTB includes a library for simulating FSMs, signal sources, and sinks, among other things.

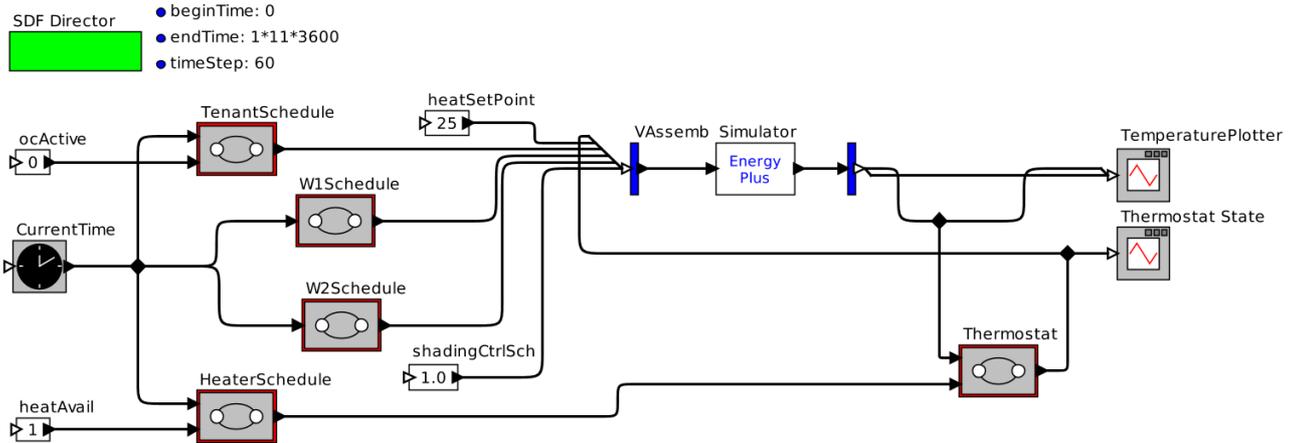


Figure 5.9: Example of a Buildings Controls Virtual Test Bed (BCVTB) simulation model using the Ptolemy II graphical environment. Actors simulate occupancy, the state of windows ( $wof_1, wof_2$ ) and heater. The EnergyPlus simulation actor (at the center) outputs a vector of temperatures at each step of the simulation.

a vector of inputs  $\{o, ws, h, wof, \dots\}$  which condition the building’s energy balance equations. The actor then makes available, also during the simulation time, the simulation outputs  $\{T_a, T_{in}, Q_s, \dot{m}_v, \dots\}$  that are used by other actors. Behaviors such as the feedback control loop associated with the operation of the thermostat, occupancy, and the opening and closing of windows are implemented by different actors, such as as FSMs, which simulate schedules and other conditioned actions, such as changing the state of the heater, depending on  $T_{in}$ .

By changing the conditions that define *context* during the simulation time with EnergyPlus, we can observe the effects of these changes on  $T_{in}$  in Figure 5.10. The output for a 24-hour simulation example starting on the 1st day of January shows how temperature  $T_{in}$  depends on the evolution of temperature  $T_a$  throughout the day, and how it changes when the heater is switched *on*, the windows are *open*, and the space is *unoccupied*. As intuitively expected, *context* has a significant impact on the evolution of  $T_{in}$ . We can observe, for example, that when windows are *open*,  $T_{in}$  tends to approximate  $T_a$  as heat is lost through natural ventilation (which depends on wind behavior). This interrupts the previous behavior where  $T_{in}$  increases at a certain rate, with the heater *on*. Using a context-based model, each of these behaviors is modeled separately.

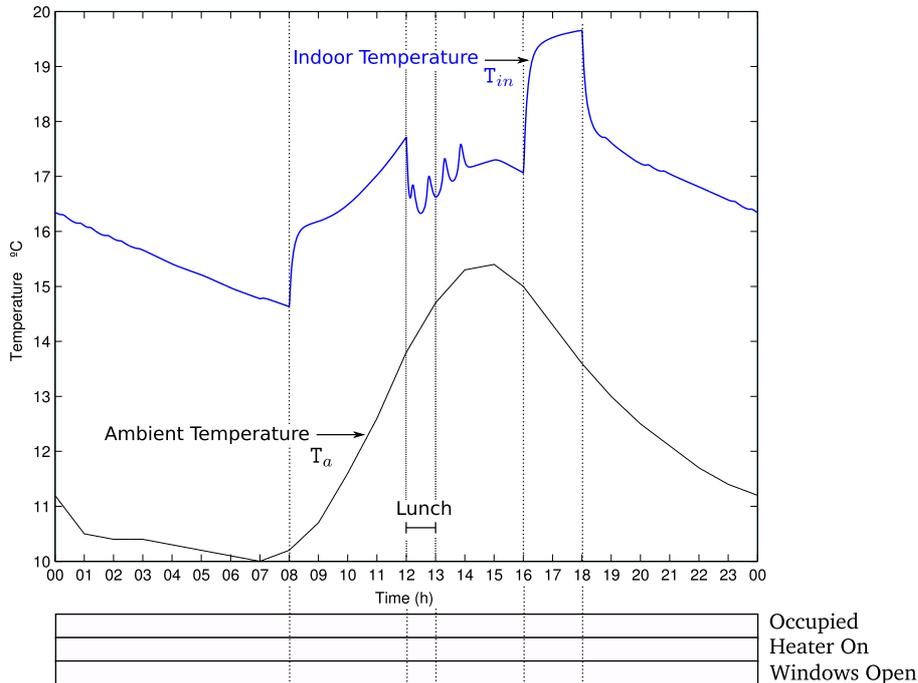


Figure 5.10: EnergyPlus 24-hour simulation output for the thermal zone showing how temperatures  $T_a$  and  $T_{in}$  evolve in different contexts depending on the heater being *on*, the windows *open*, and occupancy (simulation for January 1st, in Lisbon).

### 5.2.3 Building Site and Construction

The meteorological data of the typical meteorological year of the location<sup>1</sup> in Lisbon, Portugal, was used as the input for the building zone simulation. The theoretical building is exposed to solar radiation and wind and encloses a single TZ. It has four external walls, a floor area  $A_{floor} = 144 \text{ m}^2$  (12 m west and south walls), and a volume of  $432 \text{ m}^3$  (zone height of 3 m). The building has a window ( $4 \times 1.5 \text{ m}$  each) on each wall and no internal partitions. The envelope, made by the surrounding walls and roof, is composed of stucco over common brick and gyp-board, and the floor is composed of a single-layer of concrete. The properties of the various construction elements, used in EnergyPlus and in our RC model, are shown in Table 5.2. **Each layer of a material is defined by its thickness  $z$ , thermal conductivity  $\lambda$ , specific heat capacity  $C_p$ , and density  $\rho$ .**

<sup>1</sup>INETI, GPS: 38.73,-9.15

Element	Layering	Thickness (m)	Thermal conductivity (W m <sup>-1</sup> K <sup>-1</sup> )	Specific heat capacity (J kg <sup>-1</sup> K <sup>-1</sup> )	Density (kg m <sup>-3</sup> )
External	Stucco	0.0254	0.690	837	1858.0
Walls and Roof	Brick	0.1000	0.730	837	1922.0
	Plaster	0.0190	0.726	837	1602.0
Floor (144 m <sup>2</sup> )	Concrete	0.2000	1.730	837	2242.6
Windows (24 m <sup>2</sup> )	Glass	0.0025	0.700		
Shades	Insulation	0.0200	0.050		

Table 5.2: Properties of the constructive elements for the box-shaped building presented in Figure 5.7.

#### 5.2.4 Parameters of the RC Model

Because model RC parameters have intuitive physical meaning, all values can be determined from the material properties and geometry of the surfaces surrounding the simulated building. For a surface area  $A$ , the **conductive thermal resistance** and the **thermal capacitance** that characterizes a specific layer  $i$  in a steady-state regime are given, respectively, by:

$$R_i = \frac{\text{R-Value}}{A} \quad (\text{K W}^{-1}) \quad (5.5)$$

$$C_i = \rho C_p z A, \quad (\text{J K}^{-1}) \quad (5.6)$$

where R-Value =  $z/\lambda$  is the unit thermal resistance of the material in that layer.

Walls exchange thermal energy with the surrounding air in each TZ. Convection is the heat transfer between these two different mediums. The heat flow between a wall surface at temperature  $T_{wall}$ , and the surrounding air at temperature  $T_a$ , is given by Fourier's Law [256]:

$$Q_{conv} = h_c A (T_{wall} - T_a),$$

where  $h_c$  is the convective heat transfer coefficient. The convective resistance is given by:

$$R_c = \frac{1}{h_c A}. \quad (5.7)$$

The following sections describe how the parameters of the full-scale lumped RC thermal model were calculated using these expressions.

### Walls and Roof

The total resistance and capacity of a multi-layer construction are calculated, taking into account the resistance and capacitance of each of the  $n$  layers, as follows:

$$R = \sum_{i=1}^n R_i + R_{ext} + R_{int} \quad (5.8)$$

$$C = \sum_{i=1}^n C_i \quad (5.9)$$

where the walls and roof are assumed to be composed of the same materials. Therefore, a single 3R4C model represents the combination of these constructions for their total combined surface area, subtracting the area of the windows:

$$A_{wr} = A_{walls} + A_{roof} - A_{wind} = 264 \text{ m}^2$$

where  $A_{walls}$ ,  $A_{roof}$  and  $A_{wind}$  represent the total area of the building envelope, roof and windows, respectively.

Model parameters  $R_1, R_2, R_3, C_2, C_3$  are obtained from the 3R2C model for the multi-layer construction, as described by Fraisse et al. [240] (see appendix A). Therein, the authors calculate the truncated second-order transmission matrix of a wall composed by a single-layer material with an equivalent resistance and capacity given by  $R$  and  $C$ , and equate this matrix with the transmission matrix of the 3R2C wall model to extract its parameters. They then propose the 3R4C model to avoid having temperature pairs  $(T_a, T_1)$  and  $(T_4, T_{in})$  instantly coupled, by transferring 5% of each of the two internal capacities  $C_2$  and  $C_3$ , to  $C_1$  and  $C_4$ , respectively. However, to account for the effects of solar gains, simulation results showed that this value can go up to 70%.

### Floor

The entire floor area of the building is composed of a single layer of concrete with resistance  $R_g$  and capacitance  $C_g$ , modeled using a 2R1C network with the following resistances:

$$R_{g1} = R_{g2} = \frac{R_g}{2}.$$

Surface	Heat transfer coefficient ( $\text{W m}^{-2} \text{K}^{-1}$ )
<b>Walls and roof</b>	
Exterior ( $h_{ext}$ )	10.22
Interior ( $h_{int}$ )	3.076
<b>Floor</b>	
Interior ( $h_g$ )	1.000

Table 5.3: Convective heat transfer coefficients.

### Convective Resistances

Exterior and interior convective resistances  $R_{ext}$  and  $R_{int}, R_{gint}$  are calculated using the convective heat transfer coefficients for exterior and interior surfaces of the envelope, and the interior surface of the floor, denoted respectively by  $h_{ext}$ ,  $h_{int}$ , and  $h_g$ , listed in Table 5.3.

### Zone Air

The thermal capacitance  $C_Z$  of the total volume of air inside the TZ is calculated with (5.6), considering the following specific heat and density:

$$\begin{aligned} C_{air} &= 1.005 \text{ kJ kg}^{-1} \text{K}^{-1} \\ \rho_{air} &= 1.205 \text{ kg m}^{-3} \quad (\text{at } \approx 20 \text{ }^\circ\text{C}). \end{aligned}$$

### Windows

Windows are composed of a single layer of glass with the following thermal transmittance and Solar Heat Gain Coefficient:

$$\begin{aligned} \text{U-value} &= 3.071 \text{ (W m}^{-2} \text{K}^{-1}) \\ \text{SHGC} &= 0.258 \end{aligned}$$

The thermal resistance per unit area (R-Value =  $1/\text{U-Value}$ ) of the simplified window combines the interior and exterior surface heat transfer coefficients and the glass-to-glass resistance. The total thermal resistance of the effective window area is given by:

$$R_W = \frac{\text{R-Value}}{A_{wind}}$$

### Shades

Exterior shades with insulation are used on the outside of each window. Shades are operated with a window cover that ranges from fully *open* to fully *closed*. Figure 5.11 shows the evolution of outdoor and indoor temperatures obtained from the EnergyPlus simulation. Figure 5.11a shows the effects that  $ws$  has on the indoor temperature, considering that the building is not exposed to sun radiation. The indoor temperature is affected by the opening factor  $ws$  because shades contribute to a small increase in the thermal resistance of all window areas. Therefore,  $T_{in}$  is slightly higher when shades are *closed*. The total thermal resistance of the shades  $R_{WS}$  is calculated using  $A_{wind}$  and the physical properties given in Table 5.2, multiplied by  $ws$ . Shades also affect the amount of solar radiation entering through the windows. Thus, Figure 5.11b shows the effects of shades when the building is exposed to sun radiation. Indoor temperature  $T_{in}$  is higher during daylight hours, when shades are *open*.

### Heater

The heater, represented in Figure 5.12, consists of a 75 W *constant volume fan* with an air volume  $\dot{m}_{vh} = 1.194 \text{ kg s}^{-1}$  and an 8 kW *electric heating coil*. Resistance  $R_{ih}$  is given by:  $R_{ih} = 1/(\dot{m}_{vh}C_{air})$  and  $C_h$  is calculated from samples of  $T_h$  and  $T_{in}$ , taken from the initial heating time interval, and is given by:

$$C_h = \frac{Q_{C_h} I_s}{T_h}$$

$$Q_{C_h} = Q_h - \frac{T_h - T_{in}}{R_{ih}}$$

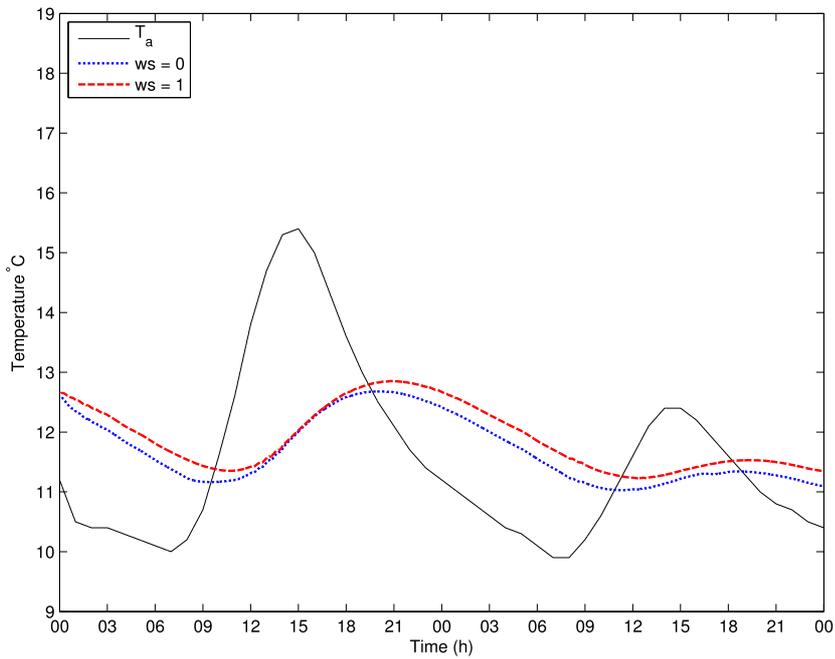
where  $I_s$  is the sampling period of the simulations, set to 60 s.

### Occupancy

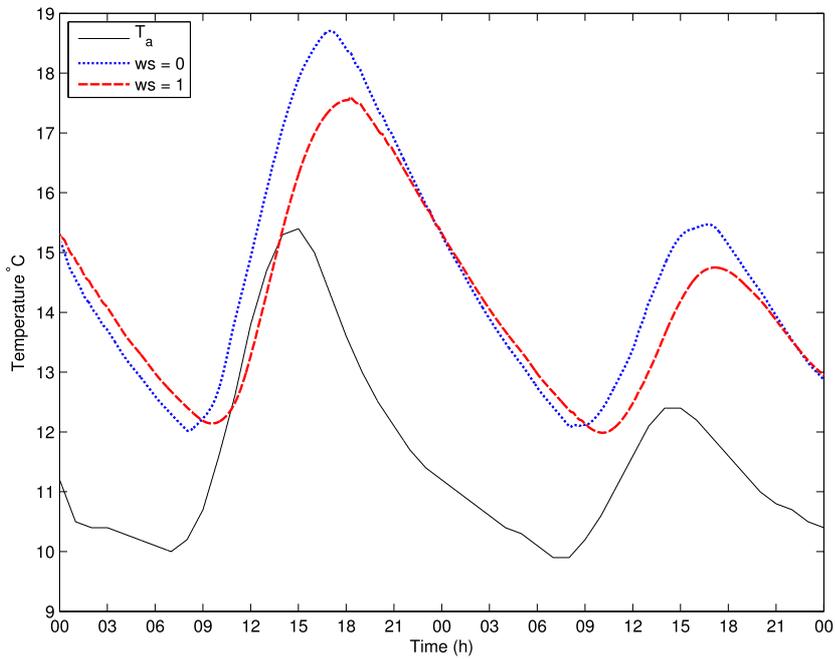
Occupancy is simulated by considering that 5 people will occupy the TZ between 8h00 and 12h00, and between 13h00 and 18h00. EnergyPlus assumes the metabolic rate to a value of  $Q_i = 132 \text{ W/person}$ .

### Solar Gains

Due to many uncertain factors that influence solar gains in real buildings (solar angles, surrounding buildings, reflections, etc.), a detailed calculation of  $Q_s$  is rather elaborate. For this experimental setup,  $Q_s$  is obtained from EnergyPlus by averaging



(a) No sun and no wind exposure.



(b) Exposed to sun and wind.

Figure 5.11: Evolution of outdoor temperature ( $T_a$ ), and indoor temperature ( $T_{in}$ ), when shades are *open* ( $ws = 0$ ) and *closed* ( $ws = 1$ ), contrasting the scenario without sun and without wind (a) to the scenario with sun and wind exposure. (b).

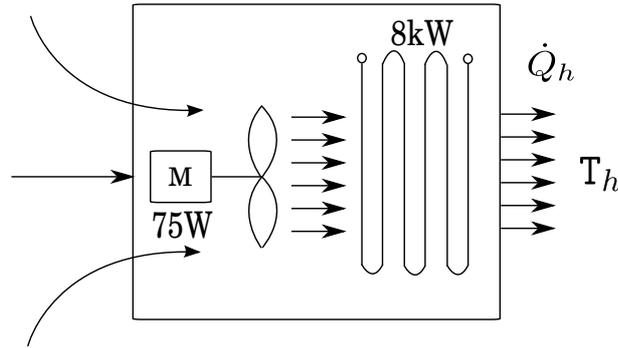


Figure 5.12: Detail of the unit heater with a constant volume fan and an electric heating coil of 8 kW, used for simulation.

the solar heat gain per area, incident on each of the four external walls and the roof surface. The total solar radiation rate transmitted through the windows,  $Q_{sw}$ , is also obtained from EnergyPlus. This radiation acts directly on the internal building mass and heat is then transmitted to the zone air. To simulate this transmission function,  $f_{sw}(Q_{sw})$  implements as a 6th order low-pass Butterworth filter to smooth  $Q_{sw}$  and simulate the mass-to-air thermal transmission delay. This filter is described by coefficient vectors  $C_A$  and  $C_B$  using the standard difference equation:

$$\begin{aligned} C_A(1)f_{sw}(n) = & C_B(1)Q_{sw}(n) + C_B(2)Q_{sw}(n-1) \dots + C_B(7)Q_{sw}(n-6) \dots \\ & -C_A(2)f_{sw}(n-1) \dots - C_A(7)f_{sw}(n-6) \end{aligned} \quad (5.10)$$

with  $C_A$  and  $C_B$  obtained using the Matlab filter design tool:

$$[C\_B, C\_A] = \text{butter}(N, W_n)$$

where  $N$  represents the order of the filter and  $W_n$  is the normalized cutoff frequency. Equation (5.10) was executed using the Matlab function:

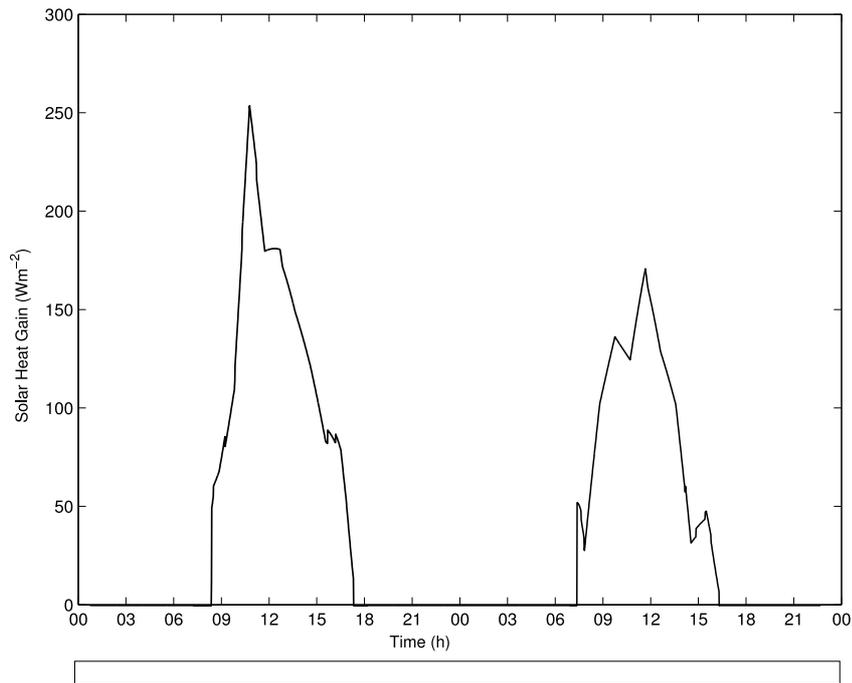
$$\text{fsw} = G * \text{filter}(C\_B, C\_A, Q_{sw})$$

with  $G$  representing the filter attenuation. Filter parameters  $N = 6$ ,  $G = 0.18$ , and  $W_n = 0.0099$ , were empirically adjusted to minimize model errors.

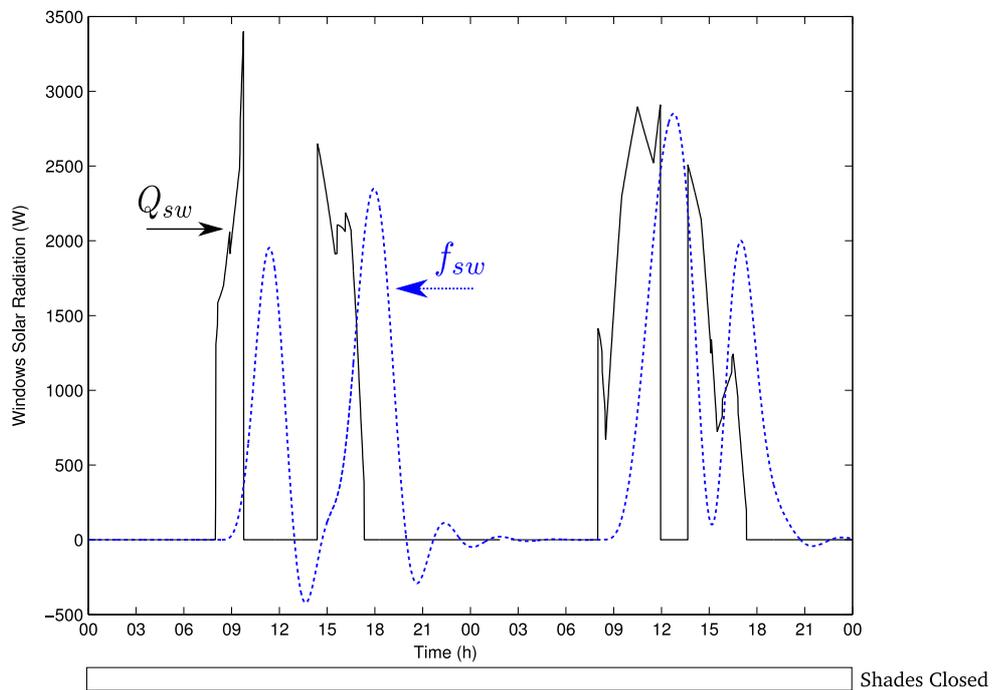
Figures 5.13a and 5.13b show  $Q_s$  and  $Q_{sw}, f_{sw}(Q_{sw})$ , during a 48-h simulation interval. Shades are *closed* when  $Q_s \geq 150$  and *open* when  $Q_s < 130$ .

### Natural Ventilation

Although the discrete sets  $W$  and  $D$  include several values for opening factors for windows and doors, in this simulation we opted to model the extreme case of on/off



(a) Averaged solar gains  $Q_s$ .



(b)  $Q_{sw}$ , and heat transmission output to zone air ( $f_{sw}$ ).

Figure 5.13: 48-h simulation of solar gains incident on the building envelope (a), and through building windows with heat transmission to zone air (b).

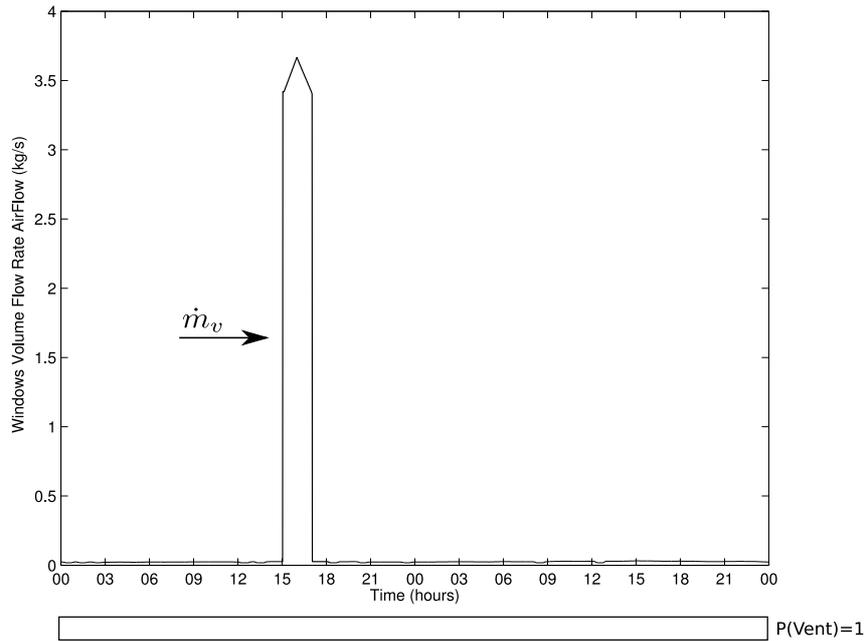


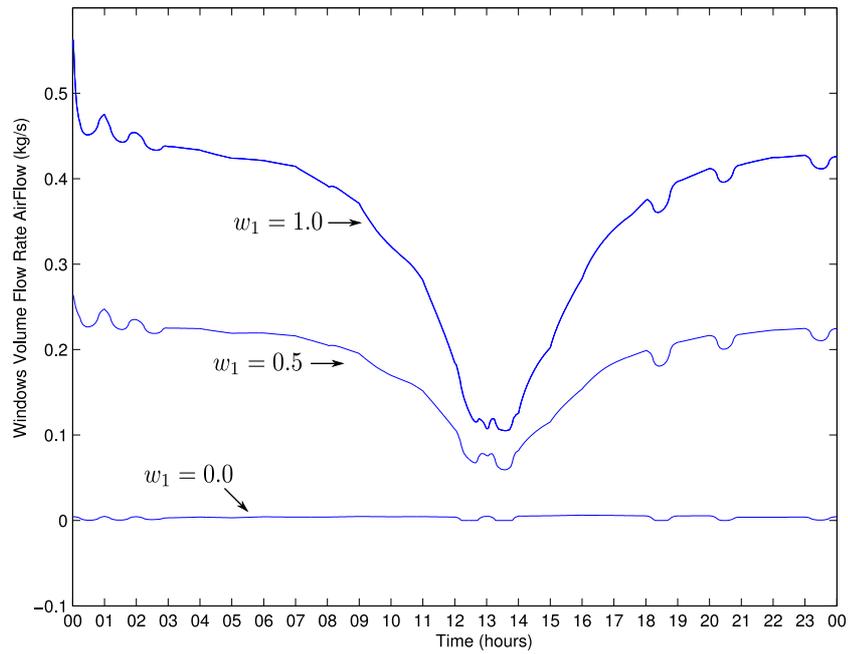
Figure 5.14: Airflow rate when two openings (door and window) are fully *open*, between 11h00 and 15h00.

modeling. Natural ventilation is activated when both window  $w_1$  and door  $D_1$  are *open*. Figure 5.14 shows the airflow rate obtained from EnergyPlus when both openings are fully *open* ( $wof_1, d_1$ ) = (1, 1) the first day between 11h00 and 15h00. The airflow is then uniformly quantized over  $\dot{M}_v$ , with  $|\dot{M}_v| = 100$ .

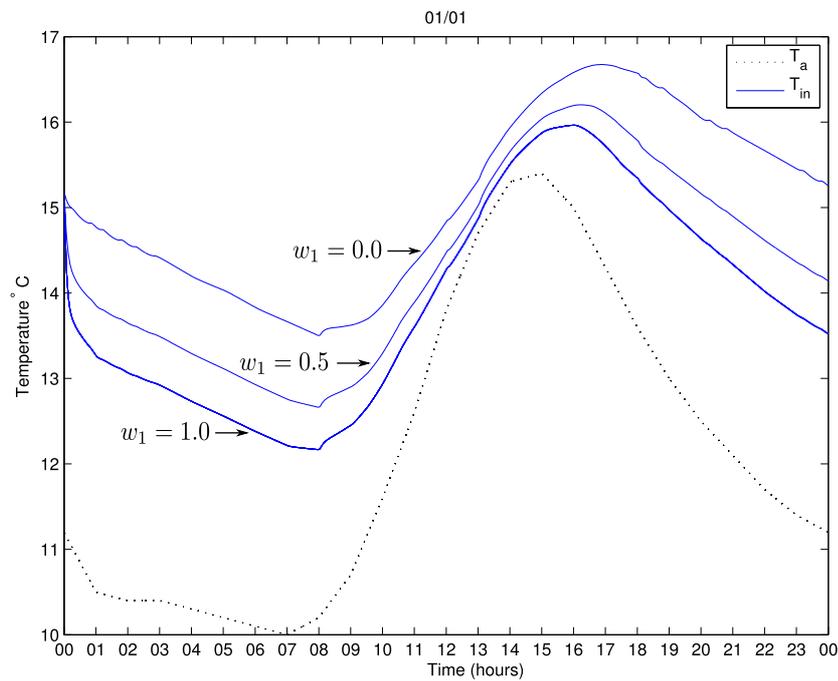
Considering a more sensible window opening, Figures 5.15a and 5.15b show the results of how the airflow rate and temperatures change with different opening factors for just window  $w_1$  (hold constant, during the entire simulation interval), with everything else *closed*. The airflow is significantly lower, and is proportional to the temperature difference between  $T_{in}$  and  $T_a$ .

### 5.2.5 Model Performance Evaluation

The simulation execution with the BCVTB and EnergyPlus, denoted by  $\chi_{EP} = (\tau, l, x)_{EP}$ , is used as validating reference for model performance evaluation. This execution is compared with the execution of the corresponding context-based model implemented using MATLAB, denoted by  $\chi_{Model} = (\tau, l, x)_{Model}$ , while sharing continuous and discrete input variables  $U = U_C \cup U_D$  between both simulations. Since the purpose of the context-based model is its prediction capability, performance is



(a) Airflow rate.



(b) Indoor ( $T_{in}$ ) and Outdoor ( $T_a$ ) temperatures.

Figure 5.15: Airflow and indoor temperature depending on different opening factors for window  $W_1$  (everything else *closed*)

evaluated on this aspect. For the execution of a hybrid automaton, the sequence of the prediction errors in continuous evolution is given by:  $\varepsilon_x(t) = x(t)_{Model} - x(t)_{EP}$ . A model is said to be “good” if it produces “small” absolute prediction errors as compared to the simulation. The simulation uses  $n$  discrete time-steps, and to quantify this measure, the mean **absolute error** (MAE) is used, given by:

$$\text{MAE} = \frac{1}{n} \sum_{t=t_1}^{t_n} \|\varepsilon(t)\|$$

considering that the hybrid time set and discrete evolution is the same for both executions with  $\tau_{EP} = \tau_{Model}$  and  $l_{EP}(t) = l_{Model}(t)$ ,  $\forall t \in \tau$ .

The mean absolute error is applied to a reduced dimension of  $x$ , by considering only the temperature  $T_{in}$  and its associated prediction error  $\varepsilon_{T_{in}}$ , due to its relevance for control. Following Lin et al. [108], this study requires that models reproduce observed data with a prescribed degree of accuracy, quantified by the following bound:

$$\text{MAX} = \|\varepsilon_{T_{in}}\|_{\infty} \leq 1.6^{\circ}\text{C}$$

with  $\varepsilon_{T_{in}} = [\varepsilon_{T_{in}}(t_1), \dots, \varepsilon_{T_{in}}(t_n)]$ . This condition derives from the fact that temperature distribution inside a room can have a spatial variation higher than 1.6 °C (3 °F) [100]. Requiring more accuracy is, therefore, unreasonable. Simulation results are presented and discussed in Section 6.2

### 5.2.6 Simulation Files

To enable the reproducibility of the results described in this Chapter, the simulation files used in this dissertation for Matalab (RL and Context-based modeling), EnergyPlus (including the simulated building “.idf” description), and the the BCVTB project, are made available at the following location:

<http://mediawiki.isr.ist.utl.pt/wiki/CtxtBasedMod>.

## 5.3 Summary

This chapter described the simulation setup environment developed to evaluate both the reinforcement learning methods proposed in this dissertation and the context-based model example given in the previous chapters. The *Bang-bang Heater* and *Set Point Heater* problems were simulated using the Matlab simulation environment.

This chapter described how the occupant was simulated (with schedules and comfort preferences) and how the results were evaluated. Two situations were described: a situation where the occupant is unrealistically persistent, and a more realistic simulation, where the occupant does not always provide the required feedback. The chapter then described an improvement to both methods, where rewards are delayed in order to account for the thermodynamic inertia of the thermal zone.

Following this set of experiments, the chapter described in detail the simulation setup for a context-based model. To evaluate the performance, the execution of the model was compared with the simulation outputs of EnergyPlus, which is a simulator that creates and uses models with much higher complexity.

# 6

## Results

*“The environment is where we all meet; where all have a mutual interest; it is the one thing all of us share.” – Lady Bird Johnson (First Lady of the United States).*

This chapter contains the simulation results of the experiments described in Chapter 5. The following sections show the results of the different experiments performed to evaluate the proposed RL methods, and the execution of the proposed context-based model. The results are discussed and conclusions are drawn for each simulation.

### 6.1 Reinforcement Learning Simulation Results

This section presents the simulation results for the simulation setup described in Section 5.1 for the *Bang-bang Heater* and *Set Point Heater* problems, with the performance metrics described in Section 5.1.3.

### 6.1.1 The Bang-Bang Heater

After executing each simulation (250 episodes) 30 times, the average performance metrics listed in Table 6.1 were obtained. Results show that the occupant is present in the TZ during an average of 100  $\Delta_t$  time instances. During this period, the occupant is comfortable 74.77% of the time and, when uncomfortable, remained in that state for an average duration of  $\Delta_{\text{mean}} = 2.40 \Delta_t$ . To guarantee this comfort level, the heater was *on* during 28% of the 24-hour (300  $\Delta_t$ ) simulation interval.

Parameter	( $\Delta_t$ )
Avg. Occupancy	100.00
<i>avg_tComf</i>	74.77
<i>avg_Δmin</i>	1.00
<i>avg_Δmax</i>	6.07
<i>avg_Δmean</i>	2.10
<i>avg_thCost<sub>1</sub></i>	83.00/ <i>N</i>

Table 6.1: Average time of occupancy and performance metrics for comfort and energy. Values for the following averages: comfort; minimum, maximum and mean size of the intervals of discomfort; total time the heater is *on*.

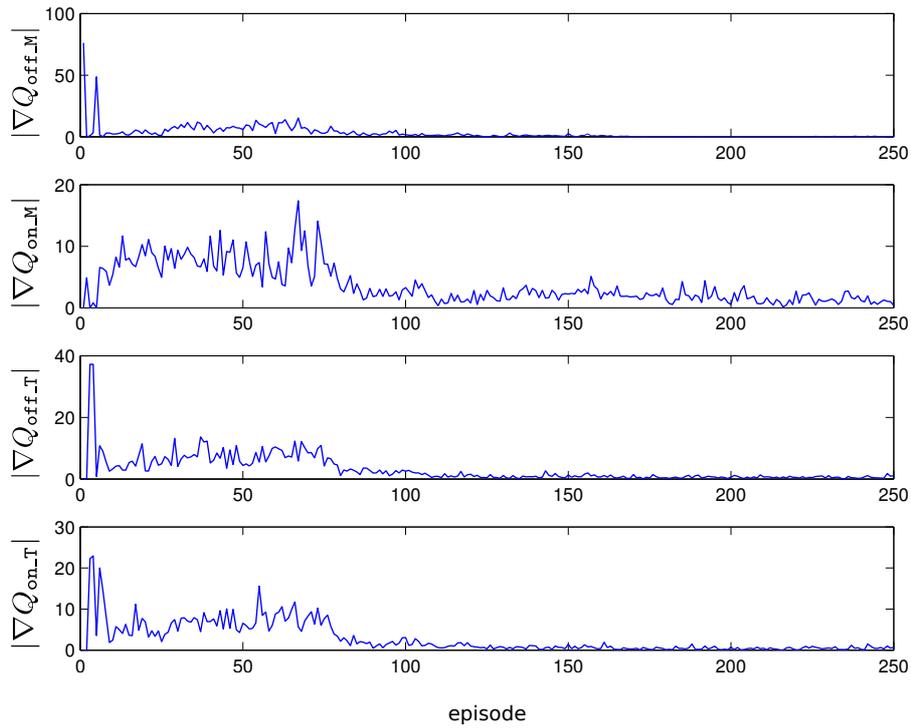
### Q-Values and Average Reward

The convergence of the Q-values, and the average reward received throughout 250 episodes are shown in Figure 6.1. The Q-value updates converge to zero, as expected and shown in Figure 6.1a, and the average reward, shown in Figure 6.1b, converges to  $\bar{r} \approx -0.05$ . Results show the average reward decreasing with exploration up to 50, when  $\varepsilon_{\text{Rate}}$  is set to zero, and increasing afterwards.

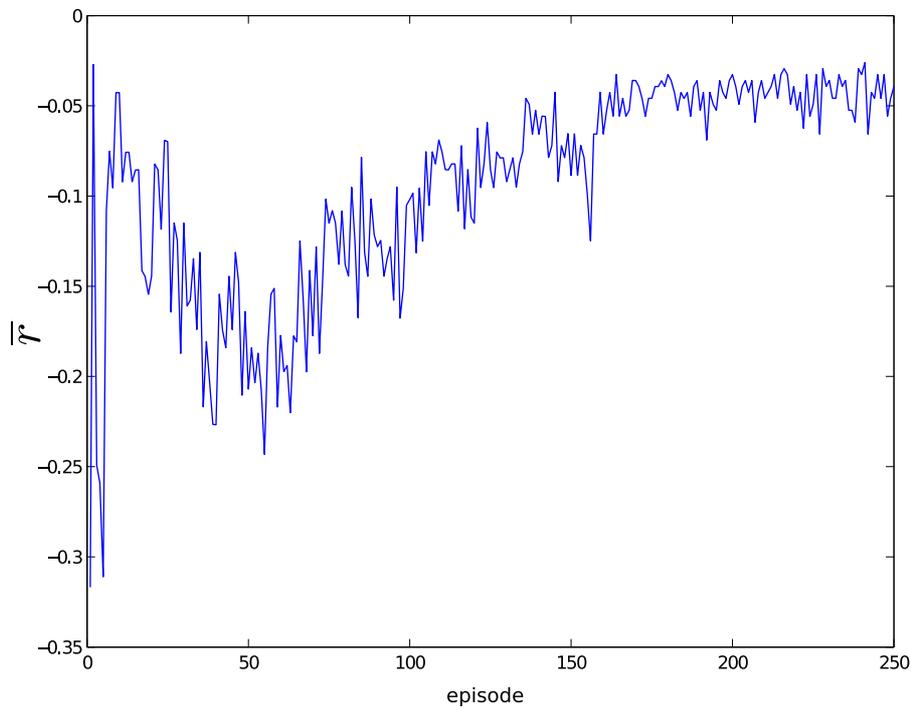
### Policy Execution

The state of the heater, the state of the occupant, and the temperature of the TZ after executing the learned policy during the last episode are shown in Figure 6.2. Line (1) represents the ideal temperature desired by the occupant. Results show that during the execution, the BAS switches the heater *off* as often as possible to minimize the use of energy, while switching it *on* to guarantee comfort during working hours

## 6.1. REINFORCEMENT LEARNING SIMULATION RESULTS



(a) The convergence of Q-values:  $|\nabla Q| \rightarrow 0$ .



(b) Average reward.

Figure 6.1: The convergence of the Q-values and the average reward received throughout 250 episodes.

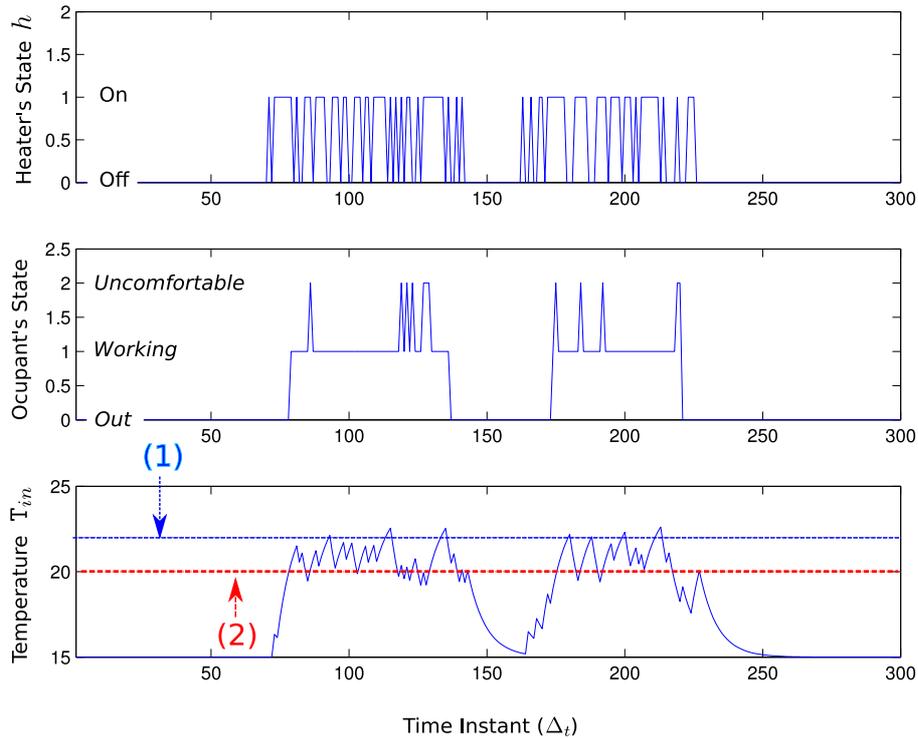


Figure 6.2: Policy application results. Line (1) represents the comfortable set point temperature (22 °C), and line (2) represents the lower temperature limit of comfort. The BAS learns how to maintain the temperatures within limits that explore the boundaries of comfort, while also saving energy when the thermal zone is unoccupied.

(anticipating the occupant's interaction). The occupant remains in the *Uncomfortable* state while the temperature is below the comfort limit, represented by the line (2).

### 6.1.2 Simulation Results With the Uncomfortable State Redefined

Performance metrics and execution results were obtained for the *Bang-bang Heater* problem with the occupant's *Uncomfortable* state redefined. Performance metrics, averaged over 30 executions, are shown in Table 6.2 and the execution of the learned policy is shown in Figure 6.3. Simulation results show that when the occupant becomes less persistent in interacting with the HVAC, there is a significant penalty in terms of comfort. There is a decrease in the average time the occupant is comfortable in the *Working* state, and an increase in the mean time the occupant is in the *Uncomfortable* state. **This concludes the idea that the amount of feedback**

received by the BAS, in many real-life situations, may not be enough to maintain comfort requirements.

For practical applications, the BAS should be able to learn from the sparse interactions that occur during the occupant’s regular schedule. However, an important limitation was identified in the current methodology that limits this capability. We verified that when the occupant interacts with the system in a certain state, and the heater is toggled to the *on* position, the selected action does not have an immediate effect on the TZ temperature due to thermal inertia. Consequently, the occupant will most likely remain uncomfortable for a certain time interval. As a result, not only does the BAS perceive that there is no immediate reward for comfort associated with the *Toggle* action, but it also receives a penalty for using energy in that state. Therefore, the BAS will eventually select the *Maintain* action in that state, thus compromising future rewards. This conflicting situation has a negative impact on learning performance.

In this section, we conclude that the reward function, in its current form, is not fully effective to optimize comfort, since it does not take into account the time delay associated with selected actions. The *Bang-bang Heater* RL algorithm needs to be modified to take this delay into account.

Parameter	$(\Delta_t)$
Avg. Occupancy	101.00
<i>avg_tComf</i>	58.60
<i>avg_Δmin</i>	1.00
<i>avg_Δmax</i>	13.53
<i>avg_Δmean</i>	4.18
<i>avg_thCost<sub>1</sub></i>	74.33/ <i>N</i>

Table 6.2: Average time of occupancy and performance metrics obtained with the occupant not interacting with the BAS so persistently, when uncomfortable (more realistic simulation). By comparing with Table 6.1, we can verify a significant penalization in comfort.

### 6.1.3 *Bang-bang Heater* Results Using Less States

The simulation results for the *Bang-bang Heater* show a significant improvement in comfort performance when using less states. Performance metrics were obtained for

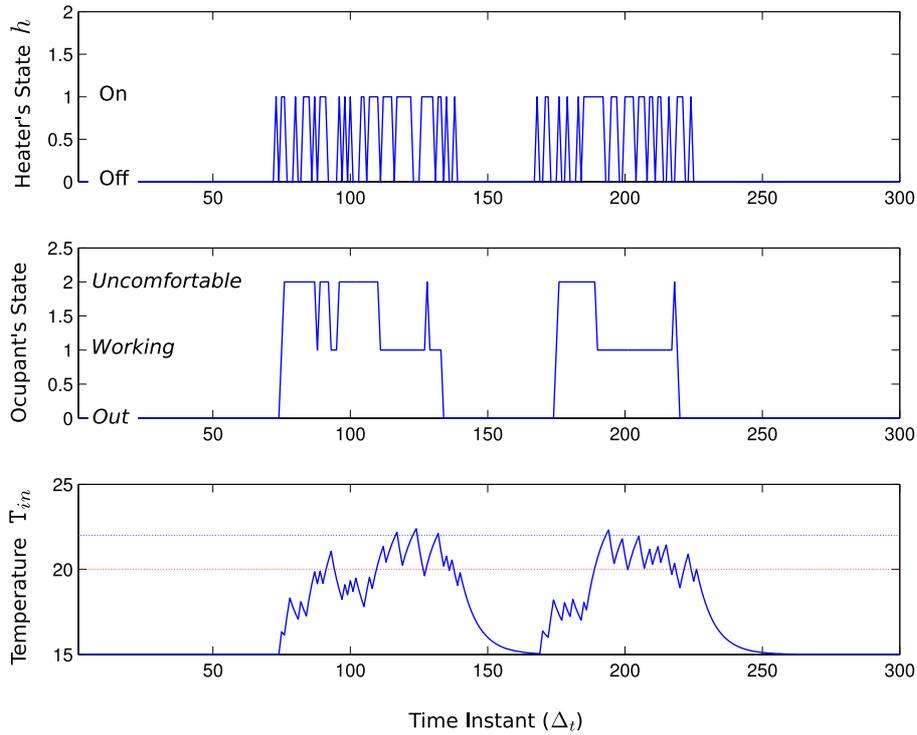


Figure 6.3: Policy application results showing a penalization in comfort, when the occupant is not persistently interacting with the heating system.

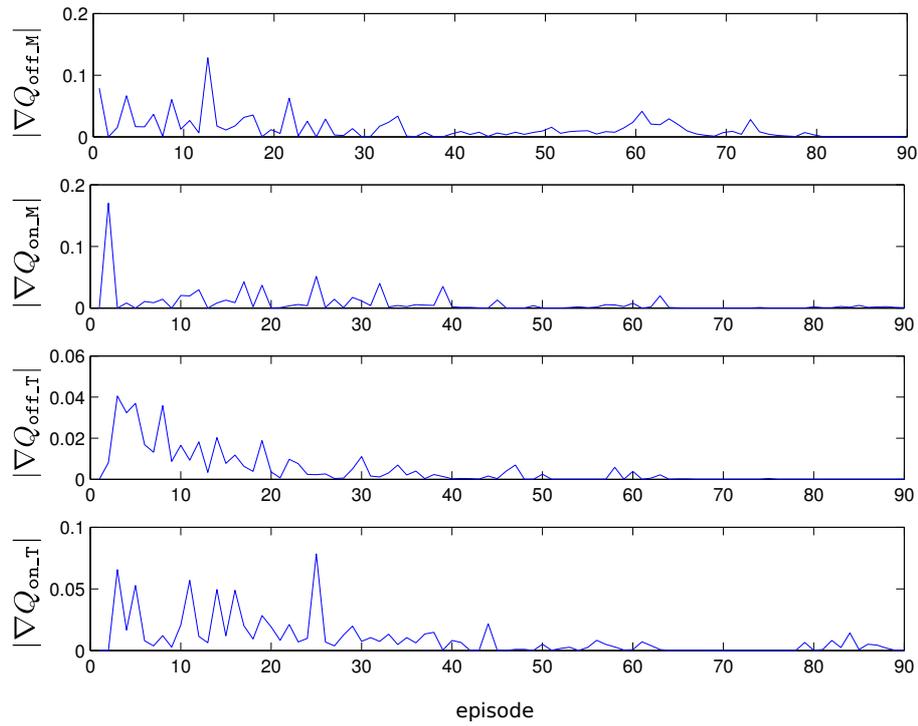
both situations: when the occupant persistently interacts with the HVAC system, and when the occupant's *Uncomfortable* state is redefined. In both cases, the results, listed in Table 6.3, show a significant increase in the amount of time the occupant is comfortable in the TZ, as compared to the previous results shown in Tables 6.1 and 6.2.

Figure 6.4 shows the convergence of Q-values and the average daily reinforcement for the worst case, when the occupant does not interact as often. Results show that the learning time decreased significantly to approximately 50 episodes. The execution of the learned policy is shown in Figure 6.5. Results show that the learned policy is effective in guaranteeing comfort conditions, while still maintaining the heater in the *off* state, during the time instances the TZ is usually vacant.

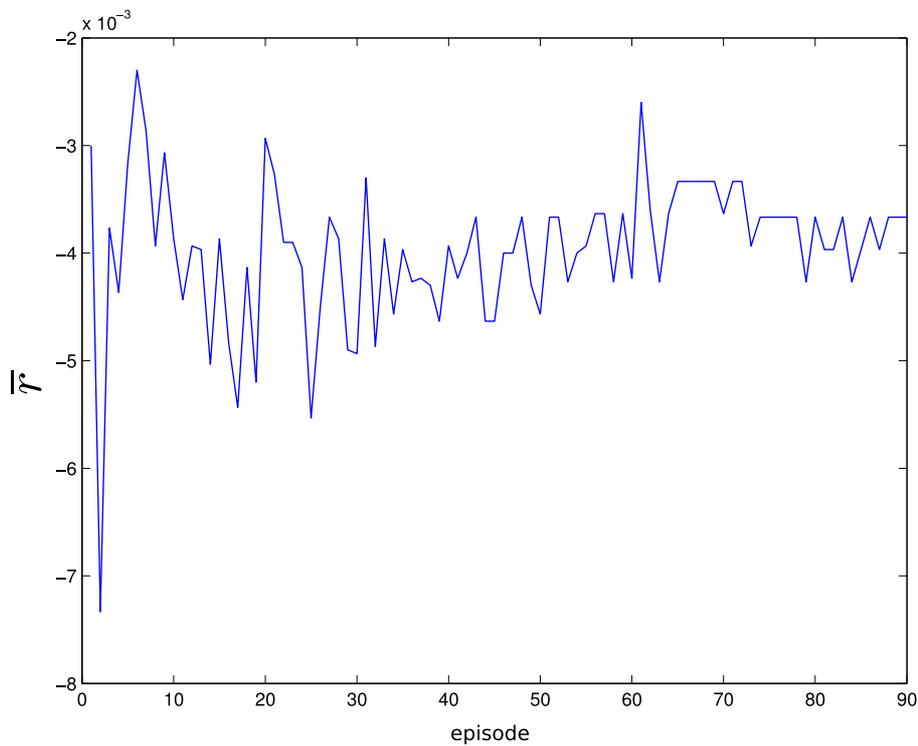
#### 6.1.4 The *Set Point Heater* Heater

This section discusses the simulation results of the *Set Point Heater*. The average reward received in each episode is shown in Figure 6.6 (converging to  $\bar{r} \approx -0.04$ ).

## 6.1. REINFORCEMENT LEARNING SIMULATION RESULTS



(a) The convergence of Q-values.



(b) Average reward.

Figure 6.4: The convergence of Q-values and the average reward, when using less states. The policy converges faster to a more stable solution.

Parameter	$(\Delta_t)$	$(\Delta_t)$
	Persistent	Less Persistent
Avg. Occupancy	100.47	103.50
$avg\_tComf$	100.47	102.90
$avg\_Delta_{min}$	0.00	0.60
$avg\_Delta_{max}$	0.00	0.60
$avg\_Delta_{mean}$	0.00	0.60
$avg\_thCost_1$	$160.66/N$	$158.66/N$

Table 6.3: Average time of occupancy and performance metrics with less states for both situations: a persistent occupant, and a less persistent occupant (as described in Section 6.1.2). Comparing metrics with Tables 6.1 and 6.2, in both cases, results show that there is a significant improvement in comfort performance when using less states.

After learning, the policy execution results are shown in Figure 6.7. Results show that the BAS learns the occupancy pattern and tries to minimize the supply of heat, while keeping the TZ temperature within the limit of comfort. Due to the lack of interactions on the part of the occupant, the BAS learns how to reduce the supply of heat during the intervals when the TZ is usually unoccupied.

### Comfort vs Heating Cost

This section shows the effects trade-off between comfort and the cost of heating. The performance metrics (averaged over 30 simulations) for the *Bang-bang Heater* and *Set Point Heater* problems (with the occupant's state *Uncomfortable* redefined) are presented in Table 6.4 for different weights values  $w_1$  and  $w_2$ . Results show that, although there no direct linear map between metrics and weights, there is a direct proportional impact in the balance between comfort and the heating cost.

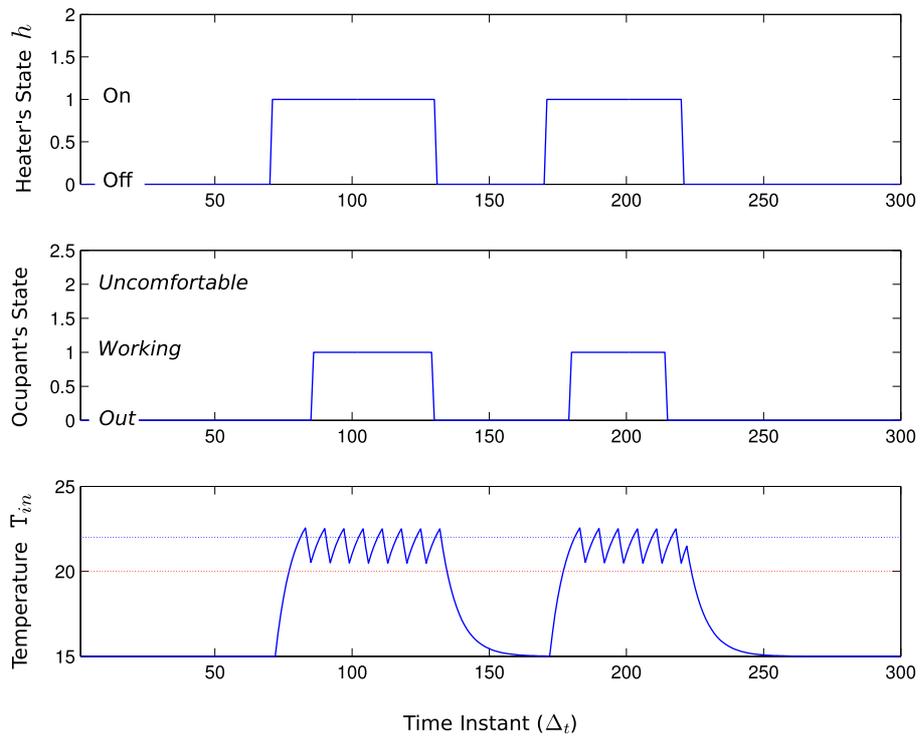


Figure 6.5: Policy application results with less states. The *Bang-bang Heater* is more effective in guaranteeing comfort conditions. The heater is *off* during the time instances the thermal zone is usually vacant.

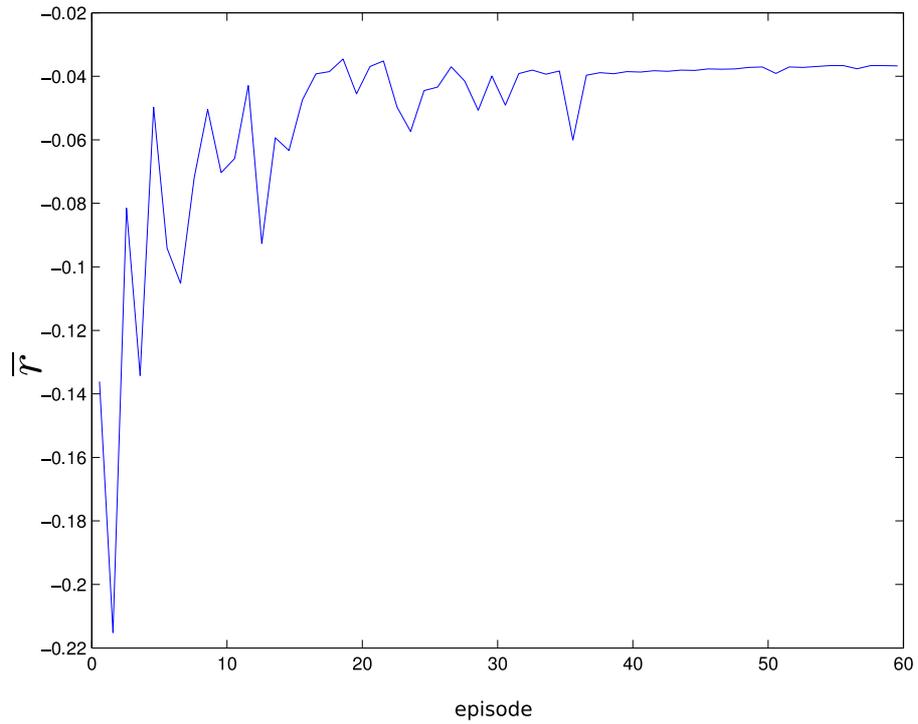


Figure 6.6: Average reward received daily with the *Set Point Heater*.

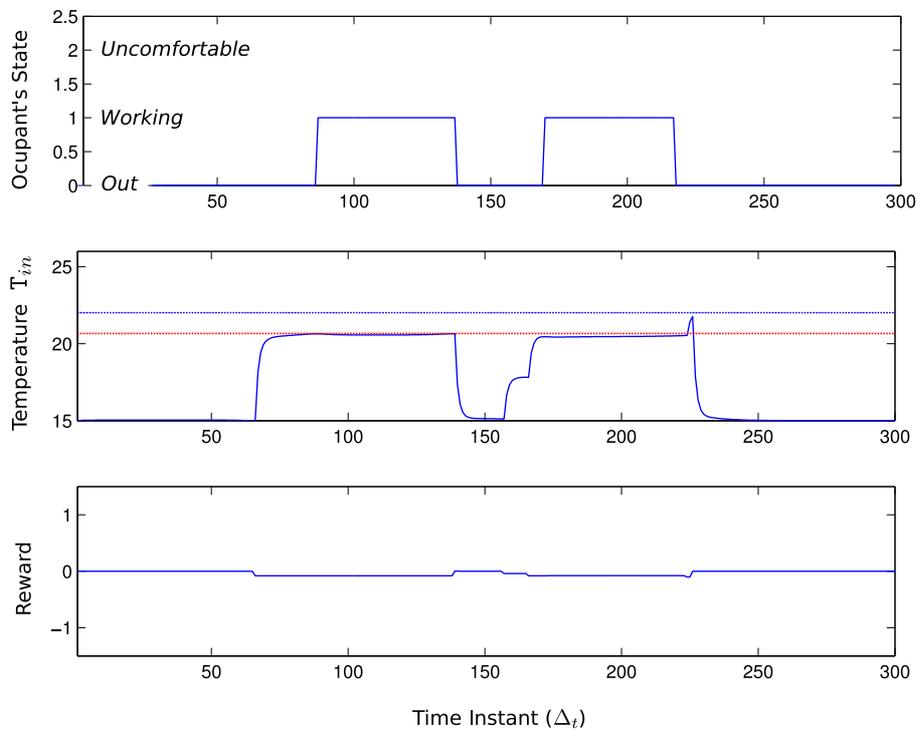


Figure 6.7: Policy execution results for the *Set Point Heater*.

6.1. REINFORCEMENT LEARNING SIMULATION RESULTS

$w_1$	$w_2$	$avg\_tComf$	$avg\_Δ_{min}$	$avg\_Δ_{max}$	$avg\_Δ_{mean}$	$avg\_thCost_1$
			$(Δ_t)$			$(Δ_t)/N$
<i>Bang-bang Heater</i>						
0.10	0.90	2.00	42.93	54.43	48.68	0.00
0.20	0.80	2.00	40.00	55.76	47.88	0.00
0.30	0.70	2.00	42.00	54.53	48.27	0.00
0.40	0.60	2.00	43.23	52.33	47.78	0.00
0.50	0.50	2.00	42.36	55.13	48.75	0.00
0.60	0.40	2.00	42.96	55.86	49.41	0.00
0.70	0.30	4.40	34.40	53.60	44.56	3.67
0.80	0.20	24.23	9.33	33.03	19.88	33.67
0.90	0.10	92.50	3.66	5.33	4.5	118.33
0.99	0.01	100.50	0.50	0.50	0.51	150.67
<i>Set Point Heater</i>						
						$thCost_2$
0.10	0.90	2.50	33.90	53.70	43.36	22.76
0.20	0.80	3.00	24.73	49.97	36.89	25.65
0.30	0.70	3.23	20.20	52.57	35.61	25.69
0.40	0.60	2.43	32.33	51.60	42.13	21.57
0.50	0.50	19.03	1.07	16.77	5.63	54.58
0.60	0.40	59.10	1.00	5.63	1.73	86.46
0.70	0.30	76.10	1.00	3.37	1.33	100.84
0.80	0.20	85.20	1.00	2.06	1.13	113.16
0.90	0.10	95.53	0.87	1.13	0.96	126.38

Table 6.4: Performance metrics obtained with different values for  $w_1$  and  $w_2$ , showing the trade-off between comfort and heating cost in both *Bang-bang Heater* and *Set Point Heater* problems.

## 6.2 Context-based Thermodynamic Modeling

The example given in Section 4.4.2 describes a context-based model with guards that are conditioned by occupancy, solar gains, the state of the heater and the airflow rate shown in Figure 5.14. Using the calculated RC values listed in Table 6.5, the execution of this example was simulated for the 1st day of January over a 48-hour time horizon between  $\tau_0 = 00\text{h}$  and  $\tau_N = 48\text{h}$ . Figure 6.8 shows the outdoor temperature  $T_a$  and the resulting indoor temperatures obtained from the model and EnergyPlus, denoted respectively by  $T_{in\_Model}$  and  $T_{in\_EP}$ . The hybrid time set and discrete trajectory are listed in Table 6.6. Results show that both trajectories for  $T_{in}$  and  $T_{in\_EP}$  overlap throughout most of the execution interval. Figure 6.9 shows the associated histogram of errors. The predictive accuracy of the model is within acceptable limits, with  $MAE = 0.1134^\circ\text{C}$  and  $MAX = 0.9018^\circ\text{C}$ .

### Discussion

The correlation shown in Figure 6.8 resulted from an iterative and incremental design where the full building model was incrementally created and simulated from its simple envelope to the full simulation example. Although the applied heat balance principles are the same as EnergyPlus and, therefore, this correlation is somehow expected, the heat transfer model presented in this dissertation is significantly more reduced than the models used by EnergyPlus. Moreover, not all of the heat balance equations used in the context-based thermodynamic model are the same as the ones used in EnergyPlus. In fact, some of these equations (e.g., the Navier-Stokes equations used to simulate computational fluid dynamics of the airflow network), cannot be expressed using simple linear state-space equations. Therefore, the complexity of the models used by EnergyPlus limit the application of most classical control techniques. The proposed context-based model, on the other hand, uses simple LTI models to simulate airflow (with all its limitations) and other thermodynamic behaviors such as the indoor solar-air transmission function, that uses a Butterworth filter as a very rough approximation to more elaborate solar radiation calculations. The resulting LTI models are much simpler and appropriate for real-time control applications.

There are not many articles in the literature that show how RC models can be used to model variations in the TZ's thermodynamic behavior. An additional procedure to validate the importance of using *context*, as a model simplification strategy, could be comparing the performance of a context-based model against the situation

Resistor	$(\text{K W}^{-1}) \times 10^{-4}$	Capacitor	$(\text{J K}^{-1}) \times 10^6$
$R_1$	3.146	$C_1$	22.188
$R_2$	3.562	$C_2$	8.629
$R_3$	8.656	$C_3$	9.506
$R_{int}$	12.314	$C_4$	19.301
$R_{ext}$	2.088	$C_g$	2.703
$R_{g1}$	20.152	$C_h$	0.234
$R_{g2}$	20.152	$C_Z$	0.523
$R_W$	238		
$R_{ih}$	8.332		
$R_{gint}$	69.444		

Table 6.5: RC values for the model that was used in the application example (represented in Figure 4.8).

Time Interval $I_i$ (h)	$l_i(\cdot)$	Time Interval $I_i$ (h)	$l_i(\cdot)$
$I_1 = [00\text{h}00, 08\text{h}00]$	$(l_1)$	$I_{11} = [15\text{h}00, 17\text{h}00]$	$(l_7)$
$I_2 = [08\text{h}00, 08\text{h}01]$	$(l_2)$	$I_{12} = [17\text{h}00, 18\text{h}00]$	$(l_4)$
$I_3 = [08\text{h}01, 09\text{h}44]$	$(l_4)$	$I_{13} = [18\text{h}00, 32\text{h}00]$	$(l_1)$
$I_4 = [09\text{h}44, 09\text{h}45]$	$(l_2)$	$I_{14} = [32\text{h}00, 32\text{h}01]$	$(l_2)$
$I_5 = [09\text{h}45, 12\text{h}00]$	$(l_3)$	$I_{15} = [32\text{h}01, 33\text{h}00]$	$(l_4)$
$I_6 = [12\text{h}00, 13\text{h}00]$	$(l_1)$	$I_{16} = [33\text{h}00, 34\text{h}00]$	$(l_5)$
$I_7 = [13\text{h}00, 13\text{h}01]$	$(l_2)$	$I_{17} = [34\text{h}00, 42\text{h}00]$	$(l_6)$
$I_8 = [13\text{h}01, 14\text{h}22]$	$(l_3)$	$I_{18} = [42\text{h}00, 42\text{h}01]$	$(l_4)$
$I_9 = [14\text{h}22, 14\text{h}23]$	$(l_2)$	$I_{19} = [42\text{h}01, 48\text{h}00]$	$(l_1)$
$I_{10} = [14\text{h}23, 15\text{h}00]$	$(l_4)$		

Table 6.6: The hybrid time set and discrete trajectory of the running example.

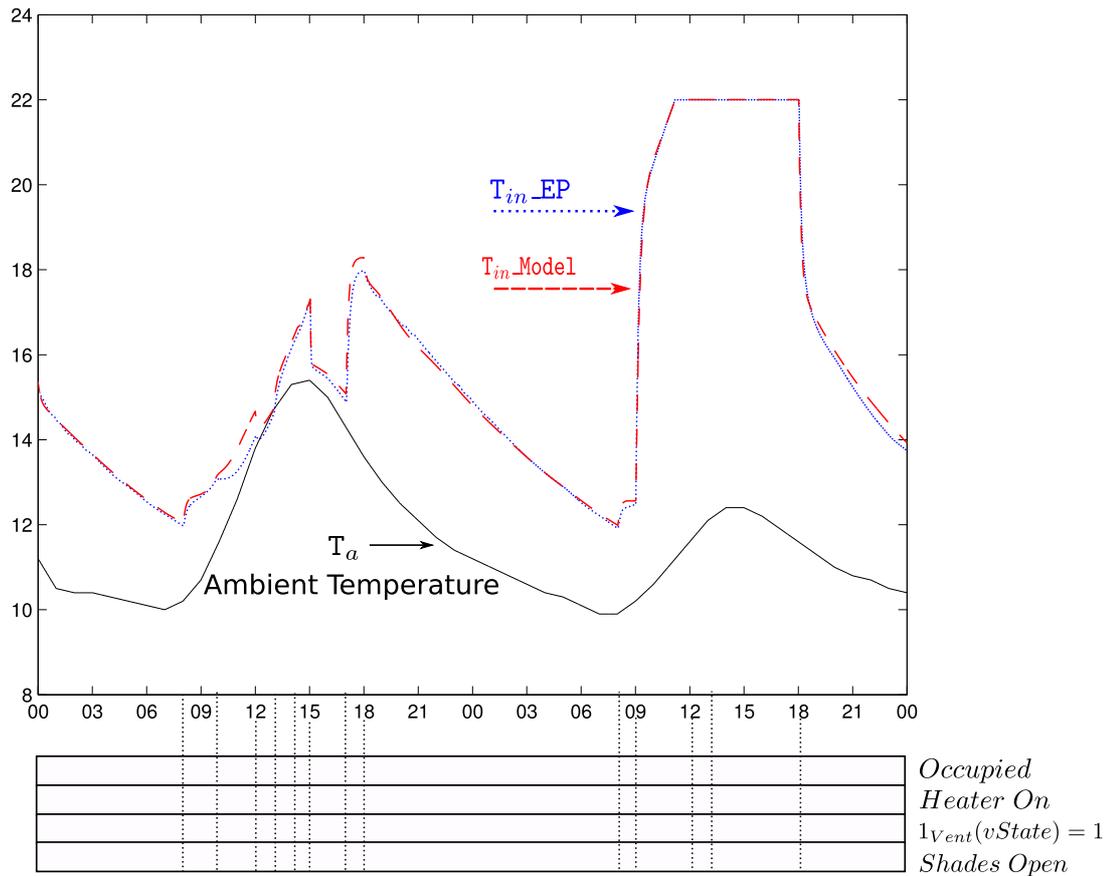


Figure 6.8: Execution and simulation results showing the outdoor temperature ( $T_a$ ) and indoor temperatures obtained using the model ( $T_{in\_Model}$ ) and with simulation using EnergyPlus ( $T_{in\_EP}$ ). Both curves  $T_{in\_Model}$  and  $T_{in\_EP}$  overlap almost exactly.

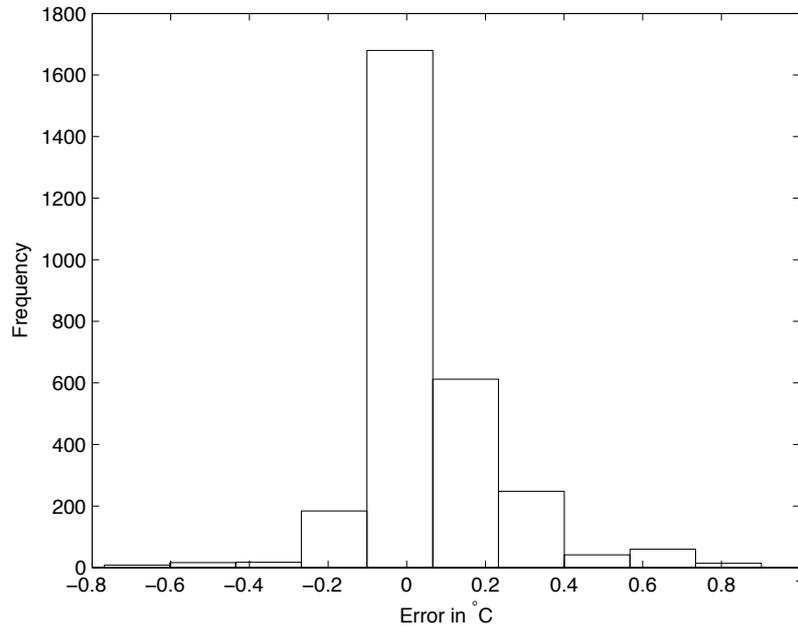


Figure 6.9: Histogram of the error distribution for the example in Figure 6.8.

where just a single RC model is used for every configuration of the TZ. However, the importance of modeling the variations in the TZ's thermodynamic behavior can be inferred directly from Figures 5.11, 5.14 and 5.15b, where model parameters are shown to have a direct implication on the evolution of the indoor temperature. A single RC model, with the same number of parameters as the LTI models used by a similar context-based model, will not correlate so well with EnergyPlus because it would have to find a compromise to model all configurations with a single set of RC parameters. Context-based models have the advantage of using several LTI models with different RC parameters adjusted to each particular *context*. To represent the TZ with a single RC model, that has the same level of accuracy as a context-based model, more complexity must be added to the RC model and in some cases (e.g., when using variable resistances and switches), this level of accuracy is not even possible. This concludes that context-based models are a more flexible and accurate solution for representing variations in the thermodynamic behavior of a TZ, when compared to the RC models used in the past literature.

### 6.3 Summary

This chapter presented the simulation results of the experiments described in Chapter 5. Results for the reinforcement learning section show that the BAS is capable of learning, for both *Bang-bang Heater* and *Set Point Heater* problems, how to minimize the supply of heat while guaranteeing comfort. In both cases, the BAS is capable of exploiting the lower limits of the temperature that the occupant is capable of supporting, without interacting with the system. Several experiments were tested: a case where the occupant is persistent, and a more realistic situation, where the occupant does not always interact when uncomfortable. Results show that, in order to increase learning performance, the BAS must take into account the average number of interactions that occur in a certain time horizon, in order to account for the thermodynamic inertia of the thermal zone. Simulation results show a trade-off between comfort and the cost of heating, depending on the weights set in the reinforcement function. By setting  $w_1$  and  $w_2$ , the learning algorithm will tend to converge to a heating strategy that minimizes the penalty (negative reward) according to those weights. After a certain learning period, the BAS adjusts the zone temperature according to the occupant's preferences and occupancy schedules. This temperature can be adjusted to a minimum set point where the occupant, usually, no longer complains.

The second section of this chapter showed the execution of the context-based thermodynamic model described in the example given in Chapter 4. The execution of the model was compared with the simulation outputs of EnergyPlus, which is a simulator that creates and uses models with much higher complexity. This fact validates the proposed context-based model, at least empirically, opening doors for future real-time control applications.

# 7

## Conclusions and Future Work

*“It seems to me that the natural world is the greatest source of excitement; the greatest source of visual beauty; the greatest source of intellectual interest. It is the greatest source of so much in life that makes life worth living.” – David Attenborough (English broadcaster and naturalist).*

People spend a substantial amount of their time inside buildings which are responsible for 40% of global energy usage and 33% of the total greenhouse gas emissions. Advances in technology for the architecture, construction, and building operations are foreseen to promote up to 70% of energy savings in this sector by the year 2030. To meet these requirements, smart buildings are expected to integrate more efficient building technologies and use state-of-the-art machine-learning techniques in their building automation systems for various building applications. These techniques will be used to model and predict building variables, and for evaluating the quality of automated decisions by controlling internal systems and effectors while trying to satisfy certain goals such as reducing energy wastage, operating costs and occupant discomfort. However, there is still a huge set of challenges for machine-learning that are not easy to address with current building technologies. Smart buildings are currently still

an open problem for research.

## 7.1 Reinforcement Learning for HVAC Control

Buildings are very large complex systems that integrate several energy systems. Since the heating, ventilation and air conditioning system is among the most energy-demanding systems in a building, this dissertation started by targeting this system and exploring how a well-known reinforcement learning algorithm - Q-learning - could be used to optimize the operation of the heating/cooling system. This research was guided by the following idea: since different building spaces can have different conditioning requirements (depending on occupancy schedules, comfort preferences, environmental conditions associated with lighting, equipment, solar gains, etc.), a building automation system should be able to automatically explore and find the optimal requirements for each space. Temperature set points for example must be set according to an optimization strategy that takes into account goals such as energy savings and comfort. Therefore, to execute this idea, a novel reinforcement learning method was proposed for learning the optimal temperature set point scheduling strategy. The heating scenario was assumed as an example, and a reward function was devised to penalize the building automation system when the heating was *on* (proportional to the amount of energy used for heating), or when an occupant acted upon the heating user control interface to signal that he/she is feeling “cold”.

Two application examples were given for two different problems of low-level heating control: *Bang-bang Heater* and *Set Point Heater*. The *Bang-bang Heater* problem assumes that zonal heating is controlled by alternately switching the activation state of a heater (there is no temperature set point control). This problem is solved with straightforward Q-Learning using discrete states and actions. The *Set Point Heater* problem on the other hand assumes that the building automation system is capable of learning how to set temperature set points. The actions and states are continuous values: the zonal temperature set point is controlled (or changes) in minutely small measures. Therefore, a continuous state and action Q-Learning algorithm was selected and customized to solve this problem. For both problems simulation results showed that the building automation system was capable of learning how to adjust the heating system for energy savings, while also taking into account the occupant’s preferences and schedules (also learned through experience and observation) in order to keep the occupant comfortable and satisfied with the building environment. When

considering the amount of energy used globally for heating, it is this researcher's belief that both control strategies can have a significant impact on reducing the global energy bill and greenhouse gas emissions.

Following these results the dissertation proceeded with a discussion on the presented reinforcement learning methodologies. This discussion included the fact that reinforcement learning algorithms present limitations when it comes to explaining the system's automated actions to humans. Since policies are encoded in Q-Value tables or in neural network weights, presenting an explanation to a human operator is not a straightforward process. Moreover, the reinforcement learning problem becomes computationally expensive to solve as more states are added to represent the building environment and the interaction with its occupants with more accuracy. Therefore, the research efforts for this dissertation became focused on finding alternative solutions to represent the state of the building environment taking into account that this environment may include an intractably large set of variables. Thereupon, the research problem changed towards the goal of finding accurate models for the building environment which can be used, among other things, for building simulation, model predictive control, and for synthesizing building automation plans (using task planning algorithms) that are explicable and predictable to humans.

## 7.2 Context-based Thermodynamic Models

Developing efficient computational models that accurately describe the thermodynamics of building spaces to optimize the energy efficiency of the heating, ventilating, and air conditioning systems is of the utmost importance. For this application, a significant part of the literature over the past twenty years has been using time-invariant reduced models to represent the thermodynamic connection between different thermal zones. A limitation of this solution lies in the fact that building environments are not time-invariant. Indeed, there are many conditions in a thermal zone that change over time. The occurrence of events such as doors, windows and blinds being opened or closed, can drastically affect the underlying processes that govern the dynamics of temperature evolution of building spaces, rendering these standard time-invariant models less effective for control and prediction. Recent literature has shown that a single reduced model cannot efficiently cover all the different contextually-dependent conditions. Such a model would have to be sufficiently complex to describe the thermodynamics of a thermal zone for every situation making it computationally ex-

pensive for real-time applications. To solve this problem this thesis describes the evolution of temperatures in building spaces by capturing discrete changes in the thermal zone through the concept of *context*. Context is defined as a discrete state conditioned by a set of variables such as the opening factor of windows, shades, natural ventilation, and other variables that affect the thermodynamics of building spaces. A context-based model contains a set with several linear time-invariant models, each effective in representing the thermal behavior of a building space in an associated set of contexts. Therefore simpler models can be selected according to certain specific conditions that are relevant for the building thermodynamics during a certain time frame. The descriptions in this thesis denote a context-based model implemented as an open hybrid automata, whose transitions between different contexts and models are described by a set of context-connecting edges, with guards and effects associated with the domain and range of the validity of each context and associated model.

The merit of this study's approach lies in the fact that many building models are expected to become conditioned by more contextual-information as building environments become more observable (e.g., to include the state of windows, doors and blinds) for computation. Therefore, it becomes important the development of models that can accurately describe the coexistence and interaction between the discrete and continuous dynamics of *context* and temperatures. Up until now, most of these environment changes were treated as disturbances in the literature. Considering the current available options for thermodynamic modeling: complex linear and nonlinear models that are accurate and suitable for off-line simulation but computationally expensive; simple lumped models that are appropriate for real-time applications but have limited accuracy. This dissertation advances the state of the art by showing that context-based models combine the best of both options. By using the concept of *context* several reduced models can be used to represent the thermodynamics of a building space while retaining similar prediction capabilities as models with higher complexity. To clarify and validate this thesis' proposal a simulated example was presented using Matlab and EnergyPlus, integrated with the Buildings Controls Virtual Test Bed software environment. Using the outputs of EnergyPlus as a ground truth, results show that the context-based model can effectively predict the evolution of temperatures in a simulated zone through different context changes if context is observable and synchronized with the model.

## 7.3 Future Work

It is this researcher's belief that this dissertation establishes a landmark connection between hybrid systems, thermodynamic modeling, and context-based reasoning. This connection opens an enormous amount of possibilities for research. For system modeling, context-based models can be extended to include other environment variables such as illumination, CO<sub>2</sub>, humidity, energy and power. These models can also be used to model other building systems such as elevators, mechanical ventilation and energy storage devices. A direction for research includes automatic identification of context-based models based on the observation of environmental variables, and the application of these models for model-predictive control. Furthermore, we speculate that by pairing models with contextual information machine-learning and reasoning algorithms can be used to find optimal conditioning plans for the thermal zones. Context-based models can be used to provide building managers and occupants with a better insight into building thermodynamics and they represent an important research step towards the possibility of providing humans with better explanations on how thermal zones are conditioned. We envision an application output that would show, for example, that by leaving a door open, a thermal zone can be heated by taking advantage of solar gains in another thermal zone; or that a certain window should be closed for energy savings. To accomplish this level of planning and reasoning, machine-learning and planning algorithms in hybrid domains need suitable representations for the building environment and we believe that context-based models are suitable candidates. The following sections include examples of recent work and options of research that can extend the work presented in this dissertation towards these objectives.

### Planning and Model Checking in hybrid domains

- Bogomolov, Sergiy and Magazzeni, Daniele and Podelski, Andreas and Wehrle, Martin. (2014) Planning as model checking in hybrid domains. In: Proceedings of twenty-eighth AAAI Conference on Artificial Intelligence and the twenty-sixth Innovative Applications of Artificial Intelligence Conference (AAAI 2014) : 27-31 July 2014, Québec City, Québec, Canada, Vol. 3. Palo Alto, Calif., S. pp 2228-2234.
- Bogomolov, Sergiy V; Magazzeni, Daniele; Minopoli, Stefano; Wehrle, Martin,

PDDL+ Planning with Hybrid Automata: Foundations of Translating Must Behavior. In: Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling (ICAPS 2015): 7-11 June 2015, Jerusalem, Israel, pp. 42-46.

- Nicolò Giorgetti, George J. Pappas and Alberto Bemporad, Bounded Model Checking in Hybrid Dynamical Systems, Proceedings of the 44th IEEE Conference on Decision and Control, and the European Control Conference, December 2005, Seville, Spain, pp. 672-677.
- Thomas A. Henzinger, Pei-Hsin Ho, Howard Wong-Toi, HyTech: A Model Checker for Hybrid Systems, Software Tools for Technology Transfer, 1997, 1, pp. 460-463.

## Control with Hybrid Systems

- Daniele Corona, Alessandro Giua, Carla Seatzu, Stabilization of switched systems via optimal control, Nonlinear Analysis: Hybrid Systems, January 2014, Volume 11, pp. 1-10.
- E.F. Camacho, D.R. Ramirez, D. Limon, D. Muñoz de la Peña, T. Alamo, Model predictive control techniques for hybrid systems, Annual Reviews in Control, April 2010, Volume 34, Issue 1, Pages 21-31.
- Vassilis M. Charitopoulos, Vivek Dua, Explicit model predictive control of hybrid systems and multiparametric mixed integer polynomial programming, AIChE Journal, July, Vol. 62, No. 9, 2016, pp. 3441-3460.

## Model Identification

- Daniel L. Ly, Hod Lipson, Learning Symbolic Representations of Hybrid Dynamical Systems, Journal of Machine Learning Research, Dec 2012, Vol. 3. 1, pp 3585-3618.
- Kersting, S.; Buss, M., Online identification of piecewise affine systems. In: Proceedings of the International Conference on Control (CONTROL), July 2014, UKACC, pp. 86-91.

- M. Tabatabaei-Pour, M. Gholami, H. R. Shaker and B. Moshiri, Recursive Identification of Piecewise Affine Hybrid Systems, 9th International Conference on Control, Automation, Robotics and Vision, Singapore, 2006, pp. 1-6.
- Seyed Mojtaba Tabatabaeipour, K. Salahshoor, Behzad Moshiri, A k-plane Clustering Algorithm for Identification of Hybrid Systems, Proceedings of the 8th WSEAS Int. Conference on Automatic Control, Modeling and Simulation, Prague, Czech Republic, March 12-14, 2006, pp 334-339.
- Daniel L. Ly, Hod Lipson, Learning Symbolic Representations of Hybrid Dynamical Systems, Journal of Machine Learning Research, 13, 2012, pp 3585-3618.

### **Context-based Reasoning**

- Gary Stein, Aveline J. Gonzalez, Learning in context: enhancing machine learning with context-based reasoning. Journal of Applied Intelligence, 2014, Vol. 41, 3, pp. 709-724.
- Duen-Ren Liu; Chih-Kun Ke; Mei-Yu Wu, Context-based knowledge support for problem-solving by rule-inference and case-based reasoning, in: 2008 International Conference on Machine Learning and Cybernetics, 12-15 July 2008, Vol. 6, pp. 3205-3210.
- Akshay Krishnamurthy, Alekh Agarwal, John Langford, Contextual-MDPs for PAC-Reinforcement Learning with Rich Observations, CoRR, 2016.



## Bibliography

- [1] United Nations Framework Convention on Climate Change. Kyoto protocol reference manual. [http://www.unfccc.int/resource/docs/publications/08\\_unfccc\\_kp\\_ref\\_manual.pdf](http://www.unfccc.int/resource/docs/publications/08_unfccc_kp_ref_manual.pdf), 2008. Online, accessed 12-september-2016.
- [2] E.U. Parliament and E.U. Council. Directive 2010/31/EU of the European Parliament and of the Council on the energy performance of buildings. Official Journal of the European Union, L153/13, May 2010.
- [3] Rod Janssen. Towards energy efficient buildings in europe. Technical report, EuroAce - The European Alliance of Companies for Energy Efficiency in Buildings, June 2004. URL <http://euroace.org>. Online, accessed 12-september-2016.
- [4] John Guckenheimer and Julio M. Ottino. Foundations for complex systems research in the physical sciences and engineering. Technical report, Report from an NSF Workshop, September 2008. URL [http://www.siam.org/about/pdf/nsf\\_complex\\_systems.pdf](http://www.siam.org/about/pdf/nsf_complex_systems.pdf). Online, accessed 12-september-2016.
- [5] Andrzej Ziębik and Krzysztof Hoinka. *Energy Systems of Complex Buildings*. Green Energy and Technology. Springer-Verlag London, 2013.
- [6] American Physical Society. Energy future: Think efficiency report. Technical report, URL <http://www.aps.org/energyefficiencyreport/report/aps-energyreport.pdf>, 2008. Online, accessed 12-september-2016.
- [7] T. Weng and Y. Agarwal. From buildings to smart buildings – sensing and actuation to improve energy efficiency. *Design Test of Computers, IEEE*, 29(4): 36–44, August 2012.
- [8] Anna Kramers and Örjan Svane. ICT applications for energy efficiency in buildings. Technical report, KTH Centre for Sustainable Communications, Stockholm, Sweden, 2011. URL <http://www.diva-portal.org/smash/get/diva2:492659/fulltext01.pdf>. Online, accessed 12-september-2016.
- [9] Paula Rocha, Afzal Siddiqui, and Michael Stadler. Improving energy efficiency via smart building energy management systems: A comparison with policy measures. *Energy and Buildings*, 88:203–213, 2015.

## BIBLIOGRAPHY

---

- [10] J.K.W. Wong, H. Li, and S.W.Wang. Intelligent building research: a review. *Automation in Construction*, 14:143–159, 2005.
- [11] Consolidated View of the ETP SG, On Research, Development & Demonstration Needs in The Horizon 2020 Work Programme 2016–2017. Smart Grids: European Technology Platform (ETP), April 2015. URL <http://www.smartgrids.eu/>. Online, accessed 12-september-2016.
- [12] J. Kleissl and Y. Agarwal. Cyber-physical energy systems: Focus on smart buildings. In *Design Automation Conference (DAC), 2010 47th ACM/IEEE*, pages 749–754, June 2010.
- [13] Pervez Hameed Shaikh, Nursyarizal Bin Mohd Nor, Perumal Nallagownden, Irraivan Elamvazuthi, and Taib Ibrahim. A review on optimized control systems for building energy and comfort management of smart sustainable buildings. *Renewable and Sustainable Energy Reviews*, 34(0):409–429, 2014.
- [14] BACnet a Data Communication Protocol for Building Automation and Control Networks (Registered trademark of American Society of HVAC Engineers). URL <http://www.bacnet.org/>. Online, accessed 12-september-2016.
- [15] Modbus (Registered trademark of Schneider Electric). URL: <http://www.modbus.org/>. Online, accessed 12-september-2016.
- [16] LonMark International. URL <http://www.lonmark.org/>. Online, accessed 12-september-2016.
- [17] ZigBee ZigBee Alliance - a standards-based wireless technology. URL: <http://www.zigbee.org/>. Online, accessed 12-september-2016.
- [18] M. Ceriotti, L. Mottola, G. P. Picco, A. L. Murphy, S. Guna, M. Corra, M. Pozzi, D. Zonta, and P. Zanon. Monitoring heritage buildings with wireless sensor networks: The torre aquila deployment. In *IPSN 2009. International Conference on Information Processing in Sensor Networks*, pages 277–288, April 2009.
- [19] Mark Weiser. The computer for the 21st century. *SIGMOBILE Mob. Comput. Commun. Rev.*, 3(3):3–11, July 1999.
- [20] Diane J. Cook, Juan C. Augusto, and Vikramaditya R. Jakkula. Ambient intelligence: Technologies, applications, and opportunities. *Pervasive and Mobile Computing*, 5(4):277–298, 2009.
- [21] Asier Aztiria, Alberto Izaguirre, and Juan Carlos Augusto. Learning patterns in ambient intelligence environments: a survey. *Artif. Intell. Rev.*, 34(1):35–51, June 2010.

- 
- [22] Carlos Ramos, Juan Carlos Augusto, and Daniel Shapiro. Ambient intelligence - the next step for artificial intelligence. *Intelligent Systems, IEEE*, 23(2):15–18, march-april 2008.
- [23] Carlos Ramos. Ambient intelligence - a state of the art from artificial intelligence perspective. In José Neves, Manuel Filipe Santos, and José Manuel Machado, editors, *Progress in Artificial Intelligence*, volume 4874 of *Lecture Notes in Computer Science*, pages 285–295. Springer Berlin Heidelberg, 2007.
- [24] Fariba Sadri. Ambient intelligence: A survey. *ACM Comput. Surv.*, 43(4):36:1–36:66, October 2011.
- [25] Nik Bessis and Ciprian Dobre, editors. *Big Data and Internet of Things: A Roadmap for Smart Environments*. Springer International Publishing, 2014.
- [26] M. I. Jordan and T. M. Mitchell. Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245):255–260, 2015.
- [27] Richard S. Sutton. *Temporal credit assignment in reinforcement learning*. PhD thesis, University of Mass. COINS TR 84-2. Amherst, MA, 1984.
- [28] A.G. Barto and P. Anandan. Pattern-recognizing stochastic learning automata. *IEEE Transactions on Systems, Man and Cybernetics*, SMC-15(3):360–375, 1985.
- [29] David H. Ackley. *Advances in neural information processing systems 1*, chapter Associative learning via inhibitory search, pages 20–28. Morgan Kaufmann Publishers Inc, San Francisco, CA, USA, 1989.
- [30] R.B. Allen. Developing agent models with a neural reinforcement technique. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, volume 1, pages 206–207, November 1989.
- [31] Jonas Magnussen. Increased Energy Efficiency in Buildings using Model Predictive Control. Master’s thesis, NTNU Department of Engineering Cybernetics, Norway, June 2011.
- [32] U.S. Energy Information Administration. International Energy Outlook. Technical Report DOE/EIA-0484(2004), 2014.
- [33] IPCC, 2013. Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Technical report, [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)] Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 1535 pp., 2013.

## BIBLIOGRAPHY

---

- [34] Direcção Geral de Energia e Geologia (DGEC). Balanço energético. Technical report, Ministério da Economia e do Emprego, Portugal, 2013.
- [35] Asif Syed. *Advanced Building Technologies for Sustainability*. John Wiley & Sons, Inc, Hoboken, New Jersey, 2012.
- [36] Abdeen Mustafa Omer. Energy, environment and sustainable development. *Renewable and sustainable energy reviews*, 12:2265–2300, December 2008.
- [37] Luis Pérez-Lombard, José Ortiz, and Christine Pout. A review on buildings energy consumption information. *Energy and Buildings*, 40(3):394–398, 2008.
- [38] International Energy Agency. Worldwide trends in energy use and efficiency - key insights from IEA indicator analysis. Technical report, IEA, 2008. URL [https://www.iea.org/publications/freepublications/publication/Indicators\\_2008.pdf](https://www.iea.org/publications/freepublications/publication/Indicators_2008.pdf). Online, accessed 12-september-2016.
- [39] Instituto Nacional de Estatística and DGEG. Instituto inquérito ao consumo de energia no sector doméstico. Technical report, Direcção Geral de Energia e Geologia (DGEC), 2010.
- [40] ADENE. Guia da Eficiência Energética. Technical report, Agência para a Energia, Portugal, 2012.
- [41] Market Study. Appliance Energy Efficiency Opportunities: China 2013. Technical report, CLASP, 2013.
- [42] Contribution of Working Group III to the Fourth Assessment Report of the Inter-governmental Panel on Climate Change. Climate change 2007: Mitigation. Technical report, [B. Metz, O.R. Davidson, P.R. Bosch, R. Dave, L.A. Meyer (eds)], Cambridge University Press, United Kingdom and New York, NY, USA, 2007.
- [43] Arthur H. Rosenfeld, Joseph J. Romm, Hashem Akbari, and Alan C. Lloyd. Painting the town white – and green. URL <http://heatisland.lbl.gov/PUBS/PAINTING/>. Online, accessed 12-september-2016.
- [44] M.H. M. Levine. Energy efficiency improvement utilizing high technology: An assessment of energy use in industry and buildings; report and case studies. *London SW1A 1HD, United Kingdom: Technical report, World Energy Council, 34 St James's Street*, 1995.
- [45] Greg Kats and Capital E. The costs and financial benefits of green buildings. Technical report, California's Sustainable Building Task Force, October 2003. URL [http://evanmills.lbl.gov/pubs/pdf/green\\_buildings.pdf](http://evanmills.lbl.gov/pubs/pdf/green_buildings.pdf). Online, accessed 12-september-2016.

- 
- [46] K.M. Fowler and E.M. Rauch. Sustainable building rating systems summary. Technical report, Pacific Northwest National Laboratory, July 2006. URL <http://www.usgbc.org/Docs/Archive/General/Docs1915.pdf>. Online, accessed 12-september-2016.
- [47] Michael Boduch and Warren Fincher. Standards of human comfort. Technical report, The University of Texas at Austin, School of Architecture, 2009.
- [48] H.A Ajimotokan, L.A Oloyede, and M.E. Ismail. Influence of indoor environment on health and productivity. *New York Science Journal*, 2009.
- [49] Pawel Wargocki, David P. Wyon, Jan Sundell, Geo Clausen, and P. Ole Fanger. The effects of outdoor air supply rate in an office on perceived air quality, sick building syndrome (sbs) symptoms and productivity. *Indoor Air*, 4(10):222–36, 2000.
- [50] William J. Fisk. *Estimates of Potential Nationwide Productivity And Health Benefits From Better Indoor Environments: An Update.* J.M. Samet, and J. F. McCarthy J. D. Spengler. *Indoor Air Quality Handbook.*, chapter 4. McGraw Hill, 1999.
- [51] Susanne Bodach. Developing bioclimatic zones and passive solar design strategies for nepal. In *30th International PLEA Conference*. CEPT University, Ahmedabad, December 2014.
- [52] Ji Eun Kang, Ki Uhn Ahn, Cheol Soo Park, and Thorsten Schuetze. Assessment of passive vs. active strategies for a school building design. *Sustainability*, 7: 15136–15151, November 2015.
- [53] Vladimir Mikler, Albert Bicol, Beth Breisnes, and Michael Labrie. Passive design toolkit. Technical report, 2009. URL <http://vancouver.ca/files/cov/passive-design-large-buildings.pdf>. Online, accessed 12-september-2016.
- [54] B. Sun, P. B. Luh, Q. S. Jia, Z. Jiang, F. Wang, and C. Song. Building Energy Management: Integrated Control of Active and Passive Heating, Cooling, Lighting, Shading, and Ventilation Systems. *IEEE Transactions on Automation Science and Engineering*, 10(3):588–602, July 2013.
- [55] Rongpeng Zhang. *Dyanmic Optimization of Integrated Active-Passive Strategies for Building Enthalpy Control*. PhD thesis, Carnegie Mellon University, May 2014.
- [56] Jianchao Zhang, Boon-Chong Seet, and Tek Tjing Lie. Building information modelling for smart built environments. *Buildings*, 5(1), 2015.

## BIBLIOGRAPHY

---

- [57] The Climate Group. SMART2020: Enabling the low carbon economy in the information age. Technical report, Global e-Sustainability Initiative (GeSI), 2008. URL <http://gesi.org/article/43>. Online, accessed 12-september-2016.
- [58] The Information Society Technologies Advisory Group (ISTAG). Scenarios for Ambient Intelligence in 2010. Technical report, European Commission Report, 2001. URL <https://cordis.europa.eu/pub/ist/docs/istagscenarios2010.pdf>. Online, accessed 12-september-2016.
- [59] The Information Society Technologies Advisory Group (ISTAG). Strategic orientations and priorities for IST in FP6. Technical report, European Commission Report, 2002. URL [https://cordis.europa.eu/pub/ist/docs/istag\\_kk4402456encfull.pdf](https://cordis.europa.eu/pub/ist/docs/istag_kk4402456encfull.pdf). Online, accessed 12-september-2016.
- [60] The Information Society Technologies Advisory Group (ISTAG). Ambient Intelligence: from vision to reality. Technical report, European Commission Report, 2003. URL [https://cordis.europa.eu/pub/ist/docs/istag-ist2003\\_consolidated\\_report.pdf](https://cordis.europa.eu/pub/ist/docs/istag-ist2003_consolidated_report.pdf). Online, accessed 12-september-2016.
- [61] M.A. Piette, J. Granderson, M. Wetter, and S. Kiliccote. Intelligent building energy information and control systems for low-energy operations and optimal demand response. *Design Test of Computers, IEEE*, 29(4):8–16, August 2012.
- [62] Johnny Wong, Heng Li, and Jenkin Lai. Evaluating the system intelligence of the intelligent building systems part1: Development of key intelligent indicators and conceptual analytical framework. *Automation in Construction*, 17:284–302, 2008.
- [63] Johnny Wong, Heng Li, and Jenkin Lai. Evaluating the system intelligence of the intelligent building systems part2: Construction and validation of analytical models. *Automation in Construction*, 17:303–321, 2008.
- [64] Clements-Croome T.D.J. What do we mean by intelligent buildings? *Automation in Construction*, 6(5):395–400, 1997.
- [65] L. Travé-Massuyès. Bridging control and artificial intelligence theories for diagnosis: A survey. *Engineering Applications of Artificial Intelligence*, 27:1–16, 2014.
- [66] Donald Frey and Vernon Smith. Advanced Automated HVAC Fault Detection and Diagnostics Commercialization Program. Technical report, Architectural Energy Corporation, prepared for: California Energy Commission, December 2008. URL <http://www.energy.ca.gov/2013publications/CEC-500-2013-054/CEC-500-2013-054.pdf>. Online, accessed 12-september-2016.

- 
- [67] M.W. Ellis and E.H. Mathews. Needs and trends in building and HVAC system design tools. *Building and Environment*, 37(5):461–470, 2002.
- [68] Anastasios I. Dounis. Artificial intelligence for energy conservation in buildings. *Advances in Building Energy Research*, 4(1):267–299, 2010.
- [69] Robert C. Ward, Jim C. Loftis, and Graham B. McBride. The “data-rich but information-poor” syndrome in water quality monitoring. *Environmental Management*, 10(3):291–297, 1986.
- [70] Tomasz Kosicki and Trygve Thomessen. Cognitive Human-Machine Interface Applied in Remote Support for Industrial Robot Systems. *International Journal of Advanced Robotic Systems*, 10(342):1–11, 2013.
- [71] Paul W. Quimby, Ritesh Khire, Francesco Leonardi, and Soumik Sarkar. A Novel Human Machine Interface for Advanced Building Controls and Diagnostics. In *Proceedings of the 3rd International High Performance Buildings Conference*. Purdue, July 2014.
- [72] Patrick Brézillon. Context in problem solving: a survey. *Knowl. Eng. Rev.*, 14(1):47–80, May 1999.
- [73] Richmond H Thomason. Representing and reasoning with context. In *Proceedings of the International Conference on Artificial Intelligence and Symbolic Computation AISC98*, pages 29–41. Springer-Verlag, Plattsburgh, New York, 1998.
- [74] John McCarthy and Sasa Buvac. Formalizing context (expanded notes). Technical report, Computer Science Department, Stanford University, Stanford, CA, USA, 1994. URL <http://www-formal.stanford.edu/jmc/mccarthy-buvac-98/>. Online, accessed 12-september-2016.
- [75] Avelino J. Gonzalez, Brian S. Stensrud, and Gilbert Barrett. Formalizing context-based reasoning: A modeling paradigm for representing tactical human behavior. *Int. J. Intell. Syst.*, 23(7):822–847, July 2008.
- [76] Marius Mikalsen and Anders Kofod-Petersen. Representing and reasoning about context in a mobile environment. *Revue D’Intelligence Artificielle RIA*, 19:479–498, 2005.
- [77] Aveline J. Gonzalez Gary Stein. Learning in context: enhancing machine learning with context-based reasoning. *Journal of Applied Intelligence*, 41(3):709–724, 2014.
- [78] Duen-Ren Liu; Chih-Kun Ke; Mei-Yu Wu. Context-based knowledge support for problem-solving by rule-inference and case-based reasoning. In *International Conference on Machine Learning and Cybernetics*, volume 6, pages 3205–3210, 12-15 July 2008.

## BIBLIOGRAPHY

---

- [79] Jong yi Hong, Eui ho Suh, and Sung-Jin Kim. Context-aware systems: A literature review and classification. *Expert Systems with Applications*, 36(4): 8509–8522, 2009.
- [80] Xin Li, Martina Eckert, José-Fernán Martínez, and Gregorio Rubio. Context aware middleware architectures: Survey and challenges. *Sensors*, 15(8):20570–20607, 2015.
- [81] Brian S. Stensrud, Gilbert C. Barrett, Viet C. Trinh, and Avelino J. Gonzalez. Context-Based Reasoning: A Revised Specification. In *Proceedings of the Seventeenth International Florida Artificial Intelligence Research Society Conference*. Miami Beach, Florida, 2004.
- [82] Patrick Brézillon and Suhayya Abu-Hakima. Using knowledge in its context: Report on the IJCAI-93 workshop, 1995.
- [83] Vanessa Tavares Nunes, Flávia Maria Santoro, and Marcos R.S. Borges. A context-based model for knowledge management embodied in work processes. *Information Sciences*, 179(15):2538–2554, 2009.
- [84] Karen Henriksen and Jadwiga Indulska. Developing context-aware pervasive computing applications: Models and approach. *Pervasive Mob. Comput.*, 2(1): 37–64, February 2006.
- [85] Antonis Bikakis, Theodore Patkos, Grigoris Antoniou, and Dimitris Plexousakis. *A Survey of Semantics-Based Approaches for Context Reasoning in Ambient Intelligence*, pages 14–23. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [86] Grigoris Antoniou, Constantinos Papatheodorou, and Antonis Bikakis. Reasoning about context in ambient intelligence environments: A report from the field. In *Proceedings of the Twelfth International Conference on the Principles of Knowledge Representation and Reasoning (KR 2010)*, pages 557–559, 2010.
- [87] I. Khajenasiri, J. Virgone, and G. Gielen. A presence-based control strategy solution for HVAC systems. In *2015 IEEE International Conference on Consumer Electronics (ICCE)*, pages 620–622, January 2015.
- [88] Hossein Mirinejad, Seyed Hossein Sadati, Maryam Ghasemian, and Hamid Torab. Control Techniques in Heating, Ventilation and Air Conditioning (HVAC) Systems. *Computer Science*, 4(9):777–783, 2008.
- [89] Anh-Tuan Nguyen, Sigrid Reiter, and Philippe Rigo. A review on simulation-based optimization methods applied to building performance analysis. *Applied Energy*, 113:1043–1058, 2014.
- [90] Ernest Davis and Gary Marcus. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Commun. ACM*, 58(9):92–103, August 2015.

- 
- [91] Giuseppe Loseto. *Knowledge Representation Methods for Smart Devices in Intelligent Buildings*. PhD thesis, Politecnico di Bari, 2009.
- [92] Mario J. Kofler, Christian Reinisch, and Wolfgang Kastner. An Intelligent Knowledge Representation of Smart Home Energy Parameters. In *World Renewable Energy Congress*, pages 921–928. Linköping, Sweden, May 2011.
- [93] Thanos G. Stavropoulos, Dimitris Vrakas, Danai Vlachava, and Nick Bassiliades. Bonsai: A smart building ontology for ambient intelligence. In *Proceedings of the 2Nd International Conference on Web Intelligence, Mining and Semantics*, WIMS '12, pages 30:1–30:12, 2012.
- [94] L. James Lo and Atila Novoselac. Localized air-conditioning with occupancy control in an open office. *Energy and Buildings*, 42(7):1120–1128, 2010.
- [95] Luis Perez-Lombard, Jose Ortiz, and Ismael R. Maestre. The map of energy flow in HVAC systems. *Applied Energy*, 88(12):5020–5031, 2011.
- [96] Christy Love. Operating room HVAC setback strategies. Technical report, The American Society for Healthcare Engineering (ASHE), 2011. URL [http://www.ashe.org/management\\_monographs/pdfs/mg2011love.pdf](http://www.ashe.org/management_monographs/pdfs/mg2011love.pdf). Online, accessed 12-september-2016.
- [97] Vahid Vakiloroyaya, Bijan Samali, Ahmad Fakhar, and Kambiz Pishghadam. A review of different strategies for HVAC energy saving. *Energy Conversion and Management*, 77:738–754, 2014.
- [98] Virginia Smith, Tamim Sookoor, and Kamin Whitehouse. Modeling Building Thermal Response to HVAC Zoning. *SIGBED Rev.*, 9(3):39–45, July 2012.
- [99] Jian Liang and Ruxu Du. Thermal comfort control based on neural network for HVAC application. In *Proceedings of 2005 IEEE Conference on Control Applications (CCA 2005)*, pages 819–824, August 2005.
- [100] Refrigerating American Society of Heating and Air Conditioning Engineers. *Thermal environmental conditions for human occupancy*. ASHRAE Standard 55, 2010.
- [101] Maarten Sourbron, Clara Verhelst, and Lieve Helsen. Building models for model predictive control of office buildings with concrete core activation. *Journal of Building Performance Simulation*, 6(3):175–198, 2013.
- [102] J.D. Álvarez, J.L. Redondo, E. Camponogara, J. Normey-Rico, M. Berenguel, and P.M. Ortigosa. Optimizing building comfort temperature regulation via model predictive control. *Energy and Buildings*, 57:361–372, 2013.

## BIBLIOGRAPHY

---

- [103] J.A. Candanedo and A.K. Athienitis. Simplified linear models for predictive control of advanced solar homes with passive and active thermal storage. In *International High Performance Buildings Conference*. Pardue, July 2010.
- [104] William O'Brien, Andreas Athienitis, and Ted Kesik. Thermal zoning and interzonal airflow in the design and simulation of solar houses: a sensitivity analysis. *Journal of Building Performance Simulation*, 4(3):239–256, 2011.
- [105] A.Rabl. Parameter estimation in buildings: methods for dynamic analysis of measured energy use. *Journal of Solar Energy Engineering*, 110:52–66, 1988.
- [106] R. Balan, S. Stan, and C. Lapusan. A model based predictive control algorithm for building temperature control. In *3rd IEEE International Conference on Digital Ecosystems and Technologies, DEST '09*, pages 540–545, 2009.
- [107] Radu Bălan, Joshua Cooper, Kuo-Ming Chao, Sergiu Stan, and Radu Donca. Parameter identification and model based predictive control of temperature inside a house. *Energy and Buildings*, 43(2–3):748–758, 2011.
- [108] Yashen Lin, Timothy Middelkoop, and Prabir Barooah. Issues in identification of control-oriented thermal models of zones in multi-zone buildings. In *IEEE 51st Annual Conference on Decision and Control (CDC)*, pages 6932–6937, December 2012.
- [109] Raad Z. Homod. Review on the HVAC System Modeling Types and the Shortcomings of Their Application. *Journal of Energy*, 2013.
- [110] Qi Luo and K.B. Ariyur. Building thermal network model and application to temperature regulation. In *2010 IEEE International Conference on Control Applications (CCA)*, pages 2190–2195, September 2010.
- [111] Siddharth Goyal and Prabir Barooah. A method for model-reduction of nonlinear thermal dynamics of multi-zone buildings. *Energy and Buildings*, (47): 332–340, 2012.
- [112] S. Goyal and P. Barooah. A method for model-reduction of nonlinear building thermal dynamics. In *American Control Conference (ACC)*, pages 2077–2082, June 2011.
- [113] Y. Gao, J.J. Roux, L.H. Zhao, and Y.Jiang. Dynamical building simulation: a low order model for thermal bridges losses. *Energy and Buildings*, (40):2236–2243, 2008.
- [114] Kun Deng, Prabir Barooah, Prashant G. Mehta, and Sean P. Meyn. Building thermal model reduction via aggregation of states. In *American Control Conference (ACC), 2010*, pages 5118–5123. Marriott Waterfront, Baltimore, MD, USA, June 2010.

- 
- [115] M.M. Gouda, S. Danaher, and C.P. Underwood. Building thermal model reduction using nonlinear constrained optimization. *Building and Environment*, 37(12):1255–1265, 2002.
- [116] S.A. Marshall. *An approximate method for reducing the order of a linear system*, pages 642–643. *Control*, 10, 1966.
- [117] Savo D. Dukić and Andrija T. Sarić. Dynamic model reduction: An overview of available techniques with application to power systems. *Serbian Journal of Electrical Engineering*, 9(2):131–169, 2012.
- [118] Tam Van Nguyen, Student Member, and Deokjai Choi. Context Reasoning Using Contextual Graph. In *2008 IEEE 8th International Conference on Computer and Information Technology Workshops*, pages 488–493. IEEE, July 2008.
- [119] Gregory D. Abowd, Anind K. Dey, Peter J. Brown, Nigel Davies, Mark Smith, and Pete Steggles. Towards a better understanding of context and context-awareness. In *Proceedings of the 1st international symposium on Handheld and Ubiquitous Computing*, HUC '99, pages 304–307. Springer-Verlag, London, UK, UK, 1999.
- [120] Klaus Kaae Andersen, Henrik Madsen, and Lars H. Hansen. Modelling the heat dynamics of a building using stochastic differential equations. *Energy and Buildings*, 31(1):13–24, 2000.
- [121] Niels Rode Kristensen, Henrik Madsen, and Sten Bay Jørgensen. Parameter estimation in stochastic grey-box models. *Automatica*, 40(2):225–237, 2004.
- [122] Anthony Fontanini, Umesh Vaidya, and Baskar Ganapathysubramanian. A stochastic approach to modeling the dynamics of natural ventilation systems. *Energy and Buildings*, 63(0):87–97, 2013.
- [123] H. Madsen and J. Holst. Estimation of continuous-time models for the heat dynamics of a building. *Energy and Buildings*, 22(1):67–79, 1995.
- [124] Anders Thavlov and Henrik W. Bindner. Thermal models for intelligent heating of buildings. In *Proceedings of the International Conference on Applied Energy (ICAE 2012)*. Suzhou, China, July 2012.
- [125] M. Castilla, J.D. Álvarez, M. Berenguel, F. Rodríguez, J.L. Guzmán, and M. Pérez. A comparison of thermal comfort predictive control strategies. *Energy and Buildings*, 43(10):2737–2746, 2011.
- [126] Samuel Prívará, Zdeněk Váňa, Eva Žáčková, and Jiří Cigler. Building modeling: Selection of the most appropriate model for predictive control. *Energy and Buildings*, 55(0):341–350, 2012.

## BIBLIOGRAPHY

---

- [127] Peder Bacher and Henrik Madsen. Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings*, 43(7):1511–1522, 2011.
- [128] Peder Bacher and Henrik Madsen. *Procedure for identifying models for the heat dynamics of buildings*. IMM-Technical report-2010-04. DTU Informatics, Technical University of Denmark, March 2010.
- [129] P.J.L. Cuijpers, M.A. Reniers, and W.P.M.H. Heemels. Hybrid transition systems. Technical report, CSR 02-12, TU/e, Eindhoven, The Netherlands, 2002. URL <http://alexandria.tue.nl/extra1/wskrap/publichtml/200212.pdf>. Online, accessed 12-september-2016.
- [130] Panos J. Antsaklis and Xenofon D. Koutsoukos. Hybrid systems control. Technical Report isis-2001-003, Dep. of Electrical Engineering, University of Notre Dame, IN 46556, 2001. URL <http://www3.nd.edu/~isis/techreports/isis-2001-003.pdf>. Online, accessed 12-september-2016.
- [131] T. Murata. Petri Nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–580, April 1989.
- [132] W. Goetzler, M. Guernsey, and J. Young. Research & development roadmap for emerging HVAC technologies. Technical report, U.S. Department of Energy, October, 2014.
- [133] Muhammad Waseem Ahmad, Monjur Mourshed, Baris Yuce, and Yacine Rezgui. Computational intelligence techniques for HVAC systems: A review. *Building Simulation*, 9(4):359–398, 2016.
- [134] G. Bovet and J. Hennebert. A distributed web-based naming system for smart buildings. In *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2014 IEEE 15th International Symposium on a*, pages 1–6, June 2014.
- [135] D.D. Šiljak. Decentralized control and computations: status and prospects. *Annual Reviews in Control*, 20(0):131–141, 1996.
- [136] E. Camponogara, D. Jia, B.H. Krogh, and S. Talukdar. Distributed model predictive control. *Control Systems, IEEE*, 22(1):44–52, feb 2002.
- [137] Dionysia Kolokotsa. The role of smart grids in the building sector. *Energy and Buildings*, 116:703–708, 2016.
- [138] J. Pan, R. Jain, S. Paul, T. Vu, A. Saifullah, and M. Sha. An internet of things framework for smart energy in buildings: Designs, prototype, and experiments. *IEEE Internet of Things Journal*, 2(6):527–537, Dec 2015.
- [139] Aqeel H. Kazmi, Michael J. O’grady, Declan T. Delaney, Antonio G. Ruzzelli, and Gregory M. P. O’hare. A review of wireless-sensor-network-enabled building

- energy management systems. *ACM Trans. Sen. Netw.*, 10(4):66:1–66:43, June 2014.
- [140] Fatima Amara, Kodjo Agbossou, Alben Cardenas, Yves Dubé, and Sousso Kelouwani. Comparison and Simulation of Building Thermal Models for Effective Energy Management. *Smart Grid and Renewable Energy*, 6:95–112, 2015.
- [141] Wandi Liu, Hai Wang, Hengyang Zhao, Shujuan Wang, Haibao Chen, Yuzhuo Fu, Jian Ma, Xin Li, and S. X. D. Tan. Thermal modeling for energy-efficient smart building with advanced overfitting mitigation technique. In *2016 21st Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 417–422, January 2016.
- [142] Alessandro Abate, Martin Fränzle, Ian Hiskens, and Martin Strelec. Modeling, Verification, and Control of Complex Systems for Energy Networks (Dagstuhl Seminar 14441). *Dagstuhl Reports*, 4(10):69–97, 2015.
- [143] Andrew Windham and Stephen Treado. A review of multi-agent systems concepts and research related to building hvac control. *Science and Technology for the Built Environment*, 22(1):50–66, 2016.
- [144] R. Kwadzogah, M. Zhou, and S. Li. Model predictive control for hvac systems – a review. In *2013 IEEE International Conference on Automation Science and Engineering (CASE)*, pages 442–447, August 2013.
- [145] V. Vakiloroyaya, S.W. SU, and Q.P. Ha. HVAC Integrated Control for Energy Saving and Comfort Enhancement. In *28th International Association for Automation and Robotics in Construction (ISARC)*, pages 245–250, 2011.
- [146] Xinqiao Jin, Zhimin Du, and Xiaokun Xiao. Energy evaluation of optimal control strategies for central VWV chiller systems. *Applied Thermal Engineering*, 27(5–6):934–941, 2007.
- [147] A.I. Dounis and C. Caraiscos. Advanced control systems engineering for energy and comfort management in a building environment - a review. *Renewable and Sustainable Energy Reviews*, 13(6–7):1246–1261, 2009.
- [148] D.Subbaram Naidu and Craig G. Rieger. Advanced control strategies for HVAC&R systems - an overview: Part 2: Soft and fusion control. *HVAC&R Research*, 17(2):144–158, 2011.
- [149] Pervez Hameed Shaikh, Nursyarizal Bin Mohd Nor, Perumal Nallagownden, and Irraivan Elamvazuthi. Intelligent Optimized Control System for Energy and Comfort Management in Efficient and Sustainable Buildings. In *Proceedings of the 4th International Conference on Electrical Engineering and Informatics (ICEEI 2013)*, volume 11, pages 99–106, 2013.

- [150] Rafael Alcalá, Jose M. Benítez, Jorge Casillas, Oscar Cordón, and Raúl Pérez. Fuzzy control of HVAC systems optimized by genetic algorithms. *Applied Intelligence*, 18(2):155–177, March 2003.
- [151] Rafael Alcalá, Jorge Casillas, Oscar Cordón, Antonio González, and Francisco Herrera. A genetic rule weighting and selection process for fuzzy control of heating, ventilating and air conditioning systems. *Engineering Applications of Artificial Intelligence*, 18(3):279–296, 2005.
- [152] Velimir Congradac, Bosko Milosavljevic, Jovan Velickovic, and Bogdan Prebiracevic. Control of the lighting system using a genetic algorithm. *Thermal Science*, 16(suppl. 1):237–250, 2012.
- [153] K.F. Fong, V.I. Hanby, and T.T. Chow. System optimization for HVAC energy management using the robust evolutionary algorithm. *Applied Thermal Engineering*, 29(11-12):2327–2334, 2009.
- [154] Jason Teeter and Mo-Yuen Chow. Application of functional link neural network to HVAC thermal dynamic system identification. *IEEE Transactions on Industrial Electronics*, 45(1):170–176, February 1998.
- [155] C.A. Hernandez S, R. Romero, and D. Giral. Optimization of the use of residential lighting with neural network. In *(CISE), 2010 International Conference on Computational Intelligence and Software Engineering*, pages 1–5, December 2010.
- [156] E. Sierra, A. Hossian, D. Rodríguez, M. García-Martínez, P. Britos, and R. García-Martínez. Optimizing building’s environments performance using intelligent systems. In *Proceedings of the 21st international conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems: New Frontiers in Applied Artificial Intelligence*, IEA/AIE ’08, pages 486–491. Springer-Verlag, Berlin, Heidelberg, 2008.
- [157] Y.K. Penya, C.E. Borges, D. Agote, and I. Fernandez. Short-term load forecasting in air-conditioned non-residential buildings. In *(ISIE), 2011 IEEE International Symposium on Industrial Electronics*, pages 1359–1364, June 2011.
- [158] Plamen Angelov. A fuzzy approach to building thermal systems optimization. In *8th International Fuzzy Systems Association World Congress*. Taipei, Taiwan, 1999.
- [159] H.-B. Kuntze and Th. Bernard. A new fuzzy-based supervisory control concept for the demand-responsive optimization of HVAC control systems. In *Proceedings of the 37th IEEE Conference on Decision and Control*, volume 4, pages 4258–4263, December 1998.

- 
- [160] Servet Soyguder and Hasan Alli. An expert system for the humidity and temperature control in HVAC systems using ANFIS and optimization with fuzzy modeling approach. *Energy and Buildings*, 41(8):814–822, 2009.
- [161] Vipul Singhvi, Andreas Krause, Carlos Guestrin, James H. Garrett, Jr., and H. Scott Matthews. Intelligent light control using sensor networks. In *Proceedings of the 3rd international conference on Embedded networked sensor systems*, SenSys '05, pages 218–229. ACM, New York, NY, USA, 2005.
- [162] E. Sierra, A. Hossian, D. Rodríguez, M. García-Martínez, P. Britos, and R. García-Martínez. Fuzzy control for improving energy management within indoor building environments. In *Electronics, Robotics and Automotive Mechanics Conference. CERMA*, pages 412–416, September 2007.
- [163] Pedro Albertos and Antonio Sala. Fuzzy logic controllers. advantages and drawbacks. In *XIII Congreso de la Asociación Chilena de Control Automático*, volume 3, pages 833–844, September 1998.
- [164] Jelena Godjevac. Comparative study of fuzzy control, neural network control and neuro-fuzzy control. Technical report, École Polytechnique Fédérale de Lausanne, February 1995.
- [165] Leemon C. Baird, Mance E. Harmon, and A. Harry Klopf. Reinforcement learning: An alternative approach to machine intelligence. Technical report, Wright Laboratory, February 1996. URL <http://leemon.com/papers/1996bhk.pdf>. Online, accessed 12-september-2016.
- [166] Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: a survey. *Journal of Artificial Intelligence Research*, 4(1): 237–285, 1996.
- [167] Nest. Enhanced auto-schedule. White paper, Nest Labs, Inc., 2014. URL <https://nest.com/downloads/press/documents/enhanced-auto-schedule-white-paper.pdf>. Online, accessed 12-september-2016.
- [168] Varick L. Erickson and Alberto E. Cerpa. Occupancy Based Demand Response HVAC Control Strategy. In *Proceedings of the 2Nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, BuildSys '10, pages 7–12. ACM, New York, NY, USA, 2010.
- [169] M. Kintner-Meyer and A. F. Emery. Optimal control of an HVAC system using cold storage and building thermal capacitance. *Energy and Buildings*, 23(1): 19–31, 1995.
- [170] Konstantinos Dalamagkidis and Dionysia Kolokotsa. Reinforcement learning for building environmental control. In C. Weber, M. Elshaw, and N. M. Mayer,

## BIBLIOGRAPHY

---

- editors, *Reinforcement Learning - Theory and Applications*, chapter 15, pages 283–294. I-Tech Publications, 2008.
- [171] International Organization for Standardization. *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*. EVS-EN ISO 7730:2006.
- [172] Povl Ole Fanger. *Thermal Comfort: Analysis and applications in environmental engineering*. McGraw-Hill, New York, 1972.
- [173] Hal Levin. Designing for people: What do building occupants really want? In *Healthy Buildings*, 2003.
- [174] J. Van Hoof. Forty years of fanger’s model of thermal comfort: comfort for all? *Indoor Air*, 18(3):182–201, 2008.
- [175] Lisje Schellen, Marcel Loomans, Wouter van Marken Lichtenbelt, Arjan Frijns, and Martin de Wit. Assessment of thermal comfort in relation to applied low exergy systems. In *Adapting to Change: New Thinking on Comfort*, April 2010.
- [176] Pablo Aparicio Ruiz, José Guadix Martín, and Luis Onieva Jesús M. Sanz. New model for the search for comfort through surveys. In *WSEAS Transactions on Circuits and systems*, volume 11, pages 125–135, 2012.
- [177] J.F. Nicol and M.A. Humphreys. Adaptive thermal comfort and sustainable thermal standards for buildings. *Energy and Buildings*, 34(6):563–572, 2002.
- [178] M A Humphreys. *Standards for Thermal Comfort*, chapter Thermal comfort temperatures and the habits of Hobbits, pages 3–13. Chapman & Hall, 1995.
- [179] E.H Mathews, D.C Arndt, C.B Piani, and E van Heerden. Developing cost efficient control strategies to ensure optimal energy use and sufficient indoor comfort. *Applied Energy*, 66(2):135–159, 2000.
- [180] TJ Williamson and P. Riordan. Thermostat strategies for discretionary heating and cooling of dwellings in temperate climates. In *5th IBPSA Building Simulation Conference*, pages 1–8, 1997.
- [181] U. Rutishauser, J. Joller, and R. Douglas. Control and learning of ambience by an intelligent building. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 35(1):121–132, January 2005.
- [182] Therese Peffer, Marco Pritoni, Alan Meier, Cecilia Aragon, and Daniel Perry. How people use thermostats in homes: A review. *Building and Environment*, 46(12):2529–2541, December 2011.

- [183] Alan Meier. Thermostat interface and usability: A survey. Technical report, Lawrence Berkeley National Laboratory, 2011. LBNL Paper LBNL-4182E.
- [184] Robert Johnson. Smart Metering: A review of Smart Metering and Survey options for Energy. Technical report, University of East Anglia – School of Environmental Sciences (LCIC), 2010.
- [185] Shengwei Wang and Xinhua Xu. Optimal and robust control of outdoor ventilation airflow rate for improving energy efficiency and IAQ. *Building and Environment*, 39(7):763–773, 2004.
- [186] Steven J. Emmerich and Andrew K. Persily. State-of-the-art review of CO<sub>2</sub> demand controlled ventilation technology and application. Technical report, National Institute of Standards and Technology, California Energy Commission, 2003.
- [187] P.J. Boait and R.M. Rylatt. A method for fully automatic operation of domestic heating. *Energy and Buildings*, 42(1):11–16, 2010. International Conference on Building Energy and Environment (COBEE 2008).
- [188] Manu Gupta, Stephen S. Intille, and Kent Larson. Adding gps-control to traditional thermostats: An exploration of potential energy savings and design challenges. In *Proceedings of the 7th International Conference on Pervasive Computing*, Pervasive '09, pages 95–114. Springer-Verlag, Berlin, Heidelberg, 2009.
- [189] MarketsandMarkets. Smart Thermostat Market by Component by Network Connectivity (Wired, Wireless), by Application (Residential, Office Building, Industrial Building, Educational Institutional, Retail, Hospitality, and Healthcare), and Geography – Forecast to 2020. Market report, September 2014.
- [190] Deborah Weinswig. The Connected Home Series 4. Energy Management. Market report, Fung Business Intelligence Centre, Gloabal Retail & Technology, January 2016.
- [191] G. Muthuselvi and Saravanan B. Real time speech recognition based building automation system. *SRPN Journal of Engineering and Applied Sciences*, 9(12): 2831–2839, December 2014.
- [192] Aaron Sloman. Review of: Affective computing. URL <http://www.cs.bham.ac.uk/research/projects/cogaff/Sloman.picard.review.pdf>. Online, accessed 12-september-2016.
- [193] Rosalind W. Picard. Affective computing for hci. In *Proceedings of HCI International (the 8th International Conference on Human-Computer Interaction) on Human-Computer Interaction: Ergonomics and User Interfaces-Volume I - Volume I*, pages 829–833. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 1999.

## BIBLIOGRAPHY

---

- [194] Norbert A. Streitz, Carsten Rocker, Thorsten Prante, Daniel van Alphen, Richard Stenzel, and Carsten Magerkurth. Designing smart artifacts for smart environments. *Computer*, 38(3):41–49, March 2005.
- [195] Anca-Diana Barbu, Nigel Griffiths, and Gareth Morton. Achieving energy efficiency through behaviour change: what does it take? Technical report, European Environment Agency (EEA) and Ricardo-AEA, 2013.
- [196] Andrea H. McMakin, Elizabeth L. Malone, and Regina E. Lundgren. Motivating residents to conserve energy without financial incentive. *Environment and Behavior Journal*, 2002.
- [197] S. Darby. The effectiveness of feedback on energy consumption. Technical report, A review for DEFRA of the literature on metering, billing, and direct displays., 2006. URL <http://www.eci.ox.ac.uk/research/energy/downloads/smart-metering-report.pdf>. Online, accessed 12-september-2016.
- [198] Eun-Ju Lee, Min-Ho Pae, Dong-Ho Kim, Jae-Min Kim, and Jong-Yeob Kim. Literature review of technologies and energy feedback measures impacting on the reduction of building energy consumption. In Seung-Deog Yoo, editor, *EKC2008 Proceedings of the EU-Korea Conference on Science and Technology*, volume 124 of *Springer Proceedings in Physics*, pages 223–228. Springer Berlin Heidelberg, 2008.
- [199] Bill Wright and Begum Nash. A study of the effects of feedback on domestic energy use (part one and two). Technical report, Sustainable Homes, Department of Energy & Climate Change, 2014.
- [200] Sergiy Bogomolov, Daniele Magazzeni, Andreas Podelski, and Martin Wehrle. Planning as model checking in hybrid domains. In *Proceedings of twenty-eighth AAAI Conference on Artificial Intelligence and the twenty-sixth Innovative Applications of Artificial Intelligence (AAAI 2014)*, volume 3, pages 2228–2234. Québec, Canada, July 2014.
- [201] Sriram Sankaranarayanan, Thao Dang, and Franjo Ivančić. Symbolic model checking of hybrid systems using template polyhedra. In C. R. Ramakrishnan and Jakob Rehof, editors, *Proceedings of the 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2008)*, pages 188–202. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [202] John Lygeros. Lecture notes on hybrid systems. Technical report, Department of Electrical and Computer Engineering, University of Patras, Patras, GR-2650, Greece, February 2004.
- [203] Thomas A. Henzinger, Pei-Hsin Ho, and Howard Wong-Toi. Hytech: A model checker for hybrid systems. In Orna Grumberg, editor, *Proceedings of the*

- 
- 9th International Conference on Computer Aided Verification*, pages 460–463. Springer Berlin Heidelberg, Berlin, Heidelberg, 1997.
- [204] Mark G. Core, H. Chad Lane, Michael Van Lent, Dave Gomboc, and Steve Solomon. Building explainable artificial intelligence systems. In *In Proceedings of the 18th Conference on Innovative Applications of Artificial Intelligence (IAAI-06)*, 2006.
- [205] Baader Franz, Horrocks Ian, and Sattler Ulrike. *Handbook of Knowledge Representation*, volume 6525, chapter 3, pages 135–179. Elsevier B.V., 2008.
- [206] Werner Nutt Franz Baader. Basic description logics. Online. URL <http://www.inf.unibz.it/~franconi/dl/course/dlhb/dlhb-02.pdf>. Online, accessed 12-september-2016.
- [207] Daniele Nardi and Ronald J. Brachman. An introduction to description logics. Online. URL <http://www.inf.unibz.it/~franconi/dl/course/dlhb/dlhb-01.pdf>. Online, accessed 12-september-2016.
- [208] Franz Baader, Ian Horrocks, and Ulrike Sattler. *Handbook on Ontologies*, chapter Description Logics, pages 21–43. International Handbooks on Information Systems, 2009.
- [209] Ian Horrocks and Peter F. Patel-Schneider. Knowledge representation and reasoning on the semantic web : OWL. Online. URL <http://www.cs.ox.ac.uk/ian.horrocks/Publications/download/2010/HoPa10a.pdf>. Online, accessed 12-september-2016.
- [210] Chih-Kun Ke and Duen-Ren Liu. Context-based knowledge support for problem-solving by rule-inference and case-based reasoning. *International journal of Innovative Computing, Information and Control*, 7(7(A)):3615–3631, July 2011.
- [211] Yu Zhang, Sarath Sreedharan, Anagha Kulkarni, and Tathagata Chakraborti. Plan explicability and predictability for robot task planning. URL <http://arxiv.org/abs/1511.08158>. Online, accessed 12-september-2016.
- [212] Magnus Boman, Paul Davidsson, Nikolaos Skarmas, Keith Clark, and Rune Gustavsson. Energy saving and added customer value in intelligent buildings, 1998. URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.41.6239&rep=rep1&type=pdf>. Online, accessed 12-september-2016.
- [213] Paul Davidsson and Magnus Boman. Saving Energy and Providing Value Added Services in Intelligent Buildings: A MAS Approach. In *Agent Systems, Mobile Agents, and Applications, LNCS*, pages 166–177. Springer-Verlag, 2000.

## BIBLIOGRAPHY

---

- [214] Bing Qiao, Kecheng Liu, and Chris Guy. A multi-agent system for building control. In *Proceedings of the IEEE/WIC/ACM international conference on Intelligent Agent Technology*, IAT '06, pages 653–659. IEEE Computer Society, Washington, DC, USA, 2006.
- [215] Ueli Rutishauser, Alain Schäfer, and A Cooperation Between. Adaptive building automation - a multi-agent approach, 2002. URL <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.10.5760>. Online, accessed 12-september-2016.
- [216] Huaglory Tianfield. A study on the multi-agent approach to large complex systems. In Vasile Palade, Robert Howlett, and Lakhmi Jain, editors, *Knowledge-Based Intelligent Information and Engineering Systems*, volume 2773 of *Lecture Notes in Computer Science*, pages 438–444. Springer Berlin / Heidelberg, 2003.
- [217] L. Martirano. A smart lighting control to save energy. In *(IDAACS), 2011 IEEE 6th International Conference on Intelligent Data Acquisition and Advanced Computing Systems*, volume 1, pages 132–138, September 2011.
- [218] X.H. Wang, D.Q. Zhang, T. Gu, and H.K. Pung. Ontology based context modeling and reasoning using OWL. In *Proceedings of the Second IEEE Annual Conference on Pervasive Computing and Communications Workshops*, pages 18–22, March 2004.
- [219] Ferit Topcu. Context modeling and reasoning techniques. URL [http://www.snet.tu-berlin.de/fileadmin/fg220/courses/SS11/snet-project/context-modeling-and-reasoning\\_topcu.pdf](http://www.snet.tu-berlin.de/fileadmin/fg220/courses/SS11/snet-project/context-modeling-and-reasoning_topcu.pdf). Online, accessed 12-september-2016.
- [220] Anind K. Dey. Understanding and using context. *Personal Ubiquitous Comput.*, 5(1):4–7, January 2001.
- [221] J. Pascoe. Adding generic contextual capabilities to wearable computers. In *ISWC*, pages 92–99, 1998.
- [222] John Mccarthy and Patrick J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, pages 463–502. Edinburgh University Press, 1969.
- [223] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach (2nd Edition)*. Prentice Hall, 2002.
- [224] Christopher Watkins. *Learning from Delayed Rewards*. PhD thesis, King's College, 1989.
- [225] Richard S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3:9–44, 1988.

- 
- [226] Leemon C. Baird III and A. Harry Klopf. Reinforcement learning with high-dimensional, continuous actions. Technical report, Wright Laboratory, 1993.
- [227] C. Gaskett. *Q-Learning for Robot Control*. PhD thesis, Australian National University, 2002.
- [228] Chris Gaskett, David Wettergreen, and Alexander Zelinsky. Q-learning in continuous state and action spaces. In *Australian Joint Conference on Artificial Intelligence*, pages 417–428. Springer-Verlag, 1999.
- [229] Hado van Hasselt and Marco A. Wiering. Reinforcement learning in continuous action spaces. In *Proceedings of the 2007 IEEE Symposium on Approximate Dynamic Programming and Reinforcement Learning*, pages 272–279, 2007.
- [230] Daniel Aronson. Overview of systems thinking. URL: <http://www.inf.unibz.it/~franconi/dl/course/dlhb/dlhb-01.pdf>, 1996. Online, accessed 12-september-2016.
- [231] Gary Bartlett. Systemic thinking a simple thinking technique for gaining systemic focus. In *The International Conference on Thinking - Breakthroughs 2001*, 2001.
- [232] Erika Ábrahám. Modeling and analysis of hybrid systems (lecture notes). Technical report, Faculty of Mathematics, Computer Science, and Natural Sciences RWTH Aachen University, April, 2012.
- [233] Edward A. Lee and Stephen Neuendorffer. Tutorial: Building Ptolemy II models graphically. Technical Report UCB/EECS-2007-129, UC Lawrence Berkely National Laboratory, October 2007.
- [234] S. Goyal, C. Liao, and P. Barooah. Identification of multi-zone building thermal interaction model from data. In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference(CDC-ECC)*, pages 181–186, December 2011.
- [235] Richard K. Strand. Modularization and simulation techniques for heat balance based energy and load calculation programs: the experience of the ASHRAE loads toolkit and EnergyPlus. In *Building Simulation. Proceedings of the 7th International IBPSA Conference on*, pages 747–753. Rio de Janeiro, Brazil, August 2001.
- [236] P. Tuomaala and Valtion teknillinen tutkimuskeskus. *Implementation and Evaluation of Air Flow and Heat Transfer Routines for Building Simulation Tools*. VTT publications: Valtion Teknillinen Tutkimuskeskus. VTT, 2002.
- [237] Hazim Awbi. *AM11 Building Performance Modelling 2015*. CIBSE, 2015.

## BIBLIOGRAPHY

---

- [238] Chartered Institution of Building Services Engineers. *CIBSE Guide A: Environmental Design New 2015*. CIBSE, 2015.
- [239] Adam Neale, Dominique Derome, Bert Blocken, and Jan Carmeliet. Determination of surface convective heat transfer coefficients by CFD. In *Building Science and Technology, 11th Canadian Conference on*, 2007.
- [240] Gilles Fraisse, Christelle Viardot, Olivier Lafabrie, and Gilbert Achard. Development of a simplified and accurate building model based on electrical analogy. *Energy and Buildings*, 34(10):1017–1031, 2002.
- [241] Hiroyasu Okuyama. Building thermal network model based on state-space system theory. In *Proceedings of the 7th International IBPSA Conference on Building Simulation*, pages 1051–1058. Rio de Janeiro, Brazil, August 2001.
- [242] Anastasios I. Dounis and Christos Caraiscos. Fuzzy comfort and its use in the design of an intelligent coordinator of fuzzy controller-agents for environmental conditions control in buildings. *Uncertain Systems*, 2(2):101–112, 2008.
- [243] Inc The MathWorks. Neural network toolbox, user’s guide, version 4. Technical report, 2002. URL [http://www.image.ece.ntua.gr/courses\\_static/nn/matlab/nnet.pdf](http://www.image.ece.ntua.gr/courses_static/nn/matlab/nnet.pdf). Online, accessed 12-september-2016.
- [244] Benjamin C. Kuo. *Automatic Control Systems (7th Ed.)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1995.
- [245] Energyplus energy simulation software. URL <http://apps1.eere.energy.gov/buildings/energyplus/>. Online, accessed 12-september-2016.
- [246] Tobias Maile, Martin Fischer, and Vladimir Bazjanac. Building energy performance simulation tools – a life-cycle and interoperable perspective. *Center for Integrated Facility Engineering (CIFE)*, (WP107), December 2007.
- [247] U.S. Department of Energy. EnergyPlus Input Output Reference –The Encyclopedic Reference to EnergyPlus Input and Output. Technical report, September 2014.
- [248] U.S. Department of Energy. EnergyPlus Engineering Reference –The Reference to EnergyPlus Calculations. Technical report, September 2014.
- [249] Michael Wetter. *A View on Future Building System Modeling and Simulation*, chapter 17. 2011.
- [250] Michael Wetter. Co-simulation of building energy and control systems with the building controls virtual test bed. *Journal of Building Performance Simulation*, 4(3):185–203, 2011.

- [251] Philip Haves and Peng Xu. The building controls virtual test bed – a simulation environment for developing and testing control algorithms, strategies and systems. In *Proceedings of the International Conference on Building Simulation (IBPSA)*, pages 1440–1446. Beijing, China, 2007.
- [252] Michael Wetter. Building Controls Virtual Test Bed, User Manual version 1.1.0. Technical report, UC Lawrence Berkely National Laboratory, January 2012.
- [253] Xiufeng Pang, P. Bhattacharya, Zheng O’Neill, Philip Haves, Michael Wetter, and Trevor Bailey. Real-time building energy simulation using EnergyPlus and the building controls virtual test bed. In *12th Conference of International Building Performance Simulation Association. Proceedings of Building Simulation 2001*, pages 2890–2896, November 2011.
- [254] Christopher X. Brooks, Edward A. Lee, and Stavros Tripakis. Exploring models of computation with Ptolemy II. In *Proceedings of the 8th International Conference on Hardware/software codesign and system synthesis, CODES/ISSS ’10*, pages 331–332. ACM, New York, NY, USA, 2010.
- [255] Edward A. Lee and Haiyang Zheng. Operational semantics of hybrid systems. In Manfred Morari and Lothar Thiele, editors, *Hybrid Systems: Computation and Control*, volume 3414 of *Lecture Notes in Computer Science*, pages 25–53. Springer Berlin Heidelberg, 2005.
- [256] Hugo S. L. C. Hens. *Building Physics: Fundamentals and Engineering Methods with Examples and Exercises*. Wiley VCH, 2007.
- [257] Patrick Lagonotte, Yves Bertin, and Jean-Bernard Saulnier. Analyse de la qualité de modèles nodaux réduits à l’aide de la méthode des quadripôles. *International Journal of Thermal Sciences*, 38(1):51–65, 1999.



# A

## Heat Transfer Through Building Constructions

This Appendix describes the 1-D Heat transfer process through the building envelope. It provides an explanation of the equations and some examples using the building envelope described in Section 5.2.

### A.1 Heat Transfer Through a Material Layer

Heat transfer by conduction through the building envelope is one of the principal components of space cooling/heating loads and energy requirements. Most building walls consist of several homogeneous layers of different materials; each material layer, depending on the thermal properties, can transfer and accumulate a certain amount of heat. To characterize these thermal properties, a material layer, represented in Figure A.1, is defined by its thickness ( $z$ -axis), section area  $A$ , thermal conductivity  $\lambda$ , specific heat capacity  $C_p$  and density  $\rho$ . The specific heat capacity is the energy required to raise a unit mass of the material one unit in temperature, and the thermal conductivity is the property of a material to conduct heat. Heat (or thermal energy) of a body with uniform properties is given by:

$$\text{Heat energy} = mC_pT \quad [\text{J}]$$

where  $m$  is the body mass and  $T$  is the body temperature.

#### A.1.1 Fourier's Law of Heat Transfer

Heat flux is usually modeled as a one-dimensional transient process with constant material properties. Heat flux from regions of higher temperature to regions of lower

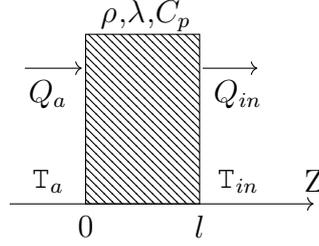


Figure A.1: Heat flux  $Q_a$  and  $Q_{in}$  at the different sections of a material layer with thickness  $l$ , section  $A$ , thermal conductivity  $\lambda$ , specific heat capacity  $C_p$ , and density  $\rho$ , placed between two temperatures  $T_a$  and  $T_{in}$  ( $z = 0$  and  $z = l$ ).

temperature is proportional to the negative temperature gradient and the conductivity of the material:

$$Q = -\lambda A \frac{\partial T}{\partial z} \quad (\text{A.1})$$

### A.1.2 Conservation of Energy

Consider that the layer in Figure A.1 is uniform ( $\lambda, C_p$  and  $\rho$  are all constant) with non-uniform temperature lying on the  $z$ -axis. Also assume that the sides are insulated and only the ends on the  $z$ -axis are exposed, with no heat source within the layer. Consider an arbitrary thin slice of the layer of width  $\Delta z$  between  $z$  and  $z + \Delta z$ , such that the temperature throughout the slice is assumed constant and represented as  $T(z, t)$ . Thus,

$$\text{Heat energy of the slice} = \rho A \Delta z C_p T(z, t).$$

By conservation of energy,

$$\begin{aligned} \text{Change of heat energy} \\ \text{of slice in time } \Delta t \end{aligned} = \begin{aligned} \text{Heat from} \\ \text{left boundary} \end{aligned} + \begin{aligned} \text{Heat from} \\ \text{right boundary} \end{aligned}$$

From (A.1) we get,

$$\rho A \Delta z C_p T(z, t + \Delta t) - \rho A \Delta z C_p T(z, t) = \Delta t \left( -\lambda A \frac{\partial T}{\partial z} \right)_z - \Delta t \left( \lambda A \frac{\partial T}{\partial z} \right)_{z+\Delta z}$$

Rearranging yields,

$$\frac{T(z, t + \Delta t) - T(z, t)}{\Delta t} = \frac{\lambda}{C_p \rho} \left( \frac{\left( \frac{\partial T}{\partial z} \right)_{z+\Delta z} - \left( \frac{\partial T}{\partial z} \right)_z}{\Delta z} \right)$$

Taking the limit  $\Delta t, \Delta z \rightarrow 0$ , gives the following Fourier continuity equation,

$$\frac{\partial T(z, t)}{\partial t} = \alpha \frac{\partial^2 T(z, t)}{\partial z^2} \quad (\text{A.2})$$

where

$$\alpha = \frac{\lambda}{\rho C_p} [\text{m}^2 \text{s}^{-1}]$$

represents the **thermal diffusivity** that measures the ability of the layer material to conduct thermal energy relative to its ability to store thermal energy.

### A.1.3 Laplace Transform

Considering the heat conduction equation given by (A.2) in the form:

$$\frac{\partial^2 \mathbb{T}(z, t)}{\partial z^2} - \frac{1}{\alpha} \frac{\partial \mathbb{T}(z, t)}{\partial t} = 0$$

with initial conditions

$$T(x, 0) = 0,$$

the Laplace transformation yields

$$\frac{\partial^2 \mathbb{T}(z, s)}{\partial z^2} - \frac{s}{\alpha} \mathbb{T}(z, s) = 0. \quad (\text{A.3})$$

The general solution of the differential equation (A.3) is

$$\mathbb{T}(z, s) = b_0 \cosh \left( \sqrt{\frac{s}{\alpha}} z \right) + \frac{b_1}{\sqrt{\frac{s}{\alpha}}} \sinh \left( \sqrt{\frac{s}{\alpha}} z \right) \quad (\text{A.4})$$

with

$$b_0 = \mathbb{T}(0, s) = T_a \quad \text{and} \quad b_1 = \frac{\partial \mathbb{T}(0, s)}{\partial z} = -\frac{Q_a}{\lambda A}$$

From (A.1) and (A.4) we obtain the following heat flux expression:

$$Q(z, s) = -\lambda A \left[ b_0 \sqrt{\frac{s}{\alpha}} \sinh \left( \sqrt{\frac{s}{\alpha}} z \right) + b_1 \cosh \left( \sqrt{\frac{s}{\alpha}} z \right) \right] \quad (\text{A.5})$$

### A.1.4 Transmission Matrix

The material layer, represented in Figure A.1 with thickness  $l$ , can be represented by the two-port network/quadrupole represented in Figure A.2 [257]. The resulting transmission equation, given by (A.4) and (A.5), relates the temperatures and heat fluxes at both sides of the layer:

$$\begin{bmatrix} \mathbb{T}_{in}(s) \\ \mathbb{Q}_{in}(s) \end{bmatrix} = \begin{bmatrix} \cosh \left( \sqrt{\frac{s}{\alpha}} l \right) & -\frac{1}{\lambda A \sqrt{\frac{s}{\alpha}}} \sinh \left( \sqrt{\frac{s}{\alpha}} l \right) \\ -\lambda A \sqrt{\frac{s}{\alpha}} \sinh \left( \sqrt{\frac{s}{\alpha}} l \right) & \cosh \left( \sqrt{\frac{s}{\alpha}} l \right) \end{bmatrix} \begin{bmatrix} \mathbb{T}_a(s) \\ \mathbb{Q}_a(s) \end{bmatrix} \quad (\text{A.6})$$

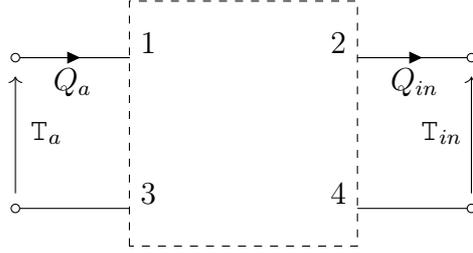


Figure A.2: Thermal two-port associated with a construction layer.

### A.1.5 Electrical Analogy

Considering the electrical analogy using the conductive thermal resistance and thermal capacitance given by (5.5) and (5.6), the resulting transmission equation, given by (A.6), becomes:

$$\begin{bmatrix} T_{in}(s) \\ Q_{in}(s) \end{bmatrix} = \begin{bmatrix} \cosh(\sqrt{sRC}) & -\frac{R}{\sqrt{sRC}} \sinh(\sqrt{sRC}) \\ -\frac{\sqrt{sRC}}{R} \sinh(\sqrt{sRC}) & \cosh(\sqrt{sRC}) \end{bmatrix} \begin{bmatrix} T_a(s) \\ Q_a(s) \end{bmatrix}. \quad (\text{A.7})$$

By rearranging (A.7) (inverting the transmission matrix), we obtain:

$$\begin{bmatrix} T_a(s) \\ Q_a(s) \end{bmatrix} = M(s) \begin{bmatrix} T_{in}(s) \\ Q_{in}(s) \end{bmatrix}, \quad (\text{A.8})$$

where  $M(s)$  is the transmission matrix in terms of Laplace variable  $s$ , given by:

$$M(s) = \begin{bmatrix} \mathcal{A}(s) & \mathcal{B}(s) \\ \mathcal{C}(s) & \mathcal{D}(s) \end{bmatrix} = \begin{bmatrix} \cosh(\sqrt{sRC}) & \frac{R}{\sqrt{sRC}} \sinh(\sqrt{sRC}) \\ \frac{\sqrt{sRC}}{R} \sinh(\sqrt{sRC}) & \cosh(\sqrt{sRC}) \end{bmatrix} \quad (\text{A.9})$$

### A.1.6 Model Reduction

Taking into account that the elements of the transmission matrix  $M$  can be expressed as Taylor series:

$$\mathcal{A}(s) = \mathcal{D}(s) = 1 + \frac{sRC}{2!} + \frac{(sRC)^2}{4!} + \frac{(sRC)^3}{6!} + \frac{(sRC)^4}{8!} + \dots$$

$$\mathcal{B}(s) = R \left[ 1 + \frac{sRC}{3!} + \frac{(sRC)^2}{5!} + \frac{(sRC)^3}{7!} + \dots \right]$$

$$\mathcal{C}(s) = sC \left[ 1 + \frac{sRC}{3!} + \frac{(sRC)^2}{5!} + \frac{(sRC)^3}{7!} + \dots \right]$$

a reduced model of  $M$  can be obtained from the truncated Taylor series of its elements [240]. The second-order limited developments of the transmission matrix are given by:

$$\begin{aligned}\mathcal{A}(s) &= \mathcal{D}(s) \approx 1 + \frac{sRC}{2} + \frac{(sRC)^2}{24} \\ \mathcal{B}(s) &\approx R \left( 1 + \frac{sRC}{6} + \frac{(sRC)^2}{120} \right) \\ \mathcal{C}(s) &\approx sC \left( 1 + \frac{sRC}{6} \right)\end{aligned}\tag{A.10}$$

## A.2 Multi-Layer Walls

Most building walls consist of more than three layers as shown in Figure A.3 and can also be represented by the two-port network represented in Figure A.2. Heat transfer through a  $n$ -layer wall can also be given by (A.8), with the transmission matrix  $M$  being the product of all layer transmission matrices including the superficial heat transfers  $M_{ext}$  and  $M_{int}$  at both sides as follows:

$$M(s) = M_{in}(s)M_n(s) \dots M_1(s)M_{ext}(s)\tag{A.11}$$

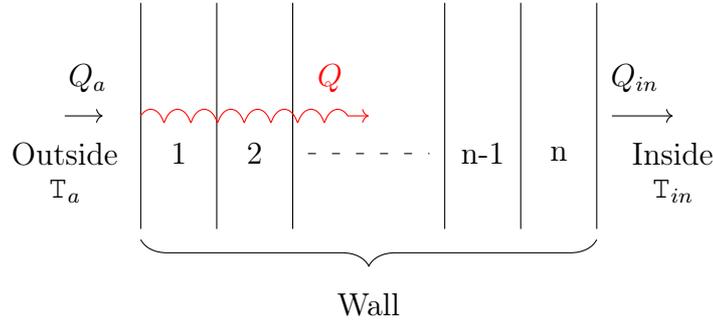


Figure A.3: Heat transfers and construction of a multi-layer wall.

The total resistance  $R$  and capacitance  $C$  of the wall, given respectively by (5.8) and (5.9), include the resistance and capacitance of each of the  $n$  layers including the exterior and interior convective resistances  $R_{ext}$  and  $R_{in}$ , calculated according to (5.7). If all transmission matrices  $M_i, i = 1, \dots, n$  are reduced and given by (A.10), then the second-order matrix  $M$  can be directly written in the form:

$$M(s) = \begin{bmatrix} 1 + sm_1 + s^2m_2 & R + sn_1 + s^2n_2 \\ sC + s^2o_2 & 1 + sm_1 + s^2m_2 \end{bmatrix}\tag{A.12}$$

and parameters  $m_1, m_2, n_1, n_2, o_2$  are calculated and obtained directly from (A.11).

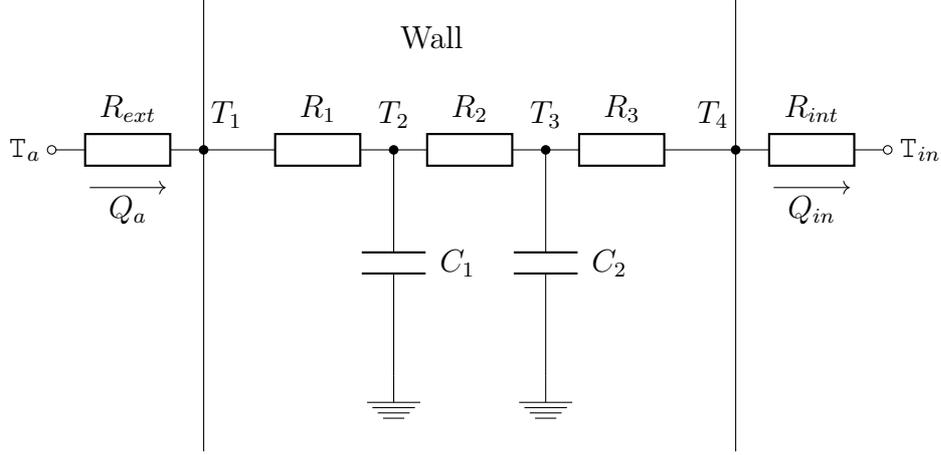


Figure A.4: 3R2C Wall model.

### A.2.1 Modeling of a Multi-Layer Wall using a 3R2C Model

A multi-layer can be modeled using two internal capacities with the 3R2C wall model illustrated in Figure A.4. Adjusting the 3R2C model to represent the heat transfer through the all the layers of the wall (with the total resistance  $R$  and capacitance  $C$ ) equates to finding parameters  $\alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2$  such that:

$$\begin{aligned}
 R_1 &= \alpha_1 R \\
 R_2 &= \alpha_2 R \\
 R_3 &= \alpha_3 R \\
 C_1 &= \beta_1 C \\
 C_2 &= \beta_2 C
 \end{aligned} \tag{A.13}$$

where in steady state conditions verify:

$$\begin{aligned}
 \alpha_1 + \alpha_2 + \alpha_3 &= 1 \\
 \beta_1 + \beta_2 &= 1
 \end{aligned} \tag{A.14}$$

Considering the two different two-port networks represented in Figure A.5 and their associated transmission matrices, the transmission matrix of the 3R2C model is given by the product of the matrices relating to each component as follows:

$$\begin{aligned}
 M_{3R2C}(s) &= \begin{bmatrix} 1 & R_1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ sC_1 & 1 \end{bmatrix} \begin{bmatrix} 1 & R_2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ sC_2 & 1 \end{bmatrix} \begin{bmatrix} 1 & R_3 \\ 0 & 1 \end{bmatrix} \\
 &= \begin{bmatrix} 1 & \alpha_1 R \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ s\beta_1 C & 1 \end{bmatrix} \begin{bmatrix} 1 & \alpha_2 R \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ s\beta_2 C & 1 \end{bmatrix} \begin{bmatrix} 1 & \alpha_3 R \\ 0 & 1 \end{bmatrix} \\
 &= \dots = \begin{bmatrix} 1 + sRCx_1 + (sRC)^2 x_2 & R + sR^2 C x_3 + s^2 R^3 C^2 x_4 \\ sC + s^2 RC^2 x_5 & 1 + sRCx_6 + (sRC)^2 x_7 \end{bmatrix}
 \end{aligned} \tag{A.15}$$

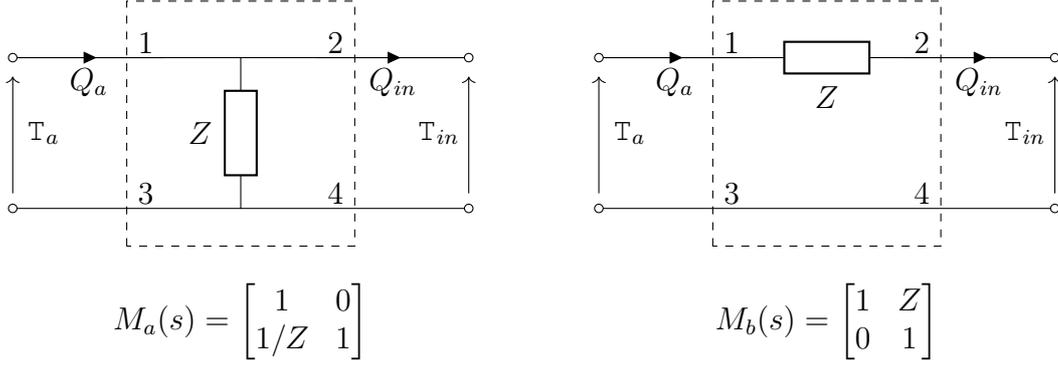


Figure A.5: Two-port network configurations with their associated transmission matrices below.

with

$$\begin{aligned}
x_1 &= \alpha_2 \beta_2 + \alpha_1 \\
x_2 &= \alpha_1 \alpha_2 \beta_1 \beta_2 \\
x_3 &= \alpha_1 \alpha_2 \beta_1 + \alpha_3 (\alpha_2 \beta_2 + \alpha_1) \\
x_4 &= \alpha_1 \alpha_2 \alpha_3 \beta_1 \beta_2 \\
x_5 &= \alpha_2 \beta_1 \beta_2 \\
x_6 &= \alpha_3 + \alpha_2 \beta_1 \\
x_7 &= \alpha_2 \alpha_3 \beta_1 \beta_2.
\end{aligned} \tag{A.16}$$

Parameters  $\alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2$  are directly obtained by identifying matrices  $M_{3R2C}(s)$  and  $M(s)$ , given by (A.12) as follows:

$$\begin{aligned}
x_1 &= x_6 = m_1/RC \\
x_2 &= x_7 = m_2/(RC)^2 \\
x_3 &= n_1/R^2C \\
x_5 &= o_2/RC^2 \\
x_4 &= n_2/R^3C^2
\end{aligned} \tag{A.17}$$

by dividing  $x_2$  by  $x_4$  we obtain:

$$\alpha_3 = n_2/(m_2R)$$

by dividing  $x_4$  by  $x_5$  we obtain:

$$\alpha_1 = n_2/(\alpha_3 R^2 o_2)$$

and  $\alpha_2$  is obtained by the condition linked to steady state given by:

$$\alpha_2 = 1 - \alpha_1 - \alpha_3.$$

Parameter  $\beta_2$  can be obtained from  $x_1$  and  $\beta_1 = 1 - \beta_2$ .

Resistors  $R_1$  and  $R_3$  include  $R_{ext}$  and  $R_{int}$ , since they were added to  $M(s)$  in (A.11). The final values of  $R_1$  and  $R_2$ , according to the full model shown in Figure A.4, are obtained by subtracting  $R_{ext}$  and  $R_{int}$ , respectively.

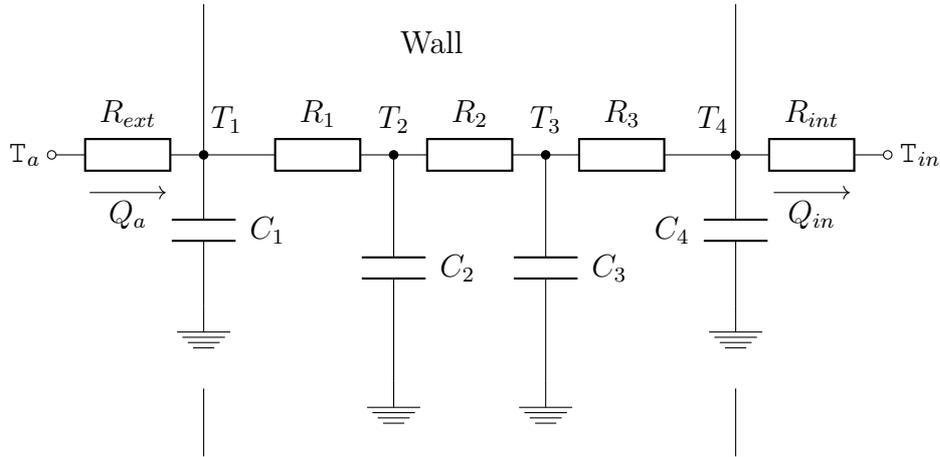


Figure A.6: 3R4C Wall model.

### A.2.2 Modeling of a Multi-Layer Wall using a 3R4C Model

For shorter simulation time steps, it is necessary to consider the effects of the thin layer on the surfaces of the wall. For this purpose, the model 3R2C must be updated so that temperatures pares  $T_a/T_1$  and  $T_4/T_{in}$  are not instantly coupled. Fraisse et al. [240] proposes transferring 5% of the two internal capacities to each of the surfaces of the wall using the 3R4C model shown in Figure A.6.

## A.3 Simulation Results

In this section a multi-layer wall is simulated for the building described in Section 5.2.3, with no doors and no windows. Walls and roof are composed by the materials listed in Table 5.2. Simulation results are presented for both the 3R2C and 3R4C models.

### A.3.1 3R2C Wall Model

The thermophysical characteristics of the multi-layer walls are modeled using the 3R2C model represented in Section A.2.1. For the entire TZ, the model is illustrated in Figure A.7. In this model temperatures  $T_1, T_2, T_3$  and  $T_4$  and  $T_{in}$  evolve according

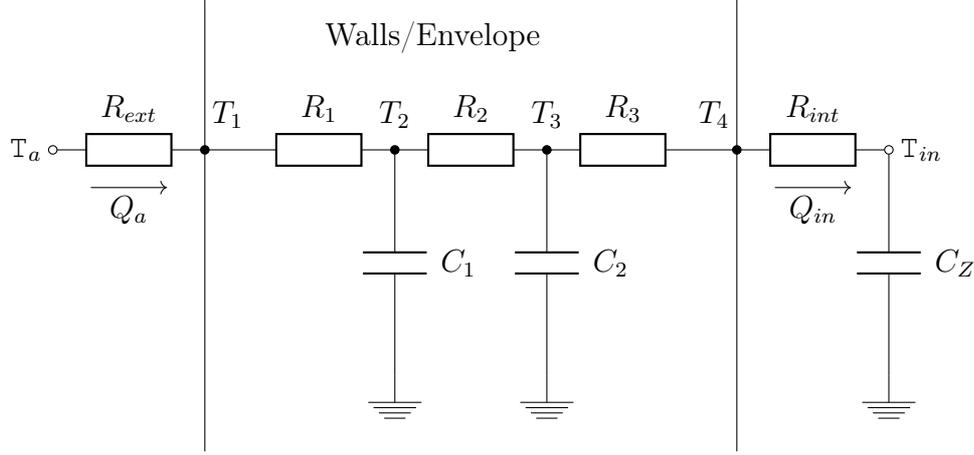


Figure A.7: TZ model using a 3R2C RC model for the envelope.

to the following set of equations:

$$\dot{T}_2 = -\left(\frac{1}{C_1(R_{ext}+R_1)} + \frac{1}{C_1R_2}\right)T_2 + \frac{1}{C_1R_2}T_3 + \frac{1}{C_1(R_{ext}+R_1)}T_a$$

$$\dot{T}_3 = \frac{1}{C_2R_2}T_2 - \left(\frac{1}{C_2R_2} + \frac{1}{C_2(R_{int}+R_3)}\right)T_3 + \frac{1}{C_2(R_{int}+R_3)}T_{in}$$

$$\dot{T}_{in} = \frac{1}{C_Z(R_{int}+R_3)}T_3 - \frac{1}{C_Z(R_{int}+R_3)}T_{in}$$

$$T_1 = \frac{R_{ext}}{R_{ext}+R_1}T_2 + \frac{R_1}{R_{ext}+R_1}T_a$$

$$T_4 = \frac{R_{int}}{R_3+R_{int}}T_3 + \frac{R_3}{R_3+R_{int}}T_{in}$$

To obtain the second-order reference matrix  $M(s)$  given by (A.12), the transmission matrix was calculated using (A.8) with the second order limited developments given by (A.10) for each transmission matrix  $M_1$ ,  $M_2$  and  $M_3$ . By following the procedure described in Section A.2.1, and using the Matlab symbolic toolbox, we obtained the following values for the 3R2C model:

$$R_1 \approx 175.428 \times 10^{-6} \text{ K W}^{-1}$$

$$R_2 \approx 350.605 \times 10^{-6} \text{ K W}^{-1}$$

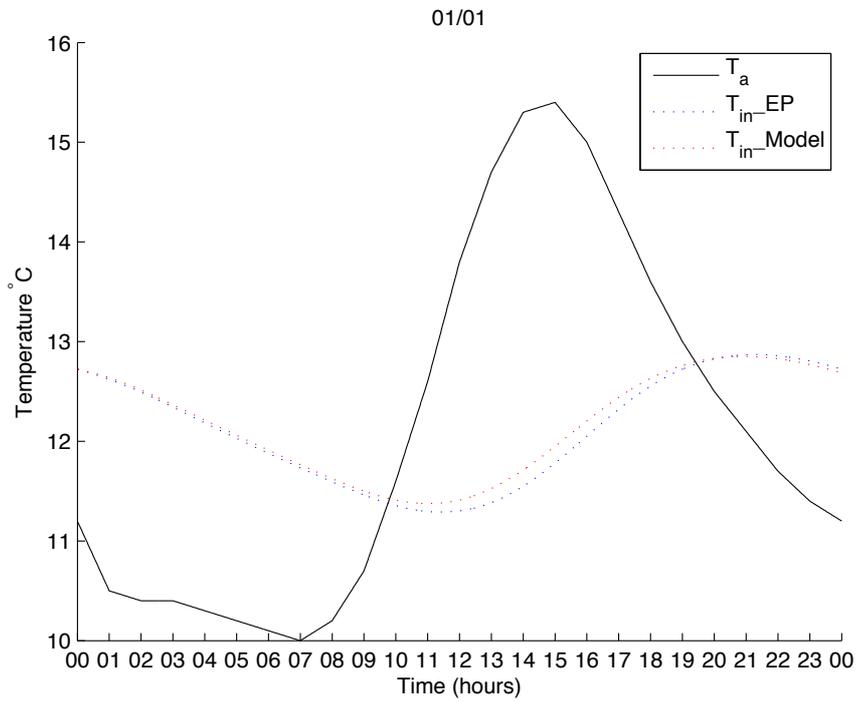
$$R_3 \approx 168.303 \times 10^{-6} \text{ K W}^{-1}$$

$$C_1 \approx 33.618 \times 10^6 \text{ J K}^{-1}$$

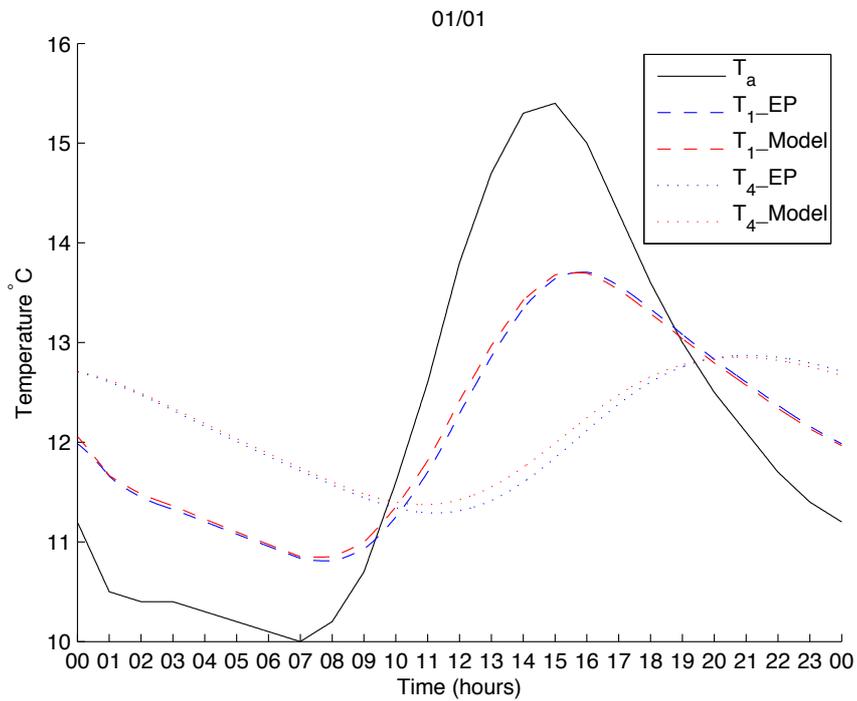
$$C_2 \approx 31.426 \times 10^6 \text{ J K}^{-1}$$

Figure A.8 shows the resulting temperatures  $T_{in}$ ,  $T_1$  and  $T_4$ , given by the 3R2C model ( $T_{in\_Model}$ ,  $T_1\_Model$ ,  $T_4\_Model$ ) and by EnergyPlus ( $T_{in\_EP}$ ,  $T_1\_EP$ ,  $T_4\_EP$ ). Figure A.9 shows the histogram of errors between  $T_{in\_Model}$  and  $T_{in\_EP}$ , were  $MAE = 0.0612^\circ\text{C}$  and  $MAX = 0.1683^\circ\text{C}$ .

APPENDIX A. HEAT TRANSFER THROUGH BUILDING CONSTRUCTIONS



(a) Outdoor and indoor temperatures.



(b) Wall surface temperatures.

Figure A.8: Temperature evolution for a multi-layer wall obtained with the 3R2C model ( $\_Model$ ) and EnergyPlus ( $\_EP$ ).

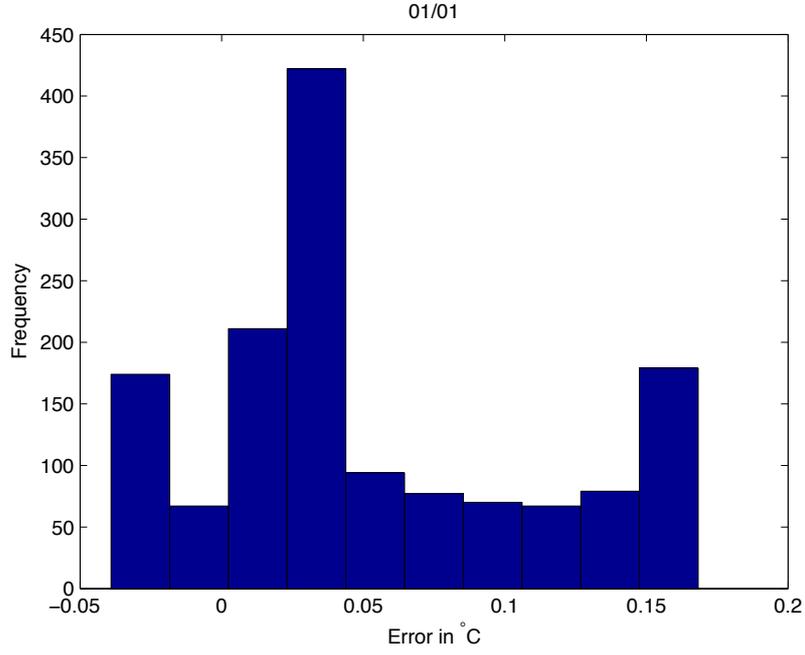


Figure A.9: Histogram of errors between  $T_{in\_Model}$  and  $T_{in\_EP}$  using the 3R2C model.

### A.3.2 3R4C Wall Model

In this section the walls are modeled using the 3R4C network shown in Figure A.6. For this purpose, we calculated the new capacitors as follows:

$$\begin{aligned}
C_1 &= 0.05 \times C\beta_1 \approx 1.681 \times 10^6 \text{ J K}^{-1} \\
C_2 &= 0.95 \times C\beta_1 \approx 31.938 \times 10^6 \text{ J K}^{-1} \\
C_3 &= 0.95 \times C\beta_2 \approx 29.855 \times 10^6 \text{ J K}^{-1} \\
C_4 &= 0.05 \times C\beta_2 \approx 1.571 \times 10^6 \text{ J K}^{-1}
\end{aligned}$$

In the 3R4C model temperatures  $T_1, T_2, T_3$  and  $T_4$  and  $T_{in}$  evolve according to the following set of equations:

$$\begin{aligned}
\dot{T}_1 &= -\left(\frac{1}{C_1 R_{ext}} + \frac{1}{C_1 R_1}\right)T_1 + \frac{1}{C_1 R_1}T_2 + \frac{1}{C_1 R_{ext}}T_a \\
\dot{T}_2 &= \frac{1}{C_2 R_1}T_1 - \left(\frac{1}{C_2 R_1} + \frac{1}{C_2 R_2}\right)T_2 + \frac{1}{C_2 R_2}T_3 \\
\dot{T}_3 &= \frac{1}{R_2 C_3}T_2 - \left(\frac{1}{C_3 R_2} + \frac{1}{C_3 R_3}\right)T_3 + \frac{1}{C_3 R_3}T_4 \\
\dot{T}_4 &= \frac{1}{C_4 R_3}T_3 - \left(\frac{1}{C_4 R_3} + \frac{1}{R_{int} C_4}\right)T_4 + \frac{1}{C_4 R_{int}}T_{in} \\
\dot{T}_{in} &= \frac{1}{C_Z R_{int}}T_4 - \frac{1}{C_Z R_{int}}T_{in}.
\end{aligned}$$

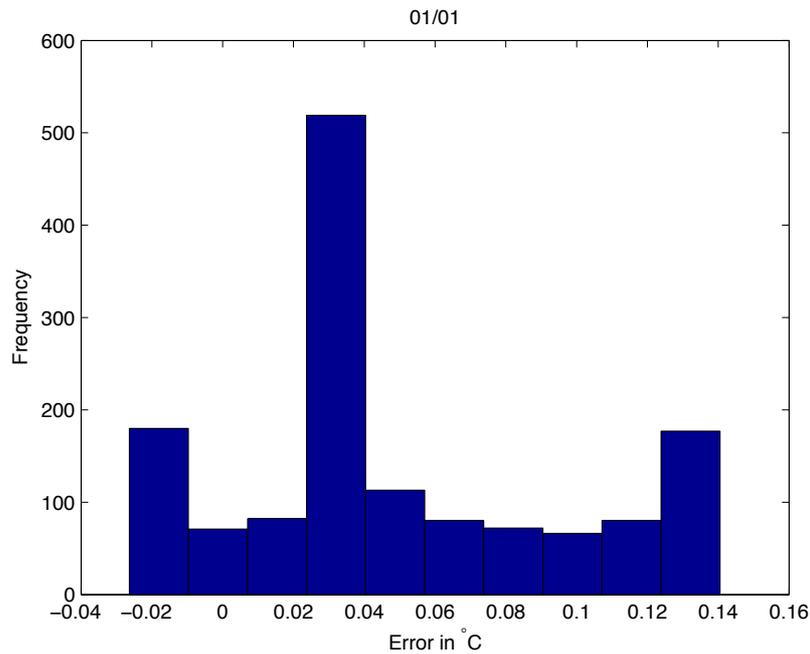
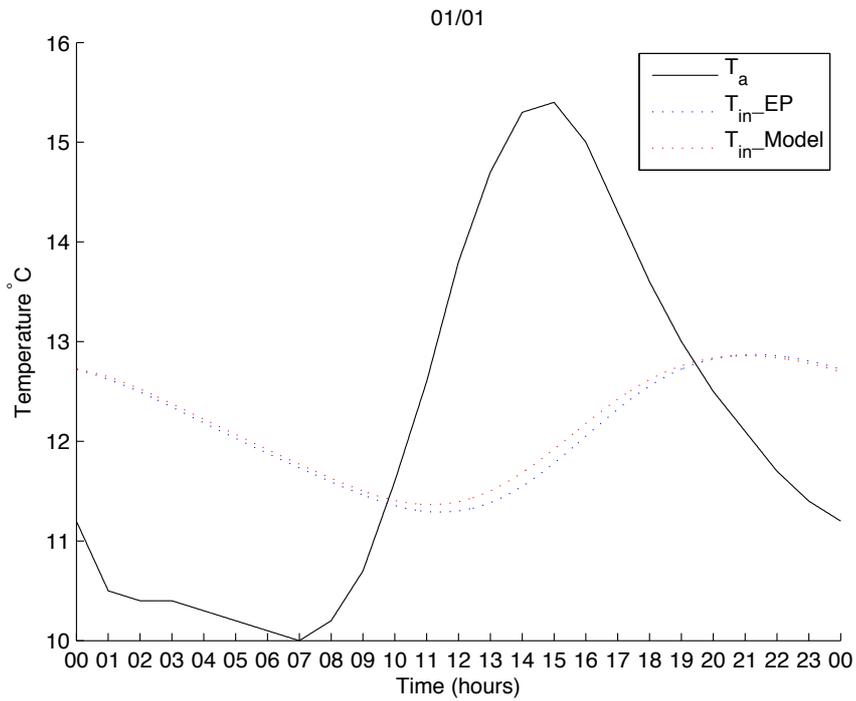
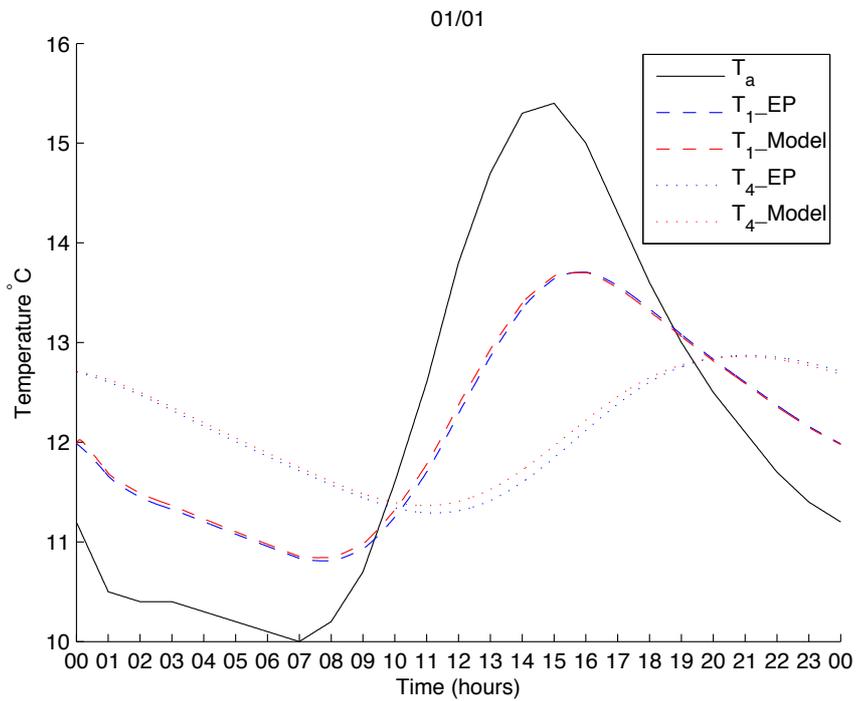


Figure A.10: Histogram of errors between  $T_{in\_Model}$  and  $T_{in\_EP}$  using the 3R4C model.

Using the 3R4C model, Figure A.11 shows the resulting temperatures  $T_{in}$ ,  $T_1$  and  $T_4$ , given by the model and by EnergyPlus. Figure A.10 shows the histogram of errors between  $T_{in\_Model}$  and  $T_{in\_EP}$ , where  $MAE = 0.0556^\circ\text{C}$  and  $MAX = 0.1405^\circ\text{C}$ . Compared to the 3R2C model, the error is slightly smaller as expected. Therefore, this model is selected to represent the building walls in the full-scale RC model described in Chapter 4.



(a) Outdoor and indoor temperatures.



(b) Wall surface temperatures.

Figure A.11: Temperature evolution for a multi-layer wall obtained with the 3R4C model ( $\_Model$ ) and EnergyPlus ( $\_EP$ ).