UNIVERSITY OF SOUTHAMPTON

Efficient Control of Domestic Space Heating Systems and Intermittent Energy Resources

by

Athanasios Aris Panagopoulos

A thesis submitted in partial fulfillment for the
degree of Doctor of Philosophy

in the
Faculty of Physical Sciences and Engineering
Electronics and Computer Science

October 2016

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING
ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Athanasios Aris Panagopoulos

Meeting the ever-growing global energy demand while reducing carbon emissions is one of the most prominent challenges of our era. In this context, efficient control of an operation, service or production process is a key tool to achieve this goal. While there are many opportunities for efficient control within the energy sustainability agenda, this work focuses on domestic space heating systems and intermittent energy resources. This is because in many countries, such as the UK and the US, the domestic sector accounts for more than 20% of the total energy consumption and over 40% of this share is related to space heating. In addition, in recent years, an increasing number of intermittent energy resources, such as photovoltaic systems and wind turbine generators are being integrated into the grid. As such, efficient control of domestic space heating systems and intermittent energy resources can lead to a major reduction in energy consumption and the corresponding CO_2 emission.

In more detail, domestic space heating automation systems (DHASs) aim to optimize the control process of domestic space heating systems with minimum user-input. Moreover, in the case of electricity-based heating, such systems can also incorporate economic control to exploit the energy buffer that heating loads provide in order to shift the heating consumption according to financial incentives, such as variable electricity import tariffs and/or the availability of cheap electricity coming from house-integrated intermittent energy resources. In the latter case, the financial benefits of economic control can be further amplified in domestic coalitions where a number of houses share their energy generation to minimize the collective energy imported from the grid.

Against this background, the first main strand of work in this thesis is to develop a new DHAS, AdaHeat, that overcomes limitations of previous approaches regarding: (i) their efficiency in dealing with the thermal dynamics of houses, (ii) their efficiency in dealing with the inherent uncertainty of the occupancy schedule in domestic settings, (iii) their usability and effectiveness in meeting the user preferences, (iv) their ability to work in conjunction with a diverse range of heating systems, and (v) their ability to efficiently consider economic control in the case of electricity-based heating, exploiting also, for the first time, the aforementioned coalition potential. The backbone of AdaHeat is an ada-

ptive model predictive control approach along with a new general heating schedule planning algorithm based on dynamic programming. In the case of economic control in the presence of house-integrated intermittent energy resources, our planning approach relies on stochastic predictions of the shared intermittent energy resource power output. To this end, we also develop a new adaptive site-specific calibration technique to improve such predictions based on Gaussian process modeling. We present a thorough evaluation of the proposed system, and show its effectiveness in terms of Pareto efficiency and usability criteria against state-of-the-art DHASs. We also show that collective economic control, in the presence of house-integrated IERs, can improve heating cost-efficiency by up to 60%, compared to independent economic control, and even more when compared to no economic control.

The second strand of work is concerned with increasing the efficiency of intermittent energy resources themselves, through efficient control. In particular, specifically for photovoltaic systems, solar tracking can be used to orient the system towards the greatest possible levels of incoming solar irradiance. This can increase the power output of a photovoltaic system by up to 100%. However, current solar tracking techniques suffer from several drawbacks: (i) they usually do not consider the forecasted or prevailing weather conditions; even when they do, they (ii) rely on complex closed-loop controllers and sophisticated instruments; and (iii) typically, they do not take the energy consumption of the trackers into account. As such, in this work, we propose PreST; a novel, low-cost and generic solar tracking approach that overcomes the above limitations, utilizing optimal control (proposed for the first time for solar tracking). In particular, our approach is able to calculate appropriate trajectories for efficient and effective day-ahead (predictive) solar tracking, based on available weather forecasts (that can come from on-line providers for free). To this end, we propose a new approximating policy iteration algorithm, suitable for large Markov decision processes, and a novel and generic solar tracking consumption model. Our simulations show that our approach can increase the power output of a photovoltaic system considerably, when compared to standard solar tracking techniques, that can lead to significant monetary gains.

As outlined above, apart from their great share in contemporary economies, both domestic space heating systems and intermittent energy resources provide considerable opportunities for energy efficient improvements through efficient control. In this work we exploit this potential and propose respective systems that improve their independent, as well as their interaction, efficiency. This can considerably reduce the respective energy consumption and the corresponding CO_2 emission towards fulfilling our goal for an energy sustainable future.

Contents

List of Figures

List of Tables

List of Main Symbols and Abbreviations

A	Action Set.
C^{Eff}	Heating system efficiency.
G_B	Beam solar irradiance.
G_D	Sky-diffuse solar irradiance.
G_R	Ground-reflected solar irradiance.
G_T	Global solar irradiance.
H	Planning horizon.
J	Optimization objective.
K	Set of azimuth orientation positions.
O	Occupancy probability.
P^{Buy}	Import tariff.
P^{Sell}	Export tariff.
R	Intermittent energy resource power output.
S	State Set.
T^{EN}	Envelope temperature.
T^{FL}	Floor temperature.
T^{IN}	Indoor air temperature.
T^{OUT}	Outside ambient temperature.

T^{SP}	Set-point temperature.
δ	Interval length.
κ	Azimuth orientation position.
λ	Slope orientation position (/weighting parameter).
λ'	SPOT+ weighting parameter.
$\mathbf{a}$	Heating control actions vector.
$\mathbf{i}$	Vector of exogenous process variables.
$\mathbf{x}$	Thermal state vector.
ϕ	Thermal leakage rate.
τ	Interaction id.
Λ	Set of slope orientation positions.
a	Action.
s	State.
AADAT	Azimuth-Altitude Dual Axis Tracking.
DHAS	Domestic Heating Automation System.
EKF	Extended Kalman Filter.
ETP	Equivalent Thermal Parameter.
GP	Gaussian Process.
HSAT	Horizontal Single Axis Tracking.
IER	Intermittent Energy Resource.
MDP	Markov Decision Process.
MPC	Model Predictive Control.
PI	Policy Iteration.
PVS	Photovoltaic System.

ST	Solar Tracking.
STPI	Solar Tracking Policy Iteration.
TSAT	Tilted Single Axis Tracking.
TTDAT	Tip-Tilt Dual Axis Tracking.
VSAT	Vertical Single Axis Tracking.
WTG	Wind Turbine Generator.

Note: In this thesis we use the typical notation of denoting vectors and matrices with bold lower-case and bold upper-case letters respectively. Moreover, when not stated otherwise, a vector is assumed to be a column vector.

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my PhD supervisors Professor Nicholas R. Jennings (Professor of Artificial Intelligence, Vice-Provost for research, Imperial College, UK, `http://www.imperial.ac.uk/people/n.jennings`) and Professor Alex Rogers (Professor of Computer Science, University of Oxford, UK, `https://www.cs.ox.ac.uk/people/alex.rogers`) for their enlightening guidance and inspiring instruction, as well as for supporting and trusting me throughout this study.

Furthermore, I would like to thank Associate Professor George Chalkiadakis (Technical University of Crete, Greece, `http://www.intelligence.tuc.gr/~gehalk`) for the valuable suggestions and fruitful collaborations during these past years. I would also like to thank Dr. Matteo Venanzi (Microsoft London, UK, `http://users.ecs.soton.ac.uk/mv1g10`), Dr. Moody Alam (University of Southampton, UK, `http://www.moodyalam.org`), Dr. Sasan Maleki (University of Southampton, UK, `http://maleki.uk`) and Dr. James Scott (Microsoft Research Cambridge, UK, `https://www.microsoft.com/en-us/research/people/jws`) for the rewarding collaborations. I am also grateful to Dr. Lampros Stavrogiannis (Crimtan, UK, `http://users.ecs.soton.ac.uk/ls8g09/`), Dr. Kostas Lagogiannis (University of Edinburgh, UK, `https://www.inf.ed.ac.uk/people/staff/Konstantinos_Lagogiannis.html`), Dr. Antonis Mytidis (University of Florida, USA, `http://www.phys.ufl.edu/ihepa/people.html`) and Assistant Professor Long Tran-Thanh (University of Southampton, UK, `http://users.ecs.soton.ac.uk/ltt08r`) for constructive comments and/or valuable discussions. At this point, I would also like to thank all the people (both academic and non-academic) in the Agents, Interaction and Complexity (AIC) group, University of Southampton, UK (`http://www.aic.ecs.soton.ac.uk`), as well as in the ORCHID project (`http://www.orchid.ac.uk`), for personal and professional development.

This work had quite high requirements in specific data and resources. I am grateful to all individuals, laboratories, agencies and firms for their help in meeting these requirements. In particular, for financing and supporting my research, I would like to acknowledge the funders and partners of the ORCHID project, i.e., EPSRC (`https://www.epsrc.ac.uk`); RCUK Global Uncertainties (`http://www.paccsresearch.org.uk`); School of Electronics and Computer Science, University of Southampton, UK (`http://www.ecs.soton.ac.uk`; Mixed Reality Lab, University of Nottingham, UK (`https://www.nottingham.ac.uk/research/groups/mixedrealitylab/index.aspx`); Pattern

Analysis and Machine Learning Research Group, University of Oxford, UK (http://www.robots.ox.ac.uk/~parg); BAE Systems (http://www.baesystems.com); Secure Meters Ltd (http://www.securetogether.com); Australian Center for Field Robotics (http://www.acfr.usyd.edu.au); Rescue Global (http://www.rescueglobal.org); as well as the Independent Scientific Panel. Notably, I would like to aknowledge the IRIDIS High Performance Computing Facility, University of Southampton, UK (http://www.southampton.ac.uk/isolutions/staff/iridis.page), and the associated support services, for granting me necessary resources for this data-intensive study. Moreover, I would like to acknowledge Microsoft Research, Cambridge, UK (https://www.microsoft.com/en-us/research/lab/microsoft-research-cambridge) and MyJoulo (http://quby.com/en/solutions/joulo) for provision of necessary data and/or respective acquisition support. Lastly, special thanks goes to Weather Underground (https://www.wunderground.com) for being a valuable online source of free-of-charge archival weather data and weather forecasting reports, as well as to the Intelligent Systems Laboratory, Technical University of Crete, Greece (http://www.intelligence.tuc.gr) for hosting valuable, free-of-charge, data acquisition platforms (c.f., http://www.intelligence.tuc.gr/prototypes.php).

Finally, I would like to thank all the members and friends of our philosophy group (considering also the new people I met over the past few years) for inspiring me, supporting me, and delving me deeper into philosophy and the scientific methodology. I would also like to thank all the old and new friends of mine for the great memories during these past years (Paga, Alex, Stelios, Xristos, Kappa, Emanouel, Dimitris, Kostas, Lampros, Amir, Chetan, Alexandros, Anton, James, Nina, Vaggelis, Valerio, Stathis, Mimi, Paulos, Nikos, Masa, Clara, Felipe, Bea, Thomas, Lucia, Valentina, Matteo, Alexandry, to name a few). Last but not least, I would like to express my love and gratitude to my family: Penelope Koulouri, Lysandros Panagopoulos, Orestis Panagopoulos and Euridiki Koulouri, for supporting me unconditionally, and for their valuable advising.

Στη Θάλασσα/Al Mar
(to the Sea—especially on a sunny day)

Declaration of Authorship

I, Aris-Athanasios Panagopoulos, declare that the thesis entitled *Efficient Control of Domestic Space Heating Systems and Intermittent Energy Resources* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published in a number of conference and journal papers (see Section 1.3.3 for a detailed list).

Signed: ....

Date: ....

Chapter 1

Introduction

Meeting the ever-growing global energy demand while reducing carbon emissions and our dependency on the practically finite[1] (and unevenly distributed around the globe) fossil fuel reserves, is one of the most prominent challenges of our era (Hoffert et al. (2002); Ramchurn et al. (2012); Abas et al. (2015)). In this context, many governments around the world have started to adopt policies to transition to a low carbon economy. For instance, in the United Kingdom (UK), the climate change act of 2008 requires at least an 80% cut in the UK's carbon emissions by 2050, compared to 1990 levels.[2] Even more broadly, the European Union (EU) has put in place legislation to reduce its emissions to 20% below 1990 levels by 2020 (Böhringer et al. (2009)), while sharing a similar objective of at least a further 60% reduction by 2050.[3] On the international scene, following the prominent initial steps of the Kyoto protocol (Oberthür and Ott (1999)), the 2015 United Nations Climate Change Conference, held in Paris, France from 30 November to 12 December 2015, negotiated a legally binding agreement on low-carbon economy policies, from all nations around the globe.[4]

In this context, *energy conservation*, *energy efficiency* and *low-carbon energy generation* have been heralded as the key tools to achieve this goal (Hoffert et al. (2002); Jaffe and Stavins (1994); Greening et al. (2000)). In particular, energy conservation refers to the goal of reducing energy consumption through a direct reduction in the products and/or services consumed (Dietz et al. (2009)). As such, energy conservation aims to

[1]Fossil fuels are formed by anaerobic decomposition of buried dead organisms (Mann et al. (2003)). Taking millions of years to be formed, they are considered as a one-time gift to human kind (Deffeyes (2008)).

[2]According to the "Climate Change Act 2008: Impact Assessment", *Department of Energy and Climate Change*, available on-line at `http://webarchive.nationalarchives.gov.uk/20090311095401/http://www.decc.gov.uk/Media/viewfile.ashx?FilePath=85_20090310164124_e_@@_climatechangeactia.pdf&filetype=4` (retrieved on 2/2016).

[3]According to th "ROADMAP 2050: A practical guide to a prosperous low-carbon Europe (Policy Report)", *European Climate Foundation*, available on-line at `http://www.roadmap2050.eu/attachments/files/Volume1_fullreport_PressPack.pdf` (retrieved on 2/2016).

[4]According to the "Adoption of the Paris agreement", *United Nations*, available on-line at `www.unfccc.int/resource/docs/2015/cop21/eng/l09r01.pdf` (retrieved on 2/2016).

reduce the consumption of energy through altering the behavior of the consumers (e.g., avoiding non-vital journeys or avoiding non-critical appliance usage). Due to this reason, energy conservation is out of the scope of this work. Such behavior changes are mainly facilitated by appropriate education programs and/or financial incentives (e.g., Kok et al. (2011); Zehner (2012)).[5] In contrast, energy efficiency refers to the goal of reducing the energy consumed while providing the same desired work (Patterson (1996)); i.e., reducing the wastage of energy in providing a service or producing a product. As such, energy efficiency aims to reduce the consumed energy without fundamentally altering the consumption behavior of the consumers. Due to this apparent advantage, energy efficiency is the main focus of this work.[6] Finally, low-carbon energy generation refers to the goal of providing the required energy based on resources with substantially lower carbon emissions, compared to conventional fossil-fuel-based generators. These resources can be intermittent energy resources (IERs), such as wind turbine generators (WTGs) and photovoltaic systems (PVSs), or non-intermittent, such as hydroelectric generators and nuclear power stations (Hoffert et al. (2002); Abas et al. (2015)). Although low-carbon energy generation is not the main focus of this work, several challenges for energy efficiency improvements emerge with respect, or within, the low-carbon energy generation agenda (considered in this work, as further discussed below).

Now, energy efficiency improvements are generally achieved by adopting more efficient technologies or more efficient production or service processes (Diesendorf (2007)). However, most of the conventional approaches require considerable investments. These include energy efficient building design and reconstruction (e.g., Yao (2013); Aksoy and Inalli (2006)) and/or development and large-scale adoption of: (i) advanced energy efficient appliances, such as energy efficient space heating systems, freezers, and ovens (Waide et al. (1997)); (ii) advanced energy efficient transportation technology, such as turbocharged engines (Hartman (2007)), electric motors,[7] and alternative fuels (Lee et al. (2014)); or (iii) advanced industrial manufacturing and resource extraction processes (e.g., Einstein et al. (2001); Farla et al. (1997); Worrell et al. (2000)). In contrast, there is an approach to energy efficiency improvements that has minimal additional cost and minimal retrofitting. This is efficient control of an operation, service or production process. In general, the aim of efficient control is to optimize the respective control process

[5] For instance, in California, US, such behavior changes are motivated via tiered energy taxation where the taxes increase along with the energy consumption (Zehner (2012)).

[6] Despite its advantages, the realized gains from energy efficiency improvements are typically less than those expected based on simple linear extrapolation approaches (Small and Van Dender (2005)). In particular, increasing the energy efficiency of services and products often makes them cheaper, and hence their consumption usually increases. For instance, energy efficiency improvements regarding personal transportation, domestic space heating and domestic space cooling have been historically accompanied by considerable respective cost reductions and significant respective usage increase (Greening et al. (2000)). This historically observed phenomenon (typically accounting for $\sim$5% to $\sim$40% of the potential gains) is well-known as the "rebound effect" and should be taken into account when planning energy efficient policies (e.g., Greening et al. (2000); Small and Van Dender (2005); Gottron (2001)).

[7] Electric motors are inherently more efficient than internal combustion engines (Sperling and Gordon (2009)). Moreover, the electricity used to supply them can come from low-carbon energy resources (Ramchurn et al. (2012)).

which can potentially lead to significant energy efficiency improvements (e.g., Oldewurtel et al. (2012); Ebert et al. (2000); Esram et al. (2007)). Given this particularity, efficient control is the particular focus of this work.

There are many opportunities and challenges for efficient control within the broad agenda for an energy sustainable future. Importantly though, in many countries, such as the UK and the US, the domestic sector accounts for more than 20% of the total energy consumption and over 40% of this share is related to space heating.[8] For this reason, improving the energy efficiency of domestic space heating systems (both electricity and non-electricity based) can lead to a major reduction in energy consumption and the corresponding CO_2 emissions. As such, improving the energy efficiency of domestic space heating systems is an indispensable, and one of the most prominent, tasks within the domestic energy efficiency agenda. Moreover, space heating systems provide considerable opportunities for energy efficiency improvements through appropriate control (as further discussed in Section 1.1). Hence, this work focuses on increasing the efficiency of domestic space heating systems, both electricity and non-electricity based, through efficient control.

In addition, within the low-carbon energy generation agenda, an increasing number of IERs are being integrated into the grid (Ramchurn et al. (2012)). Specifically, the share of renewable energy resources mix in the worldwide electricity generation has been on an upward trend since 2004 (from 17.9% in 2003 to 20.8% in 2012), while the mean annual growth rates for solar and wind power (during the period 2002-2012) have been estimated at 50.6% and 26.1%, respectively (Armaroli and Balzani (2007)).[9] This fact raises further challenges and opportunities for efficient control of, the electricity-based, domestic space heating systems, which are also considered in this work (further discussed in Section 1.1). Moreover, increasing the energy efficiency of IERs themselves, can lead to additional and considerable reductions in non-low-carbon energy consumption and the corresponding CO_2 emissions (in a similar direction to domestic space heating). Moreover, IERs also provide considerable opportunities for energy efficiency improvements through efficient control (as further discussed in Section 1.2). Due to these reasons, increasing the energy efficiency of IERs through efficient control is also considered in this work.

The rest of this chapter is structured as follows: In Section 1.1 we introduce our work on improving the energy efficiency of domestic space heating systems and in Section 1.2 we introduce our work on improving the efficiency of IERs. In each one of these sections we provide: (i) a general discussion on the particular problem; (ii) the basic

[8] According to "Annual Energy Review (AER) 2011", *US Energy Information Administration*, available on-line at `www.eia.gov/totalenergy/data/annual/archive/038411.pdf` (retrieved on 7/2014), and "United Kingdom Housing Energy Fact File", *UK Depatment of Energy and Climate Change*, available on-line at `www.gov.uk/government/uploads/system/uploads/attachment_data/file/274766/uk_housing_fact_file_2013.pdf` (retrieved on 7/2014).

[9] According to "Worldwide electricity production from renewable energy sources, 2012 edition", available on-line at `www.energies-renouvelables.org/observ-er/html/inventaire/Eng/methode.asp` (retrieved on 5/2015).

requirements/goals of our approach; and (iii) the research challenges in this respect. Then, in Section 1.3 we discuss the main contribution of our work. There, we also provide a list of the respective academic publications that emerged based on this work. Finally, in Section 1.4 we outline the rest of the thesis.

1.1 Efficient Control of Domestic Space Heating Systems

As discussed above, space heating is a significant share of domestic energy consumption and, hence, improving the energy efficiency of domestic space heating systems (both electricity and non-electricity based) can lead to a major reduction in energy consumption and the corresponding CO_2 emissions. In general, there are many ways of improving space heating energy efficiency, ranging from better insulation of the space to advanced heating technologies such as (air or ground) heat pumps and storage heating technology (e.g., Pacheco et al. (2012); Chua et al. (2010)). However, most of the conventional approaches require considerable additional investments and/or building reconstruction (Pacheco et al. (2012); Širokỳ et al. (2011)). In contrast, one approach to increasing space heating efficiency, that has minimal additional cost and minimal retrofitting, is to optimize the heating control process (Širokỳ et al. (2011)).[10]

To this end, programmable thermostats have been proposed as a means to enable more efficient heating control compared to traditional static thermostats as they enable the occupants to schedule the heating control process according to their requirements. In essence, programmable thermostats assign the task of optimizing the heating control process to the user. However, a number of field studies have shown that their potential gains are rarely realized in domestic settings, mainly due to user misconfigurations (e.g., Peffer et al. (2011); Haiad et al. (2004); Shipworth et al. (2010); Kempton (1986); Karjalainen (2009); Meier et al. (2011)). For these reasons, their Energy Star specification was suspended in 2009 (Lu et al. (2010)). In this context, energy systems research has been increasingly focused on domestic heating automation systems (DHASs) which aim to optimize the heating control process with minimum user-input, utilizing appropriate occupancy prediction technology (e.g., Mozer et al. (1997); Scott et al. (2011); Lu et al. (2010); Urieli and Stone (2013)). By limiting the human interaction, such systems aim to overcome the limitations that arise in manual scheduling of programmable thermostats in domestic settings. As such, in recent years, DHASs are also starting to emerge as commercial products, such as the smart thermostats from Nest, Honeywell and British Gas Connected Homes.[11]

[10]In this work, space heating control refers to supervisory control (typically applied manually via thermostats) and not to the low-level control used to maintain the indoor temperature close to the set-point one with minimum oscillations (Wang and Ma (2008)).

[11]`www.nest.com`; `www.honeywell.com`; and `www.hivehome.com`

Even though there are commercial products available, several engineering and theoretical challenges arise in the context of DHAS (Alan et al. (2016a,b)). In this work, we propose "AdaHeat", a new DHAS, that aims to overcome some of the limitations and shortcomings of previous approaches regarding: (i) their efficiency in dealing with the thermal dynamics of houses, (ii) their efficiency in dealing with the inherent uncertainty of the occupancy schedule in domestic settings, (iii) their usability and effectiveness in meeting the user preferences, (iv) their ability to work in conjunction with a diverse range of heating systems (that are typically employed in domestic settings), and (v) their ability to efficiently consider the economic aspects that arise in controlling electricity-based space heating systems with respect to the electricity market. In the following sections, we first provide a general discussion of research around heating automation systems (in the context of both non-domestic and domestic settings) and the respective lessons learned (Section 1.1.1). Then, in Sections 1.1.2 and 1.1.3 we present the challenges that are specific to domestic settings and the additional challenges that emerge with respect to the electricity market in the case of electricity-based space heating, respectively. There, we also provide the requirements of an appropriate DHAS. Finally, in Section 1.1.4 we discuss the shortcomings of current DHASs.

1.1.1 Heating Automation Systems

Energy research has long been preoccupied with developing heating automation systems for non-domestic buildings (e.g., commercial, industrial, offices) (e.g., Farris and Melsa (1978); Kintner-Meyer and Emery (1995); Kummert et al. (2001); Henze et al. (2004); Liu and Henze (2006)). More recently, however, with the onset of ever-increasing house instrumentation and cloud computing, experimental heating automation systems for domestic settings are also starting to emerge (e.g., Mozer et al. (1997); Scott et al. (2011); Lu et al. (2010); Urieli and Stone (2013)).

In essence, the goal of any heating automation system is to control the heating in order to balance the heating energy consumption and the occupant's thermal discomfort according to their preferences (e.g., eliminate thermal discomfort with the minimum heating energy consumption or minimize the thermal discomfort given a heating energy consumption constraint), with minimum user-input. Nevertheless, in recent years, time-varying electricity prices are being introduced in many countries around the globe, where the rates are higher during peak periods and lower during off-peak periods (e.g., Strbac (2008); Hirth (2013)). This is done to motivate the consumers to shift their consumption to off-peak periods and, as such, enhance the reliable operation of the electrical grid (Torriti et al. (2010)). For instance, in many European countries, such as the UK, Spain and Greece (Torriti et al. (2010); Strbac (2008)), night-time tariffs are usually lower than day-time tariffs, motivating the usage of storage heaters for efficient night-time electricity heating.[12] More advanced time-varying pricing programs, where the prices vary considerably during a day, are also being introduced in many European countries to maximally motivate demand shifting to off-peak periods (for more details see Torriti et al. (2010)). As such, in the case of electricity-based heating systems, usually DHASs also take into account the electricity price in optimizing the heating control process (e.g., Rogers et al. (2011); Halvgaard et al. (2012); Shann and Seuken (2014)); appropriately balancing heating cost (rather than heating energy consumption) and thermal discomfort.

Given the above discussion, one can define two families of control approaches employed in heating automation systems: (i) simple automated heating control, where the objective is to efficiently balance *heating consumption* and thermal discomfort; and (ii) simple economic control, where the objective is to efficiently balance *heating cost* and thermal discomfort. The latter is employed when a variability in the energy cost exists, commonly in the case of electricity-based space heating systems. In Section 1.1.3 we provide a more detailed discussion on economic control with respect to DHASs.

1.1.2 Towards Domestic Heating Automation Systems

As discussed above, energy systems research has long been preoccupied with developing heating automation systems for non-domestic buildings, and, more recently, more research has focused on domestic systems, DHASs, as well. Indeed, such systems are now commercial products (as discussed in Section 1.1). Nevertheless, DHASs provide additional challenges over their, more explored, non-domestic counterparts.

In particular, the thermal dynamics of domestic buildings are harder to model accurately than their non-domestic counterparts as: (i) the occupant's activity is more diverse and highly affects the thermal dynamics of the house (e.g., opening a window, operating an auxiliary heater, or cooking) (Li and Wen (2014); Fux et al. (2014)); (ii) the temperature in adjacent buildings or rooms is rarely observed and/or predicted (Li and Wen (2014); Huang et al. (2013)); and (iii) the local weather observations and forecasting reports are usually less accurate than the domestic ones due to lack of appropriate instrumentation (Li and Wen (2014); Dong et al. (2011)).[13] In this context, reliable thermal modeling of the dynamic domestic thermal characteristics is an important requirement of DHASs.

In addition, the occupancy schedule—which is an essential input to any thermal comfort model (as any comfort is experienced only when the space is occupied)—is typically

[12] That said, storage heating technology suffers from several drawbacks that typically lead to significant insufficiencies and heating ineffectiveness. In particular, storage heating technology is inherently inflexible making it hard for heating to follow the occupancy schedule and the occupants' preferences. As such, in many cases the space is heated even when it is not occupied or it is warmer (or cooler) than preferred (Roebuck (2012a)).

[13] In particular, in many cases the latter are equipped with appropriate sensors and local forecasting technology (e.g., Li and Wen (2014); Dong et al. (2011)) while the former rely on on-line general providers for both the "observations" (current local weather conditions) and forecasting reports.

unknown in domestic settings and needs to be predicted (Kleimingera et al. (2013))—in contrast to commercial buildings where it is typically less variable. In this context, all proposed predictive approaches inevitably retain an uncertainty over this schedule which is modeled in the form of probabilistic estimates (Kleimingera et al. (2013)). In the presence of this uncertainty, sacrificing thermal comfort is typically inevitable to avoid extreme heating energy consumption. For instance, eliminating expected discomfort given very small (but still non-zero) probabilistic occupancy predictions during a particular time period would require the house to be heated at the set-point temperature during that period. However, this can significantly increase the heating energy consumption with very small (or non-existent) improvements in the realized discomfort. Given this, effectively balancing heating consumption and thermal discomfort with respect to the occupant preferences becomes an essential task for DHASs. Clearly, there is a conflict between these two, and, as such, balancing them is a non-trivial two-objective optimization problem. In general, Pareto efficiency (or optimality) is crucial in non-trivial multi-objective optimization (Marler and Arora (2004)), and refers to the solution concept where none of the objective functions can be improved in value without degrading some of the other objective values. That being said, there is a potentially infinite number of Pareto optimal solutions, defining the Pareto optimal set. Hence, ideally a DHAS should be able to capture solutions in the Pareto optimal set.

Finally, domestic heating systems are much more diverse than those used in non-domestic buildings. For example, commercial buildings usually have standard HVAC (heating, ventilation and air conditioning) systems whose properties are relatively well understood, while domestic heating systems are much more diverse in type (e.g., underfloor heating, wall-mounted radiators, or fan heaters) which calls for a general DHAS (that is able to handle a variety of them). Lastly, as for any system that is intended to be used in practice, a DHAS should have a computational complexity and efficiency that allows it to be applied in real settings with limited computational resources, minimum instrumentation and operating time constrains (real time operation). For instance, if heating control decisions need to be made every 5 minutes their computation cannot exceed this time limit with the computing power typically available in domestic settings today (which may be much less than that of a typical personal computer).[14] Moreover, a DHAS should require minimum additional instrumentation installation (e.g., occupant activity sensors, temperature sensors) so that its employment is cost-effective for a typical household.

1.1.3 Integrating Advanced Economic Control

In the previous section we discussed the general requirements for DHASs. However, additional requirements emerge in the particular case of electricity-based space heating systems. In more detail, as outlined in Section 1.1.1, simple economic control emerges as

[14]Clusters and supercomputers are typically not available in domestic settings, and their employment for a DHAS, even through cloud computing, may render this technology cost-ineffective.

an additional requirement for DHASs (in the case of electricity-based heating systems) to cope with potentially time-varying electricity prices. In general, in the case of time-varying electricity prices the consumer houses would be more economically efficient if they: (i) store cheap energy purchased at off-peak periods to follow their demand or (ii) shift their electricity demand to off-peak periods. However, appropriate energy storage technology,[15] that could facilitate the former is not available or cost efficient to install at the moment in domestic settings.[16] Moreover, the latter is effective in domestic settings only when adequate energy capacity is available for shifting loads, so that the individual households do not have to drastically alter their schedules (Ramchurn et al. (2012)). Nevertheless, thermostatically controlled loads, such as refrigerators and space and water heaters, can provide such a buffer as excess energy can be retained as heat (de Nijs et al. (2015)). As such, and given the anticipated increase in electrification of heating (Hawkey (2015)), simple economic control is a basic requirement for DHASs to exploit this shifting potential. For instance, in the context of simple economic control, a house could be heated up a while before an estimated arrival of an occupant (rather than exactly before his/her arrival) which could be cheaper given the particular price variability of the day (for more information see Section 2.1.2). Nevertheless, an additional and more advanced form of economic control also arises as a requirement for DHASs in the emerging electricity market reality: *advanced* economic control.

In more detail, as discussed in the beginning of Chapter 1, within the low-carbon energy generation agenda, a large number of IERs are being integrated into the grid. In this context, many houses are now being equipped with, potentially grid-connected, intermittent energy resources (IER), such as rooftop or stand-alone photovoltaic systems or small wind turbine generators (Jacobson et al. (2015); Zahedi (2010)).[17] Moreover, in many regions, such as in several European countries and US states,[18] such houses can sell energy to the grid, but at a lower export tariff than the import tariff (i.e., the price of buying energy from the grid). In general, this price difference motivates the usage of the own produced energy before buying any additional energy from the grid (which is important for the smooth integration of IERs into the grid (Ramchurn et al. (2012))) and disables market manipulations.[19] Due to these reasons, as well as due to its economic justification (since the wholesale price is appropriately lower than the retail one), this pricing schema is expected to become a standard practice around the world

[15]Existing energy storage technologies include, among others, battery, flywheels and super-capacitors (Liserre et al. (2010)).

[16]In particular, emerging products such as Tesla's Powerwall (`www.teslamotors.com`) and BYD Auto's energy storage systems (`www.byd.com`), have received intense criticism as they are too expensive for domestic settings (Bulman (2015)).

[17]This is also supported by the ever falling cost of such installations (Sioshansi (2016)) and the numerous governmental policies to promote respective investments employed around the world (Hawkey (2015); Sioshansi (2016)).

[18]According to `www.dsireusa.org` and `www.res-legal.eu/en` (accessed on 01/2016)

[19]For instance, without appropriate control, if the export tariff is higher than the import tariff, grid entities with storage capabilities can buy energy from the grid in order to sell it back for profit. Notably, the situation is the same in a variable-tariff scenario in the case that the future export tariff is higher than the current import one.

(Sioshansi (2016)).[20] Given the above, and the fact that the power output of IERs is not human controlled, since it typically depends on the prevailing weather conditions (Boyle (2012)), these houses would be more economically efficient if they: (i) store their produced energy to follow their demand, or (ii) shift their electricity demand to follow their energy production (in a similar manner to simple economic control). In either case, they would minimize their energy import (according to the pricing motivation) and achieve financial gains. Nevertheless, as discussed above, appropriate energy storage technology that could facilitate the former is not available or cost efficient to install at the moment in domestic settings.[16] In this context, the next generation DHASs needs to also incorporate advanced economic control, in houses with electricity-based heating systems and grid-connected IERs, to exploit this shifting potential.

In contrast to simple economic control, the aforementioned benefits of advanced economic control can be further amplified in domestic coalitions, where a number of houses share their energy generation and shift their heating consumption in order to further minimize the energy imported from the grid. Having this happen, requires the market to allow consumer groups to act as an individual, and can be facilitated by forming a micro-grid of physically connected houses (Palizban et al. (2014)) or, in a more general manner, by conducting appropriate contracts through dynamic on-line energy exchange providers, utilizing the national grid and appropriate smart metering (Niese et al. (2012)). Now, smart meters are already being installed in many countries (Krishnamurti et al. (2012)) and dynamic on-line energy exchange is being explored in many countries through initiatives, such as Piclo, Clickpower, and Powershop.[21] As such, direct facilitation of such coalition formation (i.e., without a micro-grid) is likely to be available in the near future (Niese et al. (2012); Michalak et al. (2009)). In this context facilitating this coalition potential with respect to advanced economic control is an additional requirement for DHASs. This means, such systems should also incorporate an appropriate cost allocation mechanism to share the realized gains of the coalition among the members in a practical manner.

Following the above discussion (Sections 1.1.2 and 1.1.3), the basic requirements of a domestic heating automation system can be defined as follows:

1. **Minimal user-input:** Rely to the minimum extent on user-input, fulfilling the aim of being a heating *automation* system.
2. **Reliable thermal modeling:** Thermal characteristics in houses are much more dynamic than in non-domestic buildings and/or appropriate instrumentation much less intense. Reliable thermal modeling in such settings is an important requirement of DHASs.

[20] The alternative schemes used in practice, where the export tariff is equal (or even higher) than the import tariff, are generally temporary and aim to promote respective investments (Sioshansi (2016)).

[21] `www.openutility.com/piclo`; `www.clickpower.in`; `www.powershop.com.au`

3. **Dealing with occupancy uncertainty:** The occupancy schedule is usually unknown and dynamic in domestic settings and needs to be predicted. Dealing with (potentially probabilistic) occupancy estimates is a basic requirement of any DHAS.

4. **Pareto efficiency:** Balancing heating cost and thermal discomfort involves a non-trivial two-objective optimization task. That being said, efficiency in balancing cost and discomfort is a basic requirement of DHAS. Ideally a DHAS should be able to capture solutions in the Pareto optimal set.

5. **Matching the user preferences:** Matching the user preferences in balancing heating cost and thermal discomfort. In this context, three subsequent requirements of an efficient DHAS can be defined:

   (a) **Flexibility:** Be able to capture a sufficiently wide range of balancing points between heating cost and thermal discomfort (ideally within the Pareto optimal set as discussed above) that allows a variety of user preference schemes to be captured.

   (b) **Usability:** Be able to match to the user preferences in choosing one of these balancing points for its operation, via a feasible and effective human-computer interaction procedure (i.e., no population of hard to interpret mathematical equations).

   (c) **Adaptability:** Be able to adapt to potentially time-varying user preferences in trading heating cost and thermal discomfort.

6. **Generality:** Being able to work in conjunction with a diverse range of heating systems and respective technologies that are employed in domestic settings (e.g., fan heaters, underfloor heating systems, heating pumps).

7. **Applicability:** Low computational complexity and efficiency that allows the DHAS to be applicable in real settings with limited computational resources, minimum instrumentation, and operating time constrains.

8. **Integrate simple economic control:** Be able to exploit the shifting potential that arises in houses with electricity-based heating systems with respect to time-varying import tariffs.

9. **Integrate advanced economic control:** Be able to exploit the shifting potential that arises in houses with electricity-based heating systems and domestic IER generation capacity with respect to the difference in the import and export tariffs.

   (a) **Coalition potential:** Be able to exploit the coalition potential that arises with respect to advanced economic control.

   (b) **Cost allocation:** Allocating the collective gains of the coalition among the coalition members in a practical and effective manner.

1.1.4 Research Challenges

A number of DHAS have been proposed in the literature (e.g., Mozer et al. (1997); Scott et al. (2011); Lu et al. (2010); Urieli and Stone (2013); Shann and Seuken (2014)). However, they typically suffer from several drawbacks: (i) they usually rely on a simple experimental thermal model which is not reliable in practice and suitable only for proof-of-concept systems; if not (ii) they do not deal with the highly dynamic nature of house thermal characteristics; (iii) they do not provide a way of choosing the parameterizable coefficients in balancing heating energy consumption and thermal discomfort—the important challenge of matching the occupant's preferences is usually disregarded in DHASs; (iv) they usually rely on heuristic control approaches in dealing with occupancy uncertainty (without providing any guarantees or intuition regarding the performance loss from an approach that fully exploits the probabilistic estimates); if not (v) they rely on computationally expensive approaches that limit their applicability only to experimental settings; and (vi) they are usually heating-system-specific. In addition to the above limitations, there is also a lack of comparison among DHASs, as those are usually benchmarked against simple static timer programs such as "always-on" or "pre-scheduled" heating.

In addition, to date, most of the proposed DHASs focus on efficiently balancing *energy consumption* and thermal discomfort according to the occupant's preferences, based on occupancy predictions estimates. That is, simple (automated) heating control without considering any additional economic aspects other than the amount of energy consumed (e.g., Mozer et al. (1997); Gao and Keshav (2013a); Urieli and Stone (2013); Scott et al. (2011); Lu et al. (2010)). Nevertheless, in recent years, economic control is starting to emerge as an integrated part of DHASs (e.g., Shann and Seuken (2014); Rogers et al. (2011); Halvgaard et al. (2012)), where the systems appropriately balance *heating cost* and thermal discomfort, in the case of electricity-based heating. However, all these works consider simple economic control, merely accounting for variable energy import tariffs over the energy consumed, without considering grid-connected domestic IERs and energy export. In essence, all these systems aim to shift the demand to time-slots where the import tariff is favorable. Moreover, none of these works exploit the potential that coalitions provide in (advanced) economic control. In particular, the potential of domestic coalitions in such a DHAS-integrated economic control has not been exploited until now although: (i) coalition formation is a considerably active research area within the energy sustainability agenda (e.g., Alam et al. (2013); Chalkiadakis et al. (2011); Alam et al. (2015)), and (ii) several works have discussed the potential of thermostatically controlled load aggregations (including domestic space heating loads) for providing regulation services to the grid (e.g., Hao et al. (2013); Callaway (2009); de Nijs et al. (2015)). In addition, work that deals with domestic space heating system aggregations typically assumes the same preferences in balancing heating cost and thermal discomfort among the houses (e.g., Dudley and Piette (2008); Torriti et al. (2010)). However, this is impractical in realistic settings since the households have diverse, and even time-varying, preferences on this balancing (as discussed in Section 1.1.2). In Section 1.3.1 we discuss our contributions regarding DHASs against this background.

1.2 Efficient Control of Intermittent Energy Resources

As discussed in the beginning of Chapter 1, a big part of this work is dedicated to efficient control of intermittent energy resources, IERs. In particular, in recent years, a large number of IERs, such as photovoltaic systems, PVSs, and wind turbine generators, WTGs, are being integrated into the grid. Given this, increasing the efficiency of IERs can lead to a considerable reduction in non-low-carbon energy consumption and the corresponding CO_2 emissions (as discussed in the beginning of Chapter 1). In this context, there are many ways of improving IER energy efficiency, such as adopting better IER locations, advancing the IER power production technology, and adopting efficient IER control.

In more detail, the power output of IERs depends highly on the prevailing weather conditions. In particular, the power output of a PVS depends mostly on the irradiance incident to the photovoltaic (PV) module and the operating temperature of the module (Luque and Hegedus (2011)). In general, PVSs favor greater levels of incident solar irradiance and generally cooler environments which lead to lower operating temperatures (i.e., lower ambient temperature and appropriate wind speeds) (Luque and Hegedus (2011)). On the other hand, the power output of a WTG depends mostly on the prevailing wind speed and wind temperature, with the WTG generally favoring higher wind speeds and lower wind temperatures (cooler air is denser, increasing power output) (Burton et al. (2011)). As such, a straightforward way to increase IER efficiency is through better localization of the IERs, i.e, choosing the location (and altitude) for the IER installation where the weather pattern favors its operation.[22] However, installing IERs in locations that maximally favor their operation is not always possible and most of the time IERs are installed in suboptimal locations due to a variety of reasons such as space limitation, budget availability or energy transmission limitations (Kurokawa (2012)).

Another way to increase IER efficiency is through increasing the efficiency of the IER power production technology itself. In the case of PVSs, this would mainly mean improving the solar cell efficiency. At this time, shipped solar cells, commonly made from silicon, typically convert sunlight into electricity with an efficiency of around 10% to 20% (Green et al. (2015)).[23] However, considerable improvements are theoretically possible

[22] For instance, concentrating the global PVS power production to locations that maximally favor their operation, such as deserts or near-desert locations, have been proposed as a means to significantly increase their efficiency (Kurokawa (2012)). That said, centralizing IER production raises significant challenges for the grid operation and energy transmission in particular (Kurokawa (2012)).

[23] Solar cells with higher efficiency are also commercially available, however they come with significantly higher cost and, as such, are far less commonly employed (Green et al. (2015)).

when considering recent advantages in nanotechnology, biotechnology, and the materials and physical sciences used, and solar cells of higher efficiency are anticipated to be available in reasonable cost in the following years (Lewis (2007)). Regarding WTG, Betz's law gives the theoretical upper limit of the power that can be extracted from the total kinetic energy of the air flowing through a WTG at ∼59% (Betz (2014)). At this time, shipped WTG have a peak operational efficiency of around 75% to 80% of the Betz limit due to inefficiencies, such as the rotor blade friction and drag, gearbox losses, and generator and converter losses (Burton et al. (2011)). As such, further improvements in the inherent efficiency of WTG are also possible. However, using the state-of-the-art in IER energy conversion technology is highly costly (both for PVS and WTG) and, furthermore, many experimental technologies are not yet available as commercial products.

An alternative way to improve the efficiency of an IER with minimum additional investments is by efficient control. In general, IER control has two main goals: (i) to ensure the safe operation of the IER, and (ii) to maximize its energy production. In more detail, WTGs do not operate at maximum efficiency across a range of wind speeds, while high wind speeds raise safety issues for their operation. In this context, WTG control aims to ensure that the WTG is in a stall position when the wind speeds exceed the safety limit and that the WTG operates as close as possible to the maximum efficiency over all wind speeds (typically achieved through variable turbine speed control) (Burton et al. (2011)). In a similar manner, PVSs have a non-linear output efficiency which is subject to the prevailing environmental conditions and the total resistance applied to the photovoltaic (PV) module (Luque and Hegedus (2011)). As such, maximum power point tracking is equipped to maximize its efficiency by applying the proper resistance for any given environmental conditions (Luque and Hegedus (2011)). Beside these low-level control approaches, specifically for PVS, an additional higher-level control exists that can greatly increase the efficiency of the PVS. In particular, solar tracking (ST) techniques can be used to orient the system towards the greatest possible levels of incoming solar irradiance (Mousazadeh et al. (2009)). In this context, part of the solar tracking control objective is to also ensure that the PVS is in a safe position when the wind speed raises safety issues for the construction. Importantly, depending on location and season, solar tracking can increase the PVS power output by up to 100% (Mousazadeh et al. (2009)). As such, the effectiveness of the ST technique used is crucial for the overall efficiency of the PVS. Moreover, several drawbacks of current ST approaches enable us to increase its effectiveness while its nature enables the investigation of techniques that have long been the focus of AI research (as discussed in the following paragraphs). Hence, in this work we focus on increasing the PVS efficiency through advance ST control.

In particular, in this work we propose novel ST techniques that aim to improve the effectiveness and efficiency of ST with low additional cost. In the following sections, we first provide a general discussion of ST (Section 1.2.1). Then, in Section 1.2.2, we provide the requirements of appropriate ST. Finally, in Section 1.2.3 we discuss the shortcomings and limitations of current ST approaches.

1.2.1 Solar Tracking Systems

To increase the power output of a PVS, solar tracking, ST, techniques can be used to orient the system towards the greatest possible levels of incoming solar irradiance. ST can significantly improve the efficiency of a PVS (by up to 100% (Mousazadeh et al. (2009))). As such, a wide range of ST approaches have been developed over time which are generally distinguished based on their tracking architecture and the drive type.

In particular, according to the tracking architecture, two general ST categories can be defined: (i) single-axis ST, and (ii) dual-axis ST. In the former the trackers used have one degree of freedom allowing rotation over one axis (while the other axes is fixed), while in the latter the trackers have two degrees of freedom enabling rotation over two axes (that are typically normal to one another). Many single-axis and dual-axis ST architectures exist. For instance, typical representatives of single-axis ST are: (i) the horizontal single axis tracking (HSAT), where the axis of rotation is horizontal with respect to the ground; (ii) the vertical single axis tracking (VSAT), where the axis of rotation is vertical with respect to the ground; and (iii) the tilted single axis tracking (TSAT), where the axis of rotation is tilted with respect to the ground. Regarding dual-axis ST some representatives consider: (i) the tip-tilt dual axis tracking (TTDAT), where the panel array is mounted on the top of a pole which allows rotation over two axis of rotation, that are typically normal to each other; and (ii) the azimuth-altitude dual axis tracking (AADAT) which have a similar operation to TTDAT except that it typically uses a large ring on the ground (along with a series of rollers) to rotate the panel, instead of the pole mounting used in TTDAT. As such, AADAT has the advantage of distributing more evenly the panel weight and, hence, is able to consider larger arrays compared to TTDAT. For a comprehensive review on ST architectures see Roebuck (2012b).

In addition to the tracking architecture used, current ST approaches can also be distinguished based on their drive type. In this context, three main categories can be defined (Mousazadeh et al. (2009)): (i) active ST, (ii) passive ST, and (iii) manual ST. Active ST relies on, typically electrical, motors to move the PVS and is the most common ST implementation. For this reason, active ST is the main focus of our work. On the other hand, passive ST generally relies on thermal expansion effects to move the PVS. Although passive ST can demonstrate high efficiency (Mousazadeh et al. (2009)), it is far less common than active ST, and, hence, it is not considered in this work. Lastly, in some developing countries, tracking is being assigned to operators who adjust the PVS orientation manually. Manual tracking provides robustness and employment positions for the people in the proximity of the site (Mousazadeh et al. (2009)). Due to these reasons, although less common than passive ST, manual ST is also considered in this work.

Now, active ST can be further classified according to the controller type. In particular, in active tracking the motors are driven by a controller that can operate in a *closed-loop* or in an *open-loop* fashion (i.e., with or without making use of any feedback, respectively) with respect to the environmental conditions. In this aspect, typical open-loop trackers are the *chronological* trackers, which follow the sun based on a chronological model of its motion (Reda and Andreas (2004)). On the other hand, closed-loop controllers typically utilize appropriate sensors (e.g., photodiode or thermopile pyranometers (Sengupta et al. (2012))) to dynamically orient the system towards the higher level of incident solar irradiance.[24] For a comprehensive review on the ST drive types and the respective approaches see Mousazadeh et al. (2009).

1.2.2 Towards Optimal Solar Tracking Systems

Given the great efficiency improvements that can be achieved, the effectiveness of the ST system used is crucial for the overall efficiency of a PVS. In this context, an active ST system should be able to consider the weather conditions in order to orient the system towards the greatest level of solar irradiance to maximize its output. Ideally, this should be achieved with low cost, and without the need for sophisticated instrumentation and complex closed-loop controllers (that require installation by an expert). This is needed so that the ST system is available to the general public, considering also small producers with limited budget availability (e.g, house integrated PVSs or small PV plants). In the same context, the ST system should be able to operate in real settings and meet the operating time constrains with limited computational resources so that the cost of the necessary hardware is not limiting its availability. Now, since tracking itself comes with a cost, a ST system should be able to also consider the consumption cost due to tracking itself (apart from the weather conditions) in order to avoid redundant movement and the subsequent unnecessary tracking cost increase. In a similar manner, a ST system should ideally also consider the maintenance cost that generally also increases along with the active tracking operation time. Last, but not least, given the great variety of existing ST architectures, a ST technique should be generic enough to be able to consider a wide range of ST architectures typically employed in practice (e.g., HSAT; TSAT; TTDAT; or AADAT systems).

Following the above discussion the basic requirements of an effective solar tracking, ST, system can be defined as follows:

[24] That said, some interesting techniques exist which use the PV modules themselves as sensors (Dasgupta et al. (2010)).

1. **Generality:** Being able to work in conjunction with the diverse range of ST architectures that are typically employed in practice.

2. **Applicability:** Low cost, low complexity and effectiveness that allows the ST system to be widely applicable in real settings with limited budget availability, limited resources, and operating time constrains.

3. **Performance optimality:** Optimal or near-optimal performance that leads to highly efficient ST performance. In this context, three subsequent requirements can be defined:

   (a) **Considering the prevailing weather conditions:** The power output of a PVS highly depends on the incident solar irradiance, and, as such, considering the prevailing weather conditions is a basic requirement for efficient ST.

   (b) **Considering the consumption cost:** Tracking itself is costly, and, as such, avoiding redundant movement can lead to efficiency improvements.

   (c) **Considering the maintenance cost:** The maintenance cost increases with tracking. In this context, avoiding unnecessary movement can lead to additional efficiency improvements.

1.2.3 Research Challenges

A number of ST systems have been proposed in the literature (Mousazadeh et al. (2009)). However, all these approaches typically suffer from several drawbacks with respect to the requirements of an effective ST system as defined above (Section 1.2.2). In particular, open-loop active ST, although simple, does not take into account the forecasted or prevailing weather conditions (e.g., the degree of cloud coverage or humidity levels). On the other hand, closed-loop active ST, although accounting for the prevailing weather conditions, usually depends on expensive and sophisticated instruments (Mousazadeh et al. (2009)) that limits its availability to the general public. On top of that, existing closed-loop active ST approaches typically do not take into account the energy consumption caused by the tracking itself, nor do they consider the system's maintenance cost; they simply greedily turn the system towards the perceived highest irradiance values. In Section 6.4 we discuss our contribution regarding ST against this background.

1.3 Research Contributions

In this work we focus on efficient control of: i) domestic space heating systems and ii) IERs. Our respective research contributions are detailed in the following paragraphs.

1.3.1 AdaHeat: A General Adaptive Domestic Heating Automation System

Regarding efficient control of domestic space heating systems, to address the shortcomings discussed in Section 1.1.4, we propose a new general DHAS, AdaHeat, that balances heating cost and thermal discomfort in an infinite horizon optimization manner, learns an adaptive thermal model of the system under control on-line and does planning to fully exploit the occupancy probabilities. To this end, our system employs a model predictive control (MPC) (see Section 2.5) approach utilizing adaptive gray-box thermal modeling (i.e., adaptive modeling that relies on simplified physical equations—see Section 2.2) and a new general algorithm for planning that fully exploits the probabilistic occupancy estimates via dynamic programming. As such, AdaHeat: (i) is able to effectively account for the highly dynamic thermal characteristics of houses, (ii) is able to work in conjunction with both linear and non-linear optimization objectives and system models, (iii) and is general enough to consider a wide range of heating systems. Due to these reasons, AdaHeat can be considered as a general framework where specific models can be inserted to give particular characteristics. Moreover, AdaHeat adapts to the user preferences in balancing cost and discomfort as it relies on a single parametrization factor that is learned on-line. We evaluate our approach with data coming from a real house that employs underfloor heating (which constitutes a challenging testbed on the generality of our approach both in terms of thermal modeling and control) where we show the benefits of incorporating adaptive gray-box thermal modeling in DHASs as well as the effectiveness of our approach in balancing heating cost and thermal discomfort. In this context, we also run a comparison over existing heating DHASs and AdaHeat where we show that the latter leads to a more stable performance, in terms of Pareto efficiency, in various operational settings.

In addition, AdaHeat incorporates both simple economic control and advanced economic control, considering, for the first time, issues associated with domestic grid-connected IERs and export tariffs (in the case of electricity-based heating). To exploit the coalition potential that arises in advanced economic control, our respective approach is applicable to both coalitions and single houses (as an extreme case of a single-house coalition). In particular, regarding advanced economic control, we propose the formation of house coalitions that utilize IER stochastic predictions to coordinate their heating system operation, ahead of time, so as to minimize their energy import (or, more formally, to buy and sell energy to the grid when the tariffs are favorable), as an aggregate. In this context, we propose an effective heuristic heating schedule planning approach for advanced economic control in domestic coalitions. Our respective solution: (i) has a complexity that scales in a linear and parallelizable manner with the size of the coalition, and (ii) handles different preferences, in balancing heating cost and thermal discomfort among the individual households. Our planning approach relies on stochastic predictions of the shared IER power output. To achieve this we develop a new adaptive site-specific calibration technique to improve such predictions based on Gaussian process (GP) modeling. Moreover, AdaHeat also incorporates a practical cost allocation mechanism to share the realized gains of the coalition that respects individual rationality and allocation efficiency. Finally, we demonstrate the effectiveness of AdaHeat with respect to advanced economic control through real data evaluation, in a contemporary market reality with flat import and export tariffs. Specifically, we show that collective advanced economic control (within the coalition) can improve heating cost-efficiency by up to 60%, compared to independent advanced economic control, and even more when compared to AdaHeat's independent simple economic/heating control.

In more detail, we extend the state-of-the-art as follows:

- We show how adaptive gray-box thermal modeling (see Section 2.2) can be incorporated in DHASs to capture the highly dynamic nature of domestic thermal characteristics. This is the first DHAS that incorporates adaptive gray-box thermal modeling.
- We propose a new general algorithm for planning in the context of MPC (see Section 2.5), that optimally accounts for the occupancy probabilities and efficiently searches over the heating schedule space, utilizing dynamic programming.
- We evaluate our approach with data coming from a real house that employs underfloor heating (which constitutes a challenging testbed on the generality of our approach both in terms of thermal modeling and control) where we show the benefits of incorporating adaptive gray-box thermal modeling in DHASs as well as the effectiveness of our approach in balancing heating cost and thermal discomfort.
- We run a comparison over existing DHAS approaches and an improved approach that fully exploits the occupancy probabilities (i.e. AdaHeat) where we show that the latter leads to a more stable performance, in terms of Pareto efficiency, in various operational settings. In this context we also provide significant insights into the agents' usability in various settings.
- We are the first to show how advanced economic control that considers IER generation capacity of the house and export tariffs (along with import tariffs), can be incorporated in DHASs. We also show, for the first time, how to exploit the potential of coalitions in this context.
- We propose a new heuristic planning approach for collective advanced economic control, that has a complexity that scales in a linear and parallelizable manner with the coalition size. Moreover, our approach can handle the diverse household preferences in balancing heating cost and thermal discomfort.

- We propose a practical approach for advanced economic control and demonstrate its effectiveness through evaluating with real data, in a contemporary market reality. We show that collective advanced economic control can significantly improve heating cost-efficiency compared to independent advanced economic control, and even more when compared to independent simple economic/heating control.

1.3.2 PreST: A Dynamic Programming Predictive Solar Tracking Approach

Regarding efficient control of IERs, to address the shortcomings discussed in Section 1.2.2, we develop novel low-cost active (and manual) ST techniques that can be used in both an open-loop or a closed-loop manner. We do not make use of expensive equipment or sensors, but the backbone of our approach is the estimation of the optimal trajectories a day before, based on weather forecasts coming from online providers for free—hence, we name our ST approach PreST (as an abbreviation for predictive solar tracking). To this end, we employ a recently developed web tool, RENES (Panagopoulos et al. 2012), that predicts the power output of a PVS given available weather forecasts.[25] These predictions form the reward dynamics of a new policy iteration (PI) technique we devise. The technique, Solar Tracking Policy Iteration (STPI), alternatively optimizes over action sub-spaces. Although optimizing over sub-problems in an alternating fashion is a generally common concept elsewhere (Bezdek and Hathaway (2002)), this is the first time that such an optimization technique is proposed for Markov decision processes (MDPs).

Importantly, the method makes use of a novel tracking system consumption model we have developed (and which can be extended to account for maintenance and other costs). The method is appropriate for dual-axis tracking, and is shown to be much more efficient than the, also sensor-less, chronological ST. We also provide four additional control methods: a PI method specialized for single-axis tracking, two near-optimal myopic methods (one specialized for single-axis and one for dual-axis), and a method that enables us to define the next-day-optimal positioning for any fixed-orientation (yet re-adjustable) PVS operating within the geographical region of a given weather station, enabling efficient manual tracking. In particular, the efficiency of the latter is higher than positioning the system according to *yearly-optimal* fixed-orientation estimates, and the method can be easily extended to define the *weekly-optimal* PVS orientation. Moreover, our methods are shown to improve the power output of a PVS even when compared to closed-loop *sensor-based* ST. In particular, our results show that our approach outperforms all benchmark methods (i.e., chronological, sensor-based and/or fixed-orientation). Though we evaluate our approach with respect to the popular azimuth-altitude dual axis trackers, AADAT, and vertical single axis trackers, VSAT, we note that it can be used in conjunction with many other ST systems (e.g., TSAT, HSAT, or TTDAT). It is worth noting here that our

[25] `www.intelligence.tuc.gr/renes`

next-day policy comes complete with an expected PVS power output estimation. This is crucial for the smooth integration of PVSs into the electrical grid—since it is essential that short-term predicted PVS production estimates are available, notwithstanding their intermittent nature (Ramchurn et al. (2012)). Last, but not least, all our methods come with guarantees of near-optimality.

Although dynamic programming can naturally find the optimal solution to the ST problem, this work is the first to propose such an approach. This is probably due to the fact that an appropriate reward model had not been devised until now (due to the lack of free and ready-to-use power output estimates, and an appropriate consumption model). We resolve this issue, and thus contribute to the state-of-the-art, as follows: first, we employ the recent method of (Panagopoulos et al. 2012) to get PVS power output estimates;[26] and, second, we devise here, for the first time, a generic, parameterizable tracker power consumption model.

In more detail, in this work we extend the state of the art in the following ways:

- We are the first to formalize ST as a dynamic programming problem, and propose novel low-cost and generic ST techniques that come both with optimality or near-optimal performance guarantees, and complete with an expected PVS power output estimation.
- We propose, for the first time, a generic tracking system consumption model to model the ST dynamics.
- We propose a new policy iteration approximation algorithm for large state-action spaces that considers the first alternative optimization dynamic programming algorithm for MDPs.
- We run an evaluation based on real data to show that our approach outperforms all commonly employed ST benchmark techniques (i.e., chronological, sensor-based and/or fixed-orientation), which can lead to significant monetary gains.

1.3.3 Academic Publications

This work has lead to the following academic publications:

1. Muddasser Alam, **Athanasios Aris Panagopoulos**, Alex Rogers, Nicholas R. Jennings, and James Scott, (2014) *"Applying Extended Kalman Filters to Adaptive Thermal Modelling in Homes"*, poster abstract at **ACM BuildSys 2014**, Memphis, US, 05 - 06 Nov 2014. 2pp.

[26]We note here, that one could alternatively use the work of Chakraborty et al. (2012), which, however, unlike RENES, does not come with a ready-to-use web tool, and requires the availability of historical PVS-specific production output data.

2. **Athanasios Aris Panagopoulos**, Georgios Chalkiadakis, and Nicholas R. Jennings, (2015) *"Towards Optimal Solar Tracking: A Dynamic Programming Approach"*, in Proc. of the 29th AAAI Conference on Artificial Intelligence (**AAAI-2015**), Austin, TX, USA, 25 - 30 January 2015, pp. 695-701.

3. **Athanasios Aris Panagopoulos**, Moody Alam, Alex Rogers, and Nicholas R Jennings, (2015) *"AdaHeat: A general adaptive intelligent agent for domestic heating control"*, in Proc. of the 14th Int. Conference on Autonomous Agents and Multi-Agent Systems (**AAMAS-2015**), Istanbul, Turkey, 04 - 08 May 2015, pp. 1295-1303.

4. **Athanasios Aris Panagopoulos**, Sasan Maleki, Alex Rogers, Matteo Venanzi, and Nicholas R. Jennings, (2016) *"Advanced Economic Control of Electricity-based Space Heating Systems in Domestic Coalitions with Shared Intermittent Energy Resources"*, ACM Transactions on Intelligent Systems and Technology (**ACM TIST**).

1.4 Thesis Outline

The rest of this thesis is structured as follows:

- **Chapter 2:** In this chapter, we provide background material and an overview of related work regarding domestic heating automation systems, DHASs and intermittent energy resources IERs. We consider the latter with respect to our work on DHAS-integrated advanced economic control and efficient IER control, i.e. solar tracking).

- **Chapter 3:** In this chapter, we describe our general adaptive DHAS, AdaHeat. AdaHeat accounts for simple heating control and simple economic control, as well as advanced economic control, exploiting also the coalition potential that arises in the latter. In this context, here we also detail our proposed scheme for collective advanced economic control in the context of AdaHeat.

- **Chapter 4:** In this chapter, we provide a thorough evaluation of AdaHeat and a comprehensive comparison of existing state-of-the art heating automation systems. We do so by independently evaluating AdaHeat with respect to simple heating and simple economic control, and with respect to advanced economic control.

- **Chapter 5:** In this chapter, we detail our dynamic-programming-based predictive ST approach, PreST. To do so, here we also outline the necessary astronomical background with respect to ST and provide a detailed discussion on popular ST architectures that are also the main focus of our work.

- **Chapter 6:** In this chapter, we provide a detailed real-data-based evaluation of PreST against commonly employed ST techniques.

- **Chapter 7:** In this chapter, we provide a conclusion discussion considering also the limitations of this thesis and future work directions. Here, we also provide a detailed evaluation of our work against the above-stated requirements.

Chapter 2

Related Work

In this chapter we provide relevant background material and an overview of related work regarding domestic heating automation systems (DHASs) and intermittent energy resources (IERs) (the latter with respect to our work on DHAS-integrated advanced economic control and efficient IER control, i.e. solar tracking). In particular, in Section 2.1 we discuss domestic heating systems with respect to their main characteristics that introduce challenges for efficient modeling and heating control in the context of heating automation systems. In Section 2.2 we provide a general discussion of thermal modeling approaches with a particular focus on gray-box thermal modeling and its adaptive version. Then, in Section 2.4 we overview the literature of approaches for predicting the occupancy schedule. Further on, in Section 2.5 we discuss the control approaches utilized in the context of DHASs, while in Section 2.6 we provide a detailed discussion of model predictive control. Then, in Section 2.7 we review the related work on domestic heating automation systems, as well as on non-domestic heating automation systems that deal with occupancy uncertainty and hence could also be applicable in domestic settings. In Section 2.8 we review the literature of DHASs with respect to (simple and advanced) economic control.

Subsequently, in Section 2.9 we provide a review on stochastic prediction of IER power output which is an essential part of our DHAS-integrated advanced economic control approach (as discussed in Section 1.3.1). Then, in Section 2.10 we provide a general discussion on Gaussian process modeling (GP) which is the basis of our respective stochastic prediction approach. In Section 2.13 we provide a general discussion over Markov decision processes (MDPs) and dynamic programming which consider the backbone key concepts of our proposed solar tracking (ST) approach, as discussed in Section 6.4. Finally, Section 2.14 summarizes.

2.1 Domestic Heating Systems

As discussed in Section 1.1.2, the heating systems employed in domestic settings are diverse in type and the technologies used (e.g, underfloor heating, heating pumps, electric fan heaters and wall-mounted radiators). However, two particular characteristics of a heating system are of great significance in the context of DHAS as they highly influence the methods used for control and/or system modeling. These characteristics are: (i) the magnitude of the thermal lags and (ii) the variability of the heating cost over time. In the following sections we discuss each of these characteristics.

2.1.1 Magnitude of Thermal Lags

In general, the aim of a space heating system is to heat up the air in a space to a particular temperature (i.e., the set-point temperature) and retain it there for an arbitrary timespan. Depending on the thermal mechanism of the heating system employed, there is a thermal lag that is observed in terms of heat being transfered to the air after heating is switched off. In short, this effect is subject to the intermediate nodes of the underlying heat transfer mechanism of the system and the heat flow rates between these nodes.[1] As such, this effect is negligible for some heating systems such as air fan heaters or electric radiant heaters, however considerable for other heating systems such as central heating systems based on underfloor heating technology or wall-mounted radiators. Now, this effect can be taken into account via appropriate thermal modeling. However, the thermal lag magnitude also influences the efficiency of different heating control strategies.

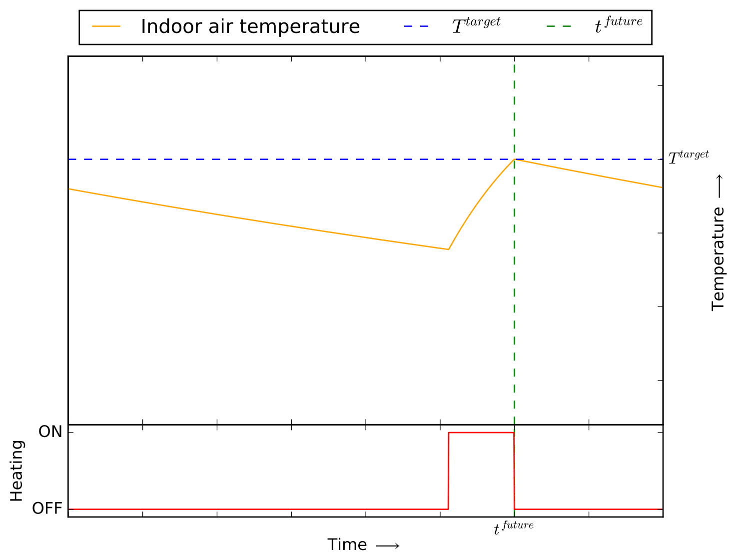

FIGURE 2.1: Heating strategy example I (*without* thermal lags, *without* cost variability)

[1]Here, the usage of the term "node" is associated with the resistor-capacitor circuit representation, usually employed when studying heat transfer mechanisms (see Deng et al. (2010) for more details).

To illustrate this, let us consider the setting where a particular indoor air temperature T^{target} needs to be achieved at a specific time instance in the future, t^{future}. In addition, we want to achieve this with the minimum heating cost. Intuitively, the optimal strategy for a heating system with negligible thermal lags would simply be to have heating switched on for the minimum time required right before t^{future}, so that T^{target} is met exactly at t^{future}, as seen in Figure 2.1.[2] In particular, given that: (i) the heated space is not a closed system (i.e., is not perfectly isolated), (ii) obeys the general laws of thermodynamics (and hence, given other things being equal, the heat transfer rate will be greater for greater temperature differences), and (iii) the surrounding environment is generally cooler, then switching on heating earlier in time would just introduce additional unnecessary heating cost. On the other hand, switching on heating any later would not allow us to meet the target temperature T^{target} at t^{future}. However, if considerable thermal lags are introduced the above defined solution concept is not optimal anymore. In such a setting, the heating window needs to be shifted earlier in time to account for the delay in the thermal response (with respect to the air temperature), deriving a heating schedule as illustrated in Figure 2.2. Now, the optimal heating schedule (i.e., the one that minimizes the heating cost) in such a setting depends on the peculiarities of the system under control such as the specific thermal lags of the heating system, the building thermal characteristics and the building thermal state.

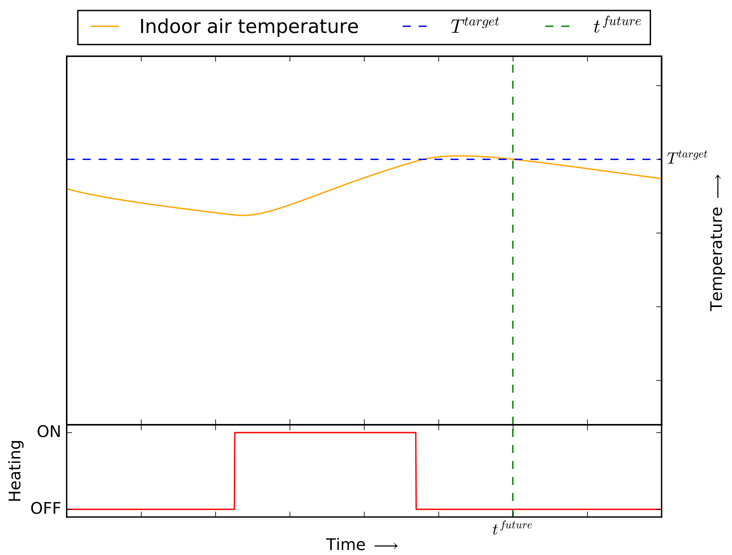

FIGURE 2.2: Heating strategy example II (*with* thermal lags, *without* cost variability)

2.1.2 Variability of Heating Cost

Another characteristic of a space heating system with great significance in the context of DHASs is the variability of the heating cost over time. In general, heating cost variability

[2]Here, we make the additional assumption that the cost of heating remains constant over time, as we further discuss in Sections 2.1.2.

mainly affects electricity-based heating systems (but not exclusively), and can be either *inherent* or *inherited* (to the heating system). In the former case, heating cost variability emerges due to a variability of the overall heating efficiency of the heating system itself. In more detail, the overall heating efficiency stands for the ratio of the useful thermal energy provided over the energy consumed by the heating system. As such, a variability of the heating system efficiency would lead to a variability of the heating cost over time. Now, most conventional heating systems have a relatively fixed efficiency over time. However, electricity-based heating systems that employ technologies that transfer heat from one place to another, i.e., *heat pumps*, have a heating efficiency that varies over time. In more detail, a heat pump's "efficiency"[3] varies with the difference in the temperature between the heat source and the heat sink (Haines and Myers (2009)).[4] This fact calls for careful heating system consumption and thermal modeling, and also influences the efficiency of different heating control strategies in the context of heating automation systems (as illustrated in the following paragraphs).

Now, heating cost variability can also be *inherited* to a heating system due to a time-varying energy cost. In the case of electricity-based heating systems this could be due to: (i) a variability in the electricity import tariffs (which leads to the DHAS requirement of simple economic control, as discussed in Section 1.1.3),[5] and/or (ii) the (time-varying) availability of cheap electricity coming from heating-system-integrated grid-connected intermittent energy resources, IERs (which leads to the DHAS requirement of advanced economic control, as discussed in Section 1.1.3)[6]. The above, call for careful heating system consumption modeling, and also influence the efficiency of different heating control strategies in the context of heating automation systems (as illustrated below).

To illustrate how heating cost variability affects the efficiency of different heating control strategies, let us consider the example introduced in Section 2.1.1 above, where a particular indoor air temperature T^{target} needs to be met at a specific time instance in the future, t^{future}, with the minimum heating cost. In such a case, we argued that, for a heating system with negligible thermal lags, heating-up the space for the minimum time required right before t^{future} could be held as the optimal solution. In this solution concept we made the additional subtle assumption that the heating system does not experience any heating cost variability over time. However, this trivial strategy is not generally optimal for heating systems with variable heating cost. In such a setting,

[3]In order to avoid any misconception, the term "efficiency" is usually avoided in the context of heating pumps, as it has a very specific thermodynamic meaning. In particular, the term coefficient of performance is used instead, to describe the ratio of useful heat movement per work input.

[4]This is also supported by the fact that systems that utilize heat pump technology usually employ additional supplementary heat sources (integrated into the heat pump system or as separate systems) in order to retain a reliable performance of the overall heating system (Haines and Myers (2009)).

[5]As discussed in Sections 1.1.1 and 1.1.3, time-varying electricity import tariffs are being introduced in many countries, as part of demand-side management programs (Ramchurn et al. (2011)), to motivate the consumers to shift their consumption to off-peak periods and enhance the reliable operation of the electrical grid.

[6]As discussed in Section 1.1.3, many houses are now being equipped with potentially grid-connected IERs within the low-carbon energy generation agenda.

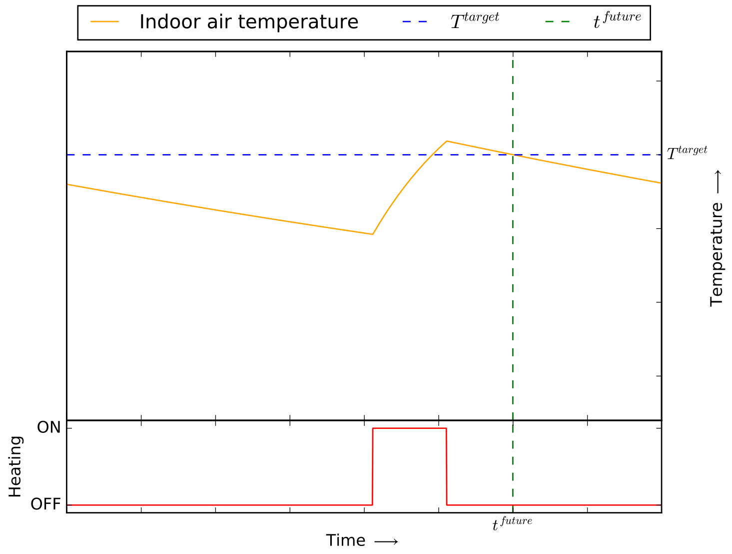

Figure 2.3: Heating strategy example III (*without* thermal lags, *with* cost variability)

heating cost varies over time and, as such, the optimal heating schedule will depend on this particular variation. For instance, if the heating cost is considerably higher during the heating window of our initial simple heating strategy, it might be more economically efficient to shift this window earlier in time (and even expand it) and just let the temperature drop to T^{target} at t^{future}, as illustrated in Figure 2.3. That said, the optimal strategy in such settings will depend on the peculiarities of the problem, including the specific variation of the heating system efficiency. Nevertheless, the above example illustrates that the initial simple policy is not always optimal in this case.

2.2 Thermal Modeling of a House

As discussed in Section 1.1.2, reliable thermal modeling is an essential part of a heating automation system and arises as a particularly prominent challenge in the highly dynamic domestic settings (considering a basic requirement of DHASs, see Section 1.1.3). Now, several thermal modeling approaches have been proposed over time and, in general, can be classified in the following categories (Prívara et al. (2013)): (i) *white-box* (physics-based), (ii) *black-box* (data-driven), and (iii) *gray-box* (combination of physics-based and data-driven) modeling approaches. The later two (i.e. black-box and gray-box approaches) can be further classified to *fixed* or *adaptive*, based on whether the modeled thermal dynamics are assumed to be fixed or time-varying. For an extended review on thermal modeling approaches see Li and Wen (2014) and Prívara et al. (2013).

In more detail, white-box approaches use detailed physics-based equations to model the building thermal dynamics. The parameters of these equations, such as thermal conductance values, heat capacity and thickness of materials, come from detailed surveys (e.g., design plans, manufacturer catalogs) or on-site measurements. A variety of mature white-box software tools are available, such as EnergyPlus[7], ESP-r[8], TRNSYS[9], and GridLAB-D[10]. However, the explicit knowledge required for white-box modeling is not always available, especially in old constructions and domestic buildings, and on-site measurements are typically time consuming and expensive. As such, these approaches are not suitable for our work here because of our requirement for generality and applicability (see Section 1.1.2).

In contrast, black-box approaches use statistical or machine learning techniques (e.g., simple polynomial curve fits, neural networks, and support vector machines) to model the thermal dynamics of a building without the need for any prior knowledge (see for example, Huang et al. (2013) and Ruano et al. (2006)). Based on whether the thermal dynamics are assumed to be fixed or time-varying, the black-box approaches can be further classified as fixed or adaptive respectively, employing different regression and/or training techniques to capture (or not) this variability (e.g., moving training windows, or sequential neural network training approaches). However, in general, black-box approaches are hard to interpret in physical terms and/or to generalize to other systems (Morel et al. (2001)). Moreover, they typically require a large amount of training data in order to demonstrate an adequate and reliable performance (Morel et al. (2001); Li and Wen (2014)). Due to these reasons, such approaches are not followed in our work.

On the other hand gray-box modeling approaches, which could be considered hybrid approaches that combine physical modeling with statistical or machine learning techniques, aim to overcome the aforementioned drawbacks (e.g, Rogers et al. (2013); Ellis et al. (2013)). Gray-box approaches use simplified physical models to capture the thermal dynamics of a building. These models are based on derived equivalent thermal parameters (ETPs), instead of parameters from surveys or on-site measurements. The ETPs are assumed to be fixed or time-varying and are estimated via appropriate statistical, or machine learning, parameter identification methods (e.g., Kalman filters or least squares methods). Depending on whether the ETPs are assumed constant or time-varying, the gray-box approaches can be further classified to fixed or adaptive respectively. In general, using simplified physical models based on ETPs, reduces the requirement of vast amount of training data, and the need for explicit knowledge or on-site measurements. As such, in our work we employ an adaptive gray-box thermal modeling approach. In the following sections we provide a more detailed discussion on gray-box thermal modeling.

[7] `http://apps1.eere.energy.gov/buildings/energyplus/energyplus_about.cfm`
[8] `www.esru.strath.ac.uk/Programs/ESP-r.htm`
[9] `http://sel.me.wisc.edu/trnsys`
[10] `www.gridlabd.org`

2.3 Gray-Box Thermal Modeling

In general, gray-box thermal modeling consists of two relatively independent tasks: (i) defining the simplified physical model (model selection), and (ii) estimating the ETPs used in the model (ETPs learning).

The task of model selection is to derive a model that strikes a balance between accuracy and complexity according to the requirements of the application (Prívara et al. (2012)). In particular, an extremely complex model would potentially be very accurate. However, a high complexity may not enable the model to generalize to other systems, and estimating the ETPs could be computationally expensive. On the other hand, a very simple model may not be able to capture essential characteristics of the thermal dynamics and may demonstrate low predictive performance. For instance, underfloor heating systems typically have considerable thermal lags over their operation. As such, a simple thermal model that considers the heat transfer to be direct between the heat source and the indoor air would not capture the thermal lags of the heating system. In contrast, a more complex model that considers the transfer of heat from the heat source to the indoor air via an intermediate node (i.e., the floor), would be able to capture the thermal lags typically observed in such systems. Typically, balancing complexity and accuracy is accomplished by an iterative procedure which starts with the simplest feasible model (and estimate the corresponding ETPs), and then iteratively refines it into a more complex one in order to identify the most suitable model (e.g., Prívara et al. (2012); Bacher and Madsen (2011); Andersen et al. (2000); Kristensen et al. (2004)).

Now, given a gray-box model, the involved ETPs are typically estimated via statistical, or machine learning, parameter identification methods. The particular parameter identification method used is subject to the complexity of the model (e.g., linear or non-linear) and whether the ETPs are assumed to be fixed or time-varying. Based on the latter assumption, gray-box modeling can be classified as either *fixed* or *adaptive*, respectively.

- **Fixed gray-box thermal modeling:** In fixed gray-box modeling, the ETPs are estimated either once or at infrequent intervals and are assumed to be constant over an arbitrary horizon. Recent representatives of this approach include the heating automation system SPOT+ where linear regression to historical data is applied in order to learn a fixed thermal model of a room in an office building (Gao and Keshav (2013b,a)). Another example considers the MyJoulo project where specially designed USB loggers are used in order to collect temperature readings from homes (Rogers et al. (2013)). Subsequently, non-linear regression is applied for thermal modeling in order to provide personalized advice on energy savings. Another example of fixed gray-box thermal modeling is MatchStick which uses a non-linear least square method to learn the thermal model of individual rooms (Ellis et al. (2013)). In addition, simpler thermal models have been used where the

ETPs are estimated based on historical averages (e.g., Scott et al. (2011); Lu et al. (2010)). All these studies assume a *fixed* model over an arbitrary horizon which is ineffective for domestic environments which can be highly dynamic, both in terms of external affect (e.g., weather or adjacent buildings effects) and occupant activity (cooking, opening a window, or operating an auxiliary heating device), as discussed in Section 1.1.2. Thinking in the limit, an extremely complex *fixed* thermal model can potentially consider all the highly dynamic thermal effects of a house. However, monitoring and considering all these effects is not feasible in realistic settings and so such approaches will not be used in this research.

- **Adaptive gray-box thermal modeling:** In contrast to *fixed* gray-box thermal modeling, in *adaptive* gray-box modeling, the ETPs are learned on-line and are assumed to be time-varying. In general, adaptive thermal modeling has been shown to be resilient and effective in highly dynamic settings. This is because the partial observation of the underlying thermal dynamics can be interpreted as non-stationarity of the assumed thermal dynamics (i.e., the derived gray-box model). In this context, several methods have been utilized for on-line estimation of time-varying ETPs such as Kalman filters, recursive least squares methods, and genetic algorithms (e.g., Fux et al. (2014); Radecki and Hencey (2012, 2013); Coley and Penman (1992); O'Neill et al. (2010); Li and Wen (2014)). For example, Fux et al. (2014) applied extended Kalman filters (EKFs) to a test building in Swiss Alpine Club, Switzerland, for adaptive thermal modeling. Their EKF models the indoor temperature, leakage rates and solar radiation and adapts to the seasonal change in heat flow introduced by the occupants. Similarly, Radecki and Hencey (2012, 2013), and O'Neill et al. (2010) use EKFs for adaptive modeling to demonstrate adequate predictions of thermal characteristics of buildings. Furthermore, Coley and Penman (1992) describe a recursive least square algorithm to identify, in real time, parameters that characterize the thermal response of a building. In general, the particular parameter estimation technique is subject to the complexity of the thermal model employed which, in turn, depends on the heating system to be controlled. Due to the above discussed reasons, in this work we employ adaptive gray-box thermal modeling.

2.4 Predicting the Occupancy Schedule

In addition to an appropriate thermal model, an additional component that is key in modeling the dynamics in the context of heating automation systems is the occupancy schedule. This is the case since any thermal comfort, or discomfort, is experienced only when the space is occupied. However, as discussed in Section 1.1.2, this schedule is usually unknown in domestic settings and needs to be predicted. To this end, several approaches to predict occupancy have been proposed over time (e.g., Krumm and Brush

(2011); Scott et al. (2011); Lu et al. (2010); Gupta et al. (2009); Krumm and Horvitz (2006); Scellato et al. (2011); Ye et al. (2009)).

Following the classification of Kleimingera et al. (2013), occupancy prediction approaches can be classified as either *schedule-based* or *context-aware*. The former predict occupancy based only on the history of the occupancy schedule (e.g., Krumm and Brush (2011); Scott et al. (2011); Lu et al. (2010)). As such, these approaches require limited infrastructures (i.e., occupancy sensors) and limited computation power. Nevertheless, these have demonstrated high prediction accuracy (Kleimingera et al. (2013)). In contrast, context-aware approaches predict the occupancy schedule based on observations of the current context of the occupants (potentially combined with the occupancy schedule history) (e.g., Gupta et al. (2009); Krumm and Horvitz (2006); Scellato et al. (2011); Ye et al. (2009)). Such approaches have the potential to demonstrate higher predictive accuracy over the schedule-based ones as additional information is being considered. However, they typically require additional instrumentation and computational power. For a review on state-of-the-art occupancy prediction approaches, along with a comprehensive evaluation, see Kleimingera et al. (2013).

In this work, we employ the schedule-based occupancy prediction approach proposed by Scott et al. (2011). This has illustrated median predictive accuracies of $\sim 80\%$ high performance in comparison to other schedule-based occupancy prediction approaches (Scott et al. (2011); Kleimingera et al. (2013)). We have chosen this approach due to the general low instrumentation needs of schedule-based approaches, and its particular efficiency. However, irrespective of the particular choice, all proposed approaches to predict the occupancy schedule inevitably retain an uncertainty over the predicted occupancy schedule which, appropriately, is modeled in the form of probabilistic occupancy estimates. As such, dealing with this uncertainty arises as a prominent challenge in heating automation systems (e.g., Mozer et al. (1997); Urieli and Stone (2013); Lu et al. (2010); Gao and Keshav (2013a); Erickson and Cerpa (2010); Dong et al. (2011)).

2.5 Control Approaches of Heating Automation Systems

As discussed in Section 1.1.1, the goal of any heating automation system is to control the heating in order to balance heating energy consumption (or heating cost) and the occupant's thermal discomfort according to their preferences, with minimum user-input. In this context, several control approaches have been investigated (for a respective review see Wang and Ma (2008) and Dounis and Caraiscos (2009)). However, in general, these approaches can be classified as either *model-based* or *model-free*, based on their reliance on a model that describes the dynamics of the system under control.[11]

[11]We note that, as far as the classification provided in Wang and Ma (2008) is concerned, here *hybrid* and *performance map-based supervisory control* approaches fall in *model-based* approaches as they also *rely* on a model of the system.

In more detail, model-free approaches aim to optimize the heating control process without a descriptive model of the system dynamics. Such control approaches range from complex expert systems and reinforcement learning approaches to simple reactive control techniques (e.g., Dalamagkidis et al. (2007); Liu and Henze (2006); Braun and Diderrich (1990); Doukas et al. (2007)). However, they typically suffer from several drawbacks, such as reliance on explicit and extensive knowledge, high computational complexity, slow convergence rate and/or instability, that typically render them unsuitable for practical purposes (Wang and Ma (2008)). Thus, such approaches will not be followed in this work.

In contrast, model-based control approaches aim to optimize the heating control process based on a model of the system dynamics which enables them to plan a heating schedule ahead of time and predict the outcome of the heating actions. Based on this model, such control approaches calculate a control law either *on-line* or *off-line*. In the former case, an appropriate policy is calculated in advance which generally defines the optimal (or near-optimal) action for any possible system condition (i.e. system state). Then the appropriate action is chosen in a straightforward way based on this mapping (and an identification of the current system state) to control the system. However, an off-line computation of an effective control law is difficult or impossible when the state space is extremely large (as is generally the case in heating automation systems), and hence, such approaches generally fail to account for essential details and/or have high computational complexity that renders them impractical in realistic settings. For instance, the work in Shann and Seuken (2014), that is based on an off-line calculation of a control law, fails to consider real-time updates of the occupancy schedule and weather condition estimates, which is crucial for the control efficiency of any heating automation system, and especially of domestic ones (as discussed in Section 1.1.1).

On the other hand, one particular family of model-based control approaches that has proven very efficient in heating automation systems, and is also employed in this work, is that of model predictive control (MPC) (e.g., Oldewurtel et al. (2010b); Široký et al. (2011); Mozer et al. (1997); Gao and Keshav (2013a); Urieli and Stone (2013); Moroşan et al. (2010); Freire et al. (2008); Oldewurtel et al. (2010a)). Along with its demonstrated effectiveness, the significant success of MPC in this context is due to its particular ability to handle control problems where an off-line computation of a control law is difficult or impossible. In the following section (Section 2.6) we provide a detailed discussion of model predictive control.

2.6 Model Predictive Control

Several control approaches have been investigated in the context of heating automation systems (as discussed in Section 2.5). However, one particular family of approaches that has been proven very efficient and has been widely used in heating automation systems is that of model predictive control, MPC, (e.g., Oldewurtel et al. (2010b); Široký et al. (2011); Mozer et al. (1997); Gao and Keshav (2013a); Urieli and Stone (2013); Moroşan et al. (2010); Freire et al. (2008); Oldewurtel et al. (2010a)). In addition to heating automation systems, MPC is widely used in many other industrial applications including also low-level heating control[12] (e.g., Qin and Badgwell (2003); Camacho and Bordons (1995); Prívara et al. (2011)). The considerable industrial success of MPC is due to its ability to handle control problems where an off-line computation of a control law is difficult or impossible, along with its ability to handle hard control constraints in a straightforward manner (Mayne et al. (2000); Qin and Badgwell (2003)). In particular, these constraints can set conditions either on *output* variables, i.e., system variables that are influenced by the control actions executed, or *input* variables, i.e., the control actions itself. Furthermore, they are "hard" in the sense that are *required* to be satisfied. For these reasons, in this work we employ an MPC approach for our domestic heating automation system.

In more detail, MPC refers to a wide family of on-line control algorithms that, more or less, share the following criteria (Camacho and Alba (2013)): (i) they make explicit use of a model that describes the dynamics of the system under control in order to predict its future state; (ii) based on this model, they calculate a sequence of actions over a finite horizon according to the optimization objective—here we refer to this process as *planing*; and, finally, (iii) they apply the first control action of the calculated sequence, and repeat the procedure, shifting the planning horizon into the future—a property known as receding horizon.[13] As such, MPC approaches require a descriptive model of the system, or models of the system components, and an appropriate optimization method for planning; i.e., a method to derive the optimal heating schedule based on the system modeling.

Now, depending on the complexity of the system model used (linear or non-linear), the complexity of the optimization objective (e.g., linear or non-linear objective function; quadratic objective function; subject to linear or non-linear constraints; subject to quadratic constraints), and the complexity of the available control actions (continuous or discrete values, subject to constraints), different optimization problems arise in planning (e.g., quadratic programming problems, linear programming problems and mixed integer programming problems). As such, considering also the computational resources available, different optimization methods are employed, such as interior-point methods,

[12]Here, low-level heating control refers to the control method employed for maintaining the inside air temperature close to a particular value (i.e., the set-point temperature) with the minimum of oscillations. To this end, the control approaches typically utilized range from simple ON/OFF control to different variations of proportional-integral-derivative (PID) control, and MPC, depending on whether a static or programmable thermostat is employed (most commonly, systems with static thermostats employ simple ON/OFF control) (Wang and Ma (2008)).

[13]Due to this last property model predictive control is also well known as receding horizon control.

simplex algorithms and dynamic programming (e.g, Wang and Boyd (2010); Patrinos et al. (2011); Lee (2011); Camacho and Alba (2013)).

However, despite its considerable success, MPC is essentially a suboptimal controller. The receding horizon technique allows us to tackle difficult control problems, but, generally, this comes with a cost that is mainly reflected in the theoretical properties of: (i) planning feasibility, (ii) control stability, and (iii) performance. In particular, minimizing an objective function which is subject to hard constraints on the output variables over a finite receding prediction horizon, might drive the system outside the feasible region where the constraints can be satisfied. This fact can lead to a non-feasible optimization problem. This can happen due to a disturbance effect or due to the receding horizon optimization procedure inadvertently driving the system outside the feasible region (Scokaert and Rawlings (1999)). However, by definition, feasibility issues in meeting *soft* constraints (constraints which are preferred but not required to be satisfied) cannot emerge. Moreover, feasibility issues in meeting hard (or soft) *input* constraints cannot emerge either. In particular, at any time instance the population of these variables is independent of the previous actions executed, entirely left to the controller and, hence, can be strictly enforced to meet the conditions. As such, feasibility is generally an issue only when the objective function is subject to hard constraints on the output variables (Scokaert and Rawlings (1999)). Nevertheless, this is *not* the case in our approach (as further discussed in Chapter 3). In addition to these feasibility issues, the recursive horizon optimization technique might lead to unstable control performance, especially in non-stable systems with fast dynamics. However, stability is not an issue in typically stable systems with generally slow-dynamics such as buildings (Širokỳ et al. (2011)). Lastly, the sequence of actions that is actually executed in the context of MPC, might differ significantly from the sequence of actions that is calculated in planning at each instance. As such, planning optimization might have only a tentative connection with the optimization of the underlying real process. In this aspect, MPC is not an optimal control approach and its performance might deteriorate compared to an an optimal one (which however might have a computational complexity that renders it impractical in realistic settings, as discussed in Section 2.5).

In general, the theoretical properties of the above concepts are subject to the particular MPC design (e.g, length of the predictive horizon, or rate of control action execution) along with the complexity of: (i) the system model, (ii) the optimization objective, (iii) the control actions and (iv) the underlying system dynamics. Addressing these theoretical questions, in all cases, is an active area of research in the context of MPC (Morari and Lee (1999); Mayne et al. (2000); Camacho and Alba (2013); Mayne and Rawlings (2001); Scokaert and Rawlings (1999)). In this context, many MPC variations have been proposed to address the above issues in different complexity settings.

In particular, typically MPC refers to the commonly employed *certainty equivalence* MPC version (Bertsekas (2005)). As implied by its name, this MPC version plans using the expected values of all predicted variables of the system dynamics *as certain*; ignoring any possible uncertainty (appropriately modeled in the form of deviation). In general, this loss of information is countered by the on-line nature of MPC, as the actual values of the predicted variables are used at every iteration as the initial state in the MPC planning. However, this fact posses challenges for the theoretical guaranties of MPC regarding planning feasibility and control stability, and also affects its performance. To this end, robust MPC and stochastic MPC approaches have been proposed that handle the model uncertainties in a more concrete manner (compared to certainty equivalence MPC) (e.g, Bemporad and Morari (1999); Kothare et al. (1996); Lee and Kouvaritakis (2000); Oldewurtel et al. (2010b); Couchman et al. (2007)). These approaches can provide robustness and stability guaranties under certain conditions. However, typically these guaranties come with a trade-off in performance and/or the proposed MPC approaches are computationally expensive (Morari and Lee (1999); Camacho and Alba (2013)). Another variation of MPC that handles model uncertainties in a concrete manner and has the potential to improve the controller performance is *adaptive* MPC (e.g., Fukushima et al. (2007); Lee and Ricker (1994)). In general, adaptive MPC differs from the above defined typical MPC procedure in that the system model is updated at each iteration in order to account for any variability in the system characteristics. In general, adaptive MPC is an attractive way to further handle model uncertainties. As such, in this work we employ an adaptive MPC approach. Moreover, once again, the absence of hard output constrains in our optimization objective and the generally slow and stable dynamics of buildings, ensure that we will not face any feasibility or stability issues, which are typically further introduced particularly by the adaptation mechanism in the context of adaptive MPC (Fukushima et al. (2007)).

2.7 Heating Automation Systems that deal with Occupancy Uncertainty

Many heating automation systems have been proposed over time for both domestic and non-domestic settings (e.g., Farris and Melsa (1978); Kintner-Meyer and Emery (1995); Kummert et al. (2001); Henze et al. (2004); Liu and Henze (2006); Mozer et al. (1997); Scott et al. (2011); Lu et al. (2010); Urieli and Stone (2013)). However, dealing with occupancy uncertainty is a necessary requirement for a heating automation system to be applicable in domestic settings (as discussed in Section 1.1.2). As such, in this review (see Table 2.1) we focus only on systems that deal with this uncertainty and hence could potentially be employed in domestic settings.

In more detail, there has been an increasing amount of work on heating automation systems that deal with occupancy uncertainty. In particular, a number of domestic heating automation systems (DHASs) have been proposed (e.g., Mozer et al. (1997); Urieli and Stone (2013); Lu et al. (2010)), along with some advanced *non-domestic*

TABLE 2.1: Heating automation systems overview

Automation System	**Limitation**			
	Threshold Occupancy Probabilities	Separate Preheating & Heating Stopping	Fixed Thermal Model	Simple Heating Control
Neurothermostat	×	×	✓	✓
Smart Thermostat	×	✓	✓	✓
PreHeat	✓	✓	✓	✓
SPOT+	✓	×	✓	✓

(mainly for offices) heating automation systems (e.g, Gao and Keshav (2013a); Erickson and Cerpa (2010); Erickson et al. (2013)). We will now discuss the main relevant systems.

2.7.1 Neurothermostat

In their pioneering work on probabilistic occupancy in heating automation systems, Mozer et al. (1997) propose the DHAS *Neurothermostat*. Neurothermostat employs a general heating control method that fully exploits the occupancy probabilities and balances cost and discomfort via a single objective in an infinite horizon optimization manner. In more detail, the expected discomfort is expressed in monetary cost, through a simple static empirical formula, and is added to the cost of heating. In this context, Neurothermostat aims to minimize this unifying cost over an infinite horizon.

To this end, the following MPC approach is utilized: At every iteration an exhaustive search over all possible heating schedules is performed, looking for a schedule that minimizes the unifying cost over a finite planning horizon. Then, the first action of the derived schedule is executed and the process is repeated shifting the horizon into the future. However, the major drawback of this work is that it employs exhaustive search (for planning a heating schedule) which is extremely costly, limiting its applicability to simple proof-of-concept settings. In this context, Neurothermostat relies on a simple, fixed and, thus, impractical thermal model. Moreover, Neurothermostat employs a static empirical formula to express discomfort in monetary cost which is problematic as this equivalence varies among users and through time (Scott et al. (2011)). Finally, as outlined in Section 2.1.2, balancing heating *consumption* and thermal discomfort in the absence of energy cost variability is essentially equivalent to balancing heating *cost* and thermal discomfort. As such, although heating cost is used in Neurothermostat, none of the essential aspects of simple or advanced economic control are considered (i.e, variable electricity import tariffs and/or heating-system-integrated IERs, see Section 2.1.2). In this aspect, Neurothermostat considers *simple* (automated) heating control (as defined in Section 1.1.1).

2.7.2 Smart Thermostat

In contrast to the infinite horizon approach above, Lu et al. (2010) presented the DHAS *Smart Thermostat*. Smart Thermostat divides heating control into two relatively independent problems: (i) when to switch off heating; *heating stopping* and (ii) when to switch it on; *preheating*. The former is tackled in a reactive manner using sensors and a hidden Markov model (Dugad and Desai (1996)) to infer sleep and departure events. Hence, when the space is not occupied the heating system is switched off and the inside temperature is allowed to sink down to a "deep" setback temperature. The latter is tackled utilizing a heating-system-specific heuristic approach. In particular, heating is switched on at a particular time instance which is chosen so as to minimize the long-term expected energy waste for a variable efficiency three-stage heating system (a two-stage heat pump and a third stage electric heater) given predictions of the occupant arrival events. To this end, a simple thermal model is used based on equivalent thermal parameters, ETPs, estimated as historical averages. Now, in order to reduce the time required to recover to the set-point temperature if an occupant returns before the space is adequately heated, the space is also heated to a "shallow" setback temperature before the first possible arrival. However, the reactive approach utilized for heating stopping is not Pareto optimal for heating systems that exhibit considerable thermal lags. In particular, *early stopping*, i.e., switching off heating some time before a departure, can reduce the heating energy consumption in such settings without any comfort loss (Ellis et al. (2012)). In addition, the preheating approach utilized is extremely system specific and only searches over a sub-region of the heating schedule space. In particular, it only searches for a specific time instance to switch on heating. In this context, more complex heating policies, which could possibly be more efficient, are not considered. Most importantly though, tackling preheat and heating stopping independently is not effective in heating systems that exhibit considerable thermal lags, such as underfloor heating systems (even if early stopping is considered). In such systems, the preheating strategy can significantly affect the optimal heating stopping strategy and vice-versa. Finally, Smart Thermostat does not consider any economical aspects (other than the energy consumed) and, as such, considers also *simple* (automated) heating control.

2.7.3 PreHeat

In contrast to the above works, more recently, a particular heuristic approach is rising in popularity which deals with the probabilistic occupancy estimates in a thresholding manner (e.g. Scott et al. (2011); Gao and Keshav (2013b,a)). In particular, the probabilistic occupancy estimates are assumed binary depending on their relation to a predefined threshold; any estimate above the threshold is assumed 1, otherwise it is assumed 0. Most well known representatives of this approach consider the work of Scott et al. (2011) and Gao and Keshav (2013a) (which also consider the benchmark of our evaluation procedure)[14].

In more detail, Scott et al. (2011) propose the DHAS *PreHeat*, which plans based on a deterministic occupancy schedule, derived through the aforementioned thresholding approach. PreHeat system works in two ways:

1. When the space is considered occupied it uses predefined set point temperatures.
2. When the space is considered non-occupied, it uses a lookahead window to check if an occupancy event is imminent (according to the derived deterministic occupancy schedule) in order to: (i) heat up the space for the minimum time required right before an occupancy event, and (ii) ensure that the set-point temperature is met at the time instance of the occupancy event.

To this end (and, in particular, in order to decide when to start heating to meet the aforementioned objectives), PreHeat uses a simple fixed thermal model based on a single ETP. Specificaly, it uses the heat-rate, which is the warming rate of the house when heating is on. This rate is estimated during an initial deployment phase as a simple historical average value and is considered to be fixed thereafter. However, in the context of the derived deterministic occupancy schedules, PreHeat tackles preheating and heating stopping independently and employs reactive heating stopping. As such, it faces the respective, aforementioned limitations. Moreover, although the preheating approach employed aims to eliminate discomfort with the minimum heating cost, this is only achieved for heating systems that do not exhibit any thermal lags or heating cost variability over time. In all other cases, this trivial preheating strategy is not optimal (as illustrated in Section 2.1). On top of the above, the actual trade-off between heating cost and discomfort is determined by the threshold choice which defines the deterministic occupancy schedule in the first place. As such, cost and discomfort are balanced based on a heuristic approach and no warranties or intuition is given regarding the performance loss from optimal heating schedule planning that, ideally, would fully exploit the probabilistic occupancy estimates. Finally, PreHeat does not consider any economical aspects (other than the energy consumed) and, as such, considers *simple* (automated) heating control.

2.7.4 SPOT+

Another example of work with the thresholding technique is SPOT+ (Gao and Keshav (2013a)), a non-domestic heating automation system for office buildings. SPOT+ tackles heating control in an infinite horizon optimization manner but plans based on a

[14] We note here that both Neurothermostat and Smart Thermostat, described above, are not suitable to serve as benchmarks for our evaluation procedure as the computational complexity of the former is restrictive for our case study system and the latter is extremely system specific.

deterministic occupancy schedule. This schedule is derived through the aforementioned thresholding approach. To this end, SPOT+ uses a parameterizable unifying formula to balance heating consumption and discomfort, in a single objective optimization manner and, appropriately, uses MPC and a graph theoretical approach (shortest path finding) for planning (in order to find the heating schedule that minimizes the unifying cost over the predicting horizon). In this context, SPOT+ uses a fixed thermal model for planning, estimated through least squares regression. Specifically, SPOT+ uses the following simple thermal model for planning:

$$T_{t+1}^{IN} = T_t^{IN} + \frac{eP_t^{hvac} - k(T_t^{IN} - T_t^{OUT})}{C} \tag{2.1}$$

where P_t^{hvac} is the power of the heating system at time t, and T^{OUT} is the outside temperature (T^{IN} stands for the indoor temperature). Moreover, e, k, and C stand for the efficiency, the conduction factor, and the building heat capacity, respectively, and consider the ETPs which are estimated through least squares regression during an initial training phase, and assumed fixed thereafter.

The formula used in SPOT+ for planning, weights heating consumption and thermal discomfort according to a user specified weighting parameter. In more detail, the formula is:

$$J(\cdot) = \sum_{\tau=1}^{|H|} \mathrm{Cost}(\cdot_\tau) + O_\tau \lambda' \,\mathrm{Disc}(\cdot_\tau) \tag{2.2}$$

where $\lambda' \in [0, \infty)$ stands for the *user-provided* weighting parameter, $|H|$ stands for the number of intervals, τ, within the planning horizon, and $O_\tau \in \{0, 1\}$ indicates whether the space is occupied during interval τ based on the derived deterministic occupancy schedule. $\mathrm{Cost}(\cdot_\tau)$ and $\mathrm{Disc}(\cdot_\tau)$ are functions that return the heating energy consumption and the thermal discomfort, respectively, according to the consumption model and discomfort metric used. In particular, SPOT+ simply uses P_t^{hvac} to model heating consumption and estimates discomfort according to a non-linear discomfort metric based on the 7-point ASHRAE scale.[15] By doing so, the proposed discomfort metric is a variation of the predicted mean vote model (Fanger et al. (1970)), extended with an affine transformation method to account for spaces occupied by a single person, and a dead-band/neutral zone (for more details on this thermal comfort metric see Gao and Keshav (2013a)).

In this context, SPOT+ balances discomfort and heating consumption on two levels: (i) based on the threshold choice to derive with the deterministic occupancy schedule and, (ii) based on the (semi-bounded, i.e. $\in [0, \infty)$) weighting parameter used in the unifying formula. However, this balancing scheme obscures how each one of the balancing techniques affect the trade-off between consumption and discomfort making parameter choice

[15]The 7-point ASHRAE scale is {cold (-3) , cool (-2), slightly cool (-1), neutral (0), slightly warm (+1), warm (+2), and hot (+3)}

tricky. Moreover, this scheme also considers a heuristic approach and no warranties or intuition is given regarding the performance loss from optimal heating schedule planning. Lastly, although shortest path finding is mentioned for planning, the algorithmic choice is not reported nor is an appropriate algorithm provided. Finally, it should be noted that SPOT+ does not consider any economical aspects (other than the energy consumed) and, as such, also considers *simple* (automated) heating control.

2.8 Heating Automation Systems that Incorporate Economic Control

In the above section (Section 2.7) we reviewed the literature of heating automation systems that deal with occupancy uncertainty (and, hence, could potentially be employed in domestic settings). However, none of the systems considered in our review deal with (simple or advanced) economic control. Indeed, the domestic heating automation systems, DHASs, that consider economic control are currently highly experimental and none of them deal with occupancy uncertainty. In addition, as discussed in Section 1.1.4, all DHASs that consider economic control account only for *simple* economic control, while none of them considers *advanced* economic control (i.e., taking into account also domestic/heating-system-integrated IERs). Nevertheless, in the following paragraphs we provide a review on the state-of-the-art of DHASs with respect to economic control in terms of concreteness.

In more detail, in recent years, economic control is starting to emerge as an integrated part of DHASs (e.g., Rogers et al. (2011); Halvgaard et al. (2012); Shann and Seuken (2014)). In particular, the work of Rogers et al. (2011) proposes an adaptive heating algorithm that considers time-varying electricity import tariffs in balancing heating cost and thermal discomfort. In this context, it first predicts the external temperature using Gaussian process (GP) regression (for more details on GP regression see Section 2.10). Then, it computes a heating schedule using mixed-integer programming and a simple fixed thermal model. However, this work is highly experimental, and assumes that: (i) the weather predictions are absolutely accurate, (ii) the occupancy schedule is perfectly known in advance, and (iii) the simple fixed thermal model utilized perfectly describes the thermal dynamics of the house. However, all the above are somewhat unrealistic assumptions in practical settings. Nevertheless, this work illustrates the potential of DHAS-integrated simple economic control and considers the motivational work for many DHAS in this respect (including ours).

Following a similar line of research, the work of Halvgaard et al. (2012) proposes an MPC-based approach for DHAS-integrated simple economic control to account for time-varying electricity import tariffs. In this work the electricity prices are incorporated in the MPC planning optimization objective as cost coefficients. To this end, this work utilizes a white-box thermal modeling approach. Nevertheless, this work assumes a perfect white-box thermal model of the thermal dynamics without taking into account the occupants' activity. However, the occupants' activity can affect the thermal dynamics in domestic settings (as discussed in Section 1.1.2) and, hence, a thermal model that does not consider these aspects cannot be assumed to be completely accurate. Moreover, in general, white-box thermal modeling requires extensive building information that is not always available in domestic settings and/or it is expensive to obtain (as discussed in Section 2.2). Last but not least, this work also assumes that the weather predictions utilized are absolutely correct and that the occupancy schedule is perfectly known in advance (which is not the case in realistic settings, as noted above).

Finally, the work of Shann and Seuken (2014) proposes a DHAS that accounts for time-varying electricity import tariffs, aiming to overcome some of the limitations of the aforementioned approaches. In more detail, in this work, the outside temperature and the electricity import tariffs are predicted using GP regression. Then this particular stochastic modeling is utilized in calculating an off-line control law to optimize the heating process. As such, this work aims to account for the weather prediction uncertainty, and, also, incorporates an approach to predict the electricity import tariffs (accounting also for the respective uncertainty).[16] However, (as already noted in Section 2.5), this work fails to account for real-time updates of the occupancy schedule and/or the weather forecasting reports, which is crucial for the control efficiency of DHASs (as discussed in Section 1.1.1). In addition, in order to further reduce the respective optimization complexity, this work utilizes a simple fixed thermal model assuming that it perfectly describes the thermal dynamics of the house. However, as noted above, this is an unrealistic assumption in practical settings.

2.9 Predicting the Power Output of Intermittent Energy Resources

As discussed in Section 1.3.1, against the background detailed in Section 2.8, in this work we propose a DHAS that incorporates both simple and advanced economic control. Now, an essential part of our DHAS-integrated advanced economic control approach is a stochastic, short term (up to 12 hours ahead), prediction of the shared IER power output (as discussed in Section 1.3.1). As discussed in Section 1.2, the power output of IERs depends on the prevailing weather conditions. In particular, the power output of a PVS depends mostly on the irradiance incident to the photovoltaic module and the operating temperature of the module (Luque and Hegedus (2011)), while the power output of a WTG depends mostly on the prevailing wind speed and wind temperature (Burton et al. (2011)). In addition, advanced IER models exist to transform these

[16]Predicting the electricity import tariffs is crucial in settings with real-time pricing which is unknown in advance, for more details see Torriti et al. (2010).

weather variables into IER power output with high accuracy (Panagopoulos et al. (2012)). As such, predicting the power output of IERs is tightly linked with the task of predicting the main environmental variables that affect their operation.

Broadly speaking, existing approaches for IER power output prediction fall into three main categories (Soman et al. (2010); Paulescu et al. (2012)): (i) approaches based on physical numerical weather prediction (NWP) models; (ii) statistical (or machine learning) approaches; and (iii) hybrid approaches that essentially combine the two former ones. In more detail, NWP-based approaches (e.g., Soman et al. (2010); Chen et al. (2014)) consider the analysis of numerous weather parameters via complex models of the dynamics governing the motion of the fluids in the atmosphere to derive the weather variable predictions. These predictions are then transformed into power output predictions using models of the actual generator. In general, such approaches are limited by the complexity of their advanced meteorological analyses, as well as the absence of historical observations for their site-specific calibration (Chen et al. (2014)). Nevertheless, many on-line providers serve detailed weather forecasting reports based on such NWP models, free of charge, that can be used to facilitate NWP-based approaches in a straightforward manner.[17] In contrast, statistical approaches, use only historical readings of the IER power output, without including any physical model, to predict future IER power output. In this context, several statistical methods have been proposed, providing either stochastic or non-stochastic predictions, including autoregressive models (Brown et al. (1984)), autoregressive–moving average models (Kamal and Jafri (1997)), artificial neural networks (Li and Shi (2010)), support vector machines (Mohandes et al. (2004)), and Bayesian methods such as Kalman filters and Gaussian processes (Jiang et al. (2010); Chen et al. (2014)). Such approaches provide accurate very short term predictions (up to 4 hours), by learning correlated signals from the historical observations, but their mid-to-long term accuracy is generally limited compared to NWP-based approaches (Soman et al. (2010); Paulescu et al. (2012)). Last, hybrid approaches combine the aforementioned statistical and NWP-based approaches, aiming to overcome the respective limitations (Soman et al. (2010)). As such, they are also employed in this work.

2.10 Gaussian Process Regression

Gaussian process (GP) regression is a well-known state-of-the-art approach for handling non-linear regression problems that has been widely and effectively used for hybrid IER power output prediction (e.g., Chen et al. (2014); Zamo et al. (2014)). GP regression also provides stochastic predictions in a principled manner and, hence, it is also employed in

[17] That said, the accuracy of such reports generally deteriorates compared to commercial ones, while they are also limited in the environmental variables that they consider (for instance, solar radiation, which is essential for photovoltaic system power output prediction, is usually absent in free-of-charge reports), making the need for site-specific calibration more apparent.

this work. Formally, a GP can be defined as a stochastic process considering a distribution over functions $f : X \to Y \in \mathbb{R}$ such that any finite subset $X' \subset X$ is multivariate Gaussian distributed. Specifically, a GP is completely defined by a mean function, $m(.)$ and a positive semi-definite covariance function (or kernel), $K(.,.)$. If the observed values, $\mathbf{y} = y_1 \dots y_n$, and a corresponding set on inputs $\mathbf{x} = x_1 \dots x_n$, are thought of as a noisy version of the true underlying function f, i.e., $\mathbf{y} = f(x) + \epsilon$, where ϵ is assumed to be zero mean Gaussian noise, $\mathcal{N}(0, \sigma^2)$, then a GP prior distribution on f can be defined as $f \sim \mathcal{N}\left(m(\mathbf{x}|\theta), K(\mathbf{x}, \mathbf{x}|\theta) + \sigma^2 I\right)$, where θ is the set of hyper-parameters that characterize the shape of the mean and covariance functions. Given the Gaussian likelihood of the GP, at any set of test points $\mathbf{x}_*$, the predictive posterior distribution at these points can be evaluated in closed form as a multivariate Gaussian distribution:

$$p(\mathbf{y}_*|\mathbf{x}_*, \mathbf{x}, \mathbf{y}, \theta) = \mathcal{N}\left(\mathbf{y}_*; m(\mathbf{x}_*|\mathbf{y}, \mathbf{x}, \theta), \Sigma(\mathbf{x}_*|\mathbf{y}, \mathbf{x}, \theta)\right) \tag{2.3}$$
$$m(\mathbf{x}_*|\mathbf{y}, \mathbf{x}, \theta) = m(\mathbf{x}_*) + K(\mathbf{x}_*, \mathbf{x}) K_y^{-1}(\mathbf{y} - m(\mathbf{x})) \tag{2.4}$$
$$\Sigma(\mathbf{x}_*|\mathbf{y}, \mathbf{x}, \theta) = K(\mathbf{x}_*, \mathbf{x}_*) - K(\mathbf{x}_*, \mathbf{x})^T K_y^{-1} K(\mathbf{x}_*, \mathbf{x}) \tag{2.5}$$
$$K_y \triangleq K(\mathbf{x}, \mathbf{x}) + \sigma^2 I \tag{2.6}$$

The hyper-parameters, θ, are typically learned by maximizing the log marginal likelihood:

$$\theta_* = \arg\max_{\theta} \left(\ln p(\mathbf{y}|\mathbf{x}, \theta) = -\frac{1}{2} \ln |K(\mathbf{x}, \mathbf{x})| - \frac{1}{2} \mathbf{y}^T K(\mathbf{x}, \mathbf{x})^{-1} \mathbf{y}^T - \frac{n}{2} \ln(2\pi) \right)$$

where n is the number of input points. The maximizer of this function can be searched using a standard gradient-based optimization method such as conjugate gradient or Newton methods—for more details see Rasmussen and Williams (2005).

2.11 Coalition Formation and Energy Systems

As discussed in Section 1.3.1, in this work we propose a DHAS that is able to exploit the coalition potential that arises in advanced economic control. Although, this is the first work to exploit this potential in DHAS-integrated advanced economic control, coalition formation is a considerably active research area within the energy sustainability agenda (e.g., Alam et al. (2015); Chalkiadakis et al. (2011); Pudjianto et al. (2007); Ramchurn et al. (2013)). In particular, the formation of grid entities' coalitions has long been proposed as a means to provide regulation services to the grid (Goswami and Kreith (2015)). More recently, it has also gained ground as a means for achieving the cost-efficient and reliable integration of the many distributed energy resources that are starting to emerge in the grid (Ramchurn et al. (2012)). In this context, dynamic coalition formation has been heralded as a key aspect of the next generation electrical grid (Ramchurn et al. (2012)).

In more detail, and regarding the energy consuming entities of the grid, respective coalition formation has been proposed to encourage less dynamic and more predictable energy consumption profiles (Ramchurn et al. (2013)) and/or to offer regulation services to the grid via appropriate demand response (Goswami and Kreith (2015); Kota et al. (2012)). Regarding the energy producers of the grid, the formation of virtual power plants has been proposed as a means to facilitate the reliable integration of IERs into the grid. In particular, virtual power plants correspond to the notion of a large number of heterogeneous distributed energy resources, usually IERs, joining forces in order to offer electricity to the grid as an aggregate—while providing the guarantees of a single "conventional" power plant (Ramchurn et al. (2012); Pudjianto et al. (2007)). As such, virtual power plants consider coalitions that create the necessary synergies among distributed energy resources, so that the effective and efficient delivery of energy is assured, while still being able to utilize (the inherently intermittent and thus untrustworthy) IERs (Ramchurn et al. (2012)). Notably, the term virtual power plants may also refer to coalitions of prosumers, i.e, grid entities that both consume and produce energy (Ramchurn et al. (2012)), or heterogeneous coalitions (i.e., consisting of consumers, producers and/or prosumers) that come together to maximize their economic benefit (Giuntoli and Poli (2013)).

Now, with respect to thermostatically controlled loads (including domestic space heating loads), several works have discussed the potential of respective aggregations for providing regulation services to the grid (e.g., Hao et al. (2013); Callaway (2009); de Nijs et al. (2015)). Nevertheless, non of these works has accounted for this potential in the context of DHAS-integrated (advanced) economic control. In addition, as discussed in Section 1.1.4, works that deal with domestic space heating system aggregations typically assume the same preferences in balancing heating cost and thermal discomfort among the households (e.g., Dudley and Piette (2008); Torriti et al. (2010)) which is impractical in realistic settings (as discussed in Section 1.1.2). In general though, several challenges arise in the formation and management of the aforementioned coalitions with respect to the members being required to come to an agreement in a wide range of economic and/or technical aspects, such as on the allocation of the collective gains, on the plan for further investments, and/or on specific contracts and collective agreements with external costumers or utility companies (Saad et al. (2012)).

2.12 Cooperative Game Theory

As discussed in Section 2.11, several challenges arise in the formation and management of coalitions in the energy systems. To this end, cooperative game theory provides the theoretical framework to tackle some of these challenges. In more detail, game theory generally comprises two main branches (Osborne and Rubinstein (1994)): (i) non-cooperative game theory and (ii) cooperative game theory. Now, the former covers the strategic choices resulting from interaction among competing independent players, while the latter focuses on competitive scenarios where players can form cooperative groups (i.e., coalitions) in order to enhance their position in a game (Augier and Teece (2014)). Such a game is called a cooperative game and considers the theoretical framework of applications considering coalition formation. In this context, here we provide a discussion on cooperative games covering the main concepts that emerge in this work.

Formally, a cooperative game is described by a finite set of players N (that considers the grand coalition) and a characteristic function $v : 2^N \to \mathbb{R}$, where $v(\emptyset) = 0$, which describes the collective payoff (or cost) of a possible coalition. In this context, a player chooses to join a coalition based on an estimate of the way the payoff of the coalition is divided among the members. As such, a key challenge in cooperative games is to allocate the payoff according to a particular notion of fairness. In this context, solving a cooperative game considers finding a vector $x \in \mathbb{R}^{\mathbb{N}}$, which represents the payoff allocation to each player, that satisfies one or more predefined properties. Now, such a vector is called a solution concept and some significant predefined properties are (Gibbons (1992)):

1. **Efficiency:** The solution concept exactly splits the collective payoff:

$$\sum_{i \in N} x_i = v(N) \tag{2.7}$$

2. **Individual rationality:** No player is worse off when joining the coalition:

$$x_i \geq v(\{i\}), \forall i \in N \tag{2.8}$$

3. **Coalitional rationality:** No subset of the coalition is worse off when joining the coalition:

$$\sum_{i \in C} x_i \geq v(C) \ \forall C \subseteq N \tag{2.9}$$

4. **Computational ease:** The solution concept can be calculated efficiently (e.g., in linear or polynomial time with respect to $|N|$).

Importantly, a solution concept that is efficient and individually rational is called an imputation (Osborne and Rubinstein (1994)). Most solution concepts consider a subset of the imputation set and based on their properties can be classified in various categories such as the core (Shapley and Shubik (1966)), the stable set (Von Neumann and Morgenstern (2007)), the kernel (Davis and Maschler (1965)), and the nucleolus (Schmeidler (1969)). For instance, the core considers the set of payoff allocations that satisfy efficiency and both individual and coalitional rationality. In this context, the core is the most important solution concept for cooperative games and is analogous to a Nash equilibrium for non-cooperative games (Osborne and Rubinstein (1994)). Nevertheless, the core of a cooperative game might be empty (Osborne and Rubinstein (1994)). Furthermore, the stable set considers the set of imputations that satisfy the following properties:

(i) internal stability, i.e., no payoff vector in the set is dominated by another vector in the set; and (ii) external stability, i.e., all payoff vectors outside the stable set are dominated by at least one vector in the stable set. Importantly, a payoff vector x dominates a payoff vector y if $\exists\ S \neq \emptyset$ where $x_i > y_i, \forall\ i \in S$ and $\sum_{i \in S} x_i \leq v(S)$. However, a stable set may or may not exist (Osborne and Rubinstein (1994)). The particular solution concept utilized in practice depends on the requirements of an application (Gibbons (1992)). Nevertheless, despite the theoretical attractiveness of sets that fulfill such advanced properties their utilization in applied game theory is usually limited due to their computational complexity (that potentially leads to intractability) and/or lack of universal existence (Osborne and Rubinstein (1994)). Given the above, as discussed in Section 1.3.1, in this work we propose a practical cost allocation mechanism to share the realized gains of a coalition within AdaHeat+ that respects efficiency and individual rationality (and, hence, considers an imputation), as well as allocation efficiency. For more details on cooperative game theory see Osborne and Rubinstein (1994).

2.13 Markov Decision Processes

As discussed in Section 1.2, a big part of this work is dedicated to increasing the efficiency of IERs themselves and in particular in developing novel solar tracking, ST, techniques. In this context, formalizing ST as a dynamic programming problem in the context of a Markov decision process (MDP) is a key contribution of our work (as discussed in Section 1.2). In general, MDPs, named after the Russian mathematician Andrey Markov, provide a mathematical formulation for decision making under uncertainty. In its basic form, a MDP is a discrete time stochastic control process. In particular, at each time step, the process is in a state s, and the decision maker chooses an action a that is available in this state. At the next time step, the process responds by transitioning into a new state s' (according to the action followed and the transition probabilities), and giving the decision maker a corresponding reward.

More formally, a Markov decision process is a 4-tuple $< S, A, P(\cdot|\cdot,\cdot), R.(\cdot,\cdot) >$, where:

- S is the set of states.
- $A(s)$ is the set of actions available at each state.
- $P(s'|s,a)$ is the probability that taking action a in state s will lead to state s'.
- $R_a(s,s')$ is the expected immediate reward by doing action a and transitioning to state s' from state s.

The core problem of a MDP is to find the policy that indicates to the decision maker which action to choose at every state. This policy is a function π linking each state to an action $\pi(s)$. Within this context, the objective is to find a policy π that maximizes a cumulative function of the random rewards.[18] This function typically considers the expected discounted sum of random rewards over the planning horizon, H (i.e., the number of actions that the decision maker can execute):

$$J = \sum_{t=0}^{H} \gamma^t R_{\pi(s_t)}(s_t, s_{t+1}) \tag{2.10}$$

where $\gamma \in [0, 1]$ is the discount factor, that intuitively defines the importance of rewards shorter in time over those coming latter in time. In particular, a MDP can be of finite or infinite planning horizon. Meaning that in the former case H is limited, while in the latter case $H = \infty$. In this context, the discount factor is crucial for solving (i.e., calculating the optimal policy) infinite horizon problems (and is typically set close to 1), while by setting $\gamma = 0$, only the immediate reward is considered (as Equation 2.10 trivially becomes $J = R_{\pi(s_0)}(s_0, s_1)$). In the case of finite horizon MDPs, the discount factor can be set exactly to 1 without any feasibility issues emerging in calculating the optimal policy (Sutton and Barto (1998)).

Now, MDPs can be solved by a wide range of techniques and algorithms typically employing dynamic or linear programming (Sutton and Barto (1998)). That said, dynamic programming algorithms scale better to large problems (such as the one considered in this work, as we further discuss in Chapter 5) (Sutton and Barto (1998); Puterman (2014)). As such, in the following paragraphs we focus on the general family of dynamic programming approaches.

In more detail, generally, dynamic programming refers to the concept of solving a complex problem by breaking it down into a collection of simpler sub-problems (Cormen et al. (2001)). In this context, almost all dynamic programming algorithms for solving MDPs are based on estimating value functions which are used to organize and structure the search for good policies (Sutton and Barto (1998); Puterman (2014)). In their basic form, state value functions estimate how good it is for the process to be in a particular state in terms of future rewards that can be expected (Sutton and Barto (1998)). Since the expected rewards depend on the actions executed, value functions are defined with respect to particular policies; $V^\pi(s)$. Notably, a fundamental property of value functions, that essentially facilitates dynamic programming approaches, is that they can be expressed recursively. In particular, given any policy π, the value of any $s \in S$ can be expressed recursively according to the value of its successor states as:

$$V^\pi(s) = \sum_{s'} P(s'|s, \pi(s)) \left(R_{\pi(s)}(s, s') + \gamma V^\pi(s') \right) \tag{2.11}$$

[18] Because of the Markov property (i.e., the probability distribution of future states depends only on the present state and the action executed), the optimal policy can be written as a function of s alone.

Now, the optimal policy is the one that has an expected return greater or equal to that of any other policy for all $s \in S$. Although there may be more than one optimal policy, they share the same optimal value function, $V^*(s)$, which can be defined by the Bellman optimality equation (Sutton and Barto (1998)):

$$V^*(s) = \max_{a \in A(s)} \left\{ \sum_{s'} P(s'|s,a) \left(R_a(s,s') + \gamma V^*(s') \right) \right\} \tag{2.12}$$

As such, the optimal policy of a MDP, π^* is that corresponding to the Bellman optimality equation above.

In this context, almost all dynamic programming algorithms calculate the optimal value function, and the corresponding optimal policy, by typically executing two essentially independent steps of calculations: (i) policy evaluation, and (ii) policy improvement.

In more detail, in policy evaluation the state-value function of an arbitrary policy π is calculated according to Equation 2.11. In essence, Equation 2.11 is a system of $|S|$ simultaneous linear equations in $|S|$ unknowns (the $V^\pi(s), s \in S$) (Sutton and Barto (1998)). As such, policy evaluation may be formulated and solved as a set of linear equations, or solved iteratively as illustrated in Algorithm 1 (Sutton and Barto (1998)), below.

Algorithm 1 Policy Evaluation

```
procedure POLICY EVALUATION(π)
  Initialize V(s) = 0, for all s ∈ S
  do
    for every s ∈ S do
      v ← V(s)
      V(s) ← Σ_{s'} P(s'|s, π(s)) (R_{π(s)}(s, s') + γV(s'))
      Δ ← max{Δ, v − V(s)}
  while Δ ≥ θ (a very small number)
  return V ≃ V^π
```

The input of the above algorithm is the policy to be evaluated, while the output considers the estimated corresponding value function. Formally, iterative policy evaluation converges only in the limit. As such, in practice it is halted shortly before and, typically, when the quantity $\max_{s \in S}\{V_{k+1}(s) - V_k(s)\}$ becomes sufficiently small (as per the above algorithm), where k indicates iteration (Sutton and Barto (1998)).

In policy improvement an initial policy π is improved according to its value function, V^π, to derive an improved policy, π' (unless the initial policy is already optimal). More formally, the improved policy, π', is calculated as:

$$\pi'(s) = \arg\max_{a \in A(s)} \left\{ \sum_{s'} P(s'|s,a) \left(R_a(s,s') + \gamma V^\pi(s') \right) \right\} \tag{2.13}$$

Now, executing the above two steps in an alternating fashion ensures convergence to the optimal value function and the optimal policy (Sutton and Barto (1998)). Importantly, policy iteration (Howard (1960)) is a dynamic programming algorithm that builds directly on this principle, where the above steps are repeated in an alternating fashion until convergence (i.e., until the policy is stable) as illustrated in Algorithm 2.

Algorithm 2 Policy Iteration

procedure POLICY ITERATION(π)
  Initialize $V(s) \in \mathbb{R}$ and $\pi(s) \in A(s)$, for all $s \in S$
  do
    $V \leftarrow$ Policy Evaluation(π)
    policy-stable $\leftarrow$ True
    for every $s \in S$ **do**
      $b \leftarrow \pi(s)$
      $\pi(s) \leftarrow \arg\max_{a \in A(s)} \left\{ \sum_{s'} P(s'|s,a) \left(R_a(s,s') + \gamma V(s') \right) \right\}$
      if $b \neq \pi(s)$ **then**
        policy-stable $\leftarrow$ False
  while policy-stable = False

As such, policy iteration iteratively improves over an initial policy. In particular, each policy evaluation is carried out over the value function of the previous policy which typically leads to a progressive increase in the speed of convergence (as stated in Sutton and Barto (1998) this is "presumably because the value function changes little from one policy to the next"). Given the above, it becomes clear that the convergence speed of policy iteration highly depends on the initial policy used and how close this is to the optimal.

Finally it should be noted that the policy evaluation step can be truncated in several ways without losing the optimality convergence guarantees. As such, in other dynamic programming approaches the above general steps (i.e., policy evaluation and policy improvement) are being repeated, either independently or combined together, in various orders, to estimate the optimal policy. For instance, in modified policy iteration (Puterman and Shin (1978); Van Nunen (1976)), policy evaluation is repeated several times within each iteration, while in value iteration (Bellman (1957)) policy evaluation is stopped after just one sweep. That said, each dynamic programming algorithm comes with its own limitations and advantages and the choice of the appropriate algorithm highly depends on the characteristics of the MDP to be solved (Sutton and Barto (1998)). Now, due to its particular ability to improve on an initial policy in a straightforward way (as discussed above) in this work we focus on policy iteration. In particular, as discussed in Section 6.4, in this work we propose a new modified policy iteration schema suitable for very large action-state space MDPs (like the one considered in our work).

2.14 Summary

In this chapter we provided relevant background material and an overview of related work regarding DHASs and IERs. In more detail, in Section 2.1 we discussed domestic heating systems with respect to their main characteristics that introduce challenges for efficient modeling and control in the context of heating automation systems. In this context, we showed how variability of the (overall) heating system efficiency, energy cost variability, and thermal lags, can independently affect the efficiency of intuitive and simple heating strategies that are utilized in heating control. By doing so we illustrated the particular challenges that arise in meeting our generality and Pareto efficiency requirements (as discussed in Section 1.1.2). In addition, we discussed the challenges that these characteristics introduce for the reliable thermal and/or consumption modeling of heating systems (as per the reliable modeling requirements in Section 1.1.2).

Subsequently, in Section 2.2 we reviewed the literature of thermal modeling and we provided a general discussion over the pros and cons of different approaches. By doing so, we pointed out the limitations and shortcomings of white-box approaches that render them unsuitable for incorporation in our domestic heating automation system (i.e., reliance on detailed, and commonly unavailable, construction information; and/or need for time consuming and expensive on-site measurements). In addition, we also discussed the respective shortcomings and limitations of black-box approaches (i.e., need for vast amount of training data, hard to ensure reliability), and justified our decision to incorporate an adaptive gray-box approach in our domestic heating automation system. To this end, we reviewed the literature of both fixed and adaptive gray-box thermal modeling to show how the latter has been shown to be resilient and effective in highly dynamic settings, such as houses. This makes them a suitable choice for our DHAS with respect to our reliable thermal modeling, applicability, generality and minimal user-input requirements (as discussed in Section 1.1.2).

Further on, in Section 2.4, we reviewed the literature of occupancy prediction approaches with respect to their efficiency, computational power needs and infrastructure requirements. In so doing, we showed how to choose one of these approaches for incorporation into our DHAS in order to rely to the minimum extent on user-input and deal with occupancy uncertainty (as per the respective requirements discussed in Section 1.1.2). Subsequently, in Section 2.6 we provided a detailed discussion over the wide family of the commonly employed MPC control approaches, with respect to their theoretical properties. Specifically, we showed how we choose a particular MPC approach for our DHAS, based on its theoretical properties and the particular characteristics of our application domain, with respect to the Pareto efficiency and applicability requirements of our system (Section 1.1.2). Then, in Section 2.7 we reviewed the literature of heating automation systems that deal with occupancy uncertainty (and hence could potentially be employed in domestic settings), positioning our work against the state-of-the-art and choosing our benchmark for our evaluation procedure that best meet our work requirements. There, we also discussed the shortcomings of current approaches with respect to the basic requirements of a domestic heating automation system as detailed in Section 1.1.2. Finally, in Section 2.8 we reviewed the literature of DHASs with respect to the requirements of economic control (Section 1.1.3). There, we illustrated that the DHASs that deal with economic control are highly experimental and possess considerable limitations with respect to real setting scenarios. In addition, we showed that none of the DHASs that consider economic control deal with advanced economic control. As such, in Sections 2.7 and 2.8 we detailed the limitations of current approaches (summarized in Section 1.1.4) that consider the motivational aspect for the development of AdaHeat.

Subsequently, in Section 2.9 we detailed the literature on predicting the power output of IERs and we provided a general discussion over the pros and cons of different approaches. By doing so, we pointed out the limitations and shortcomings of NWP-based approaches that make them unsuitable for incorporation in our advanced economic control schema (i.e., high complexity and absence of site-specific calibration). In addition, we discussed the respective shortcomings and limitations of statistical approaches (i.e limited mid-to-long term accuracy) and justified our decision to incorporate a hybrid approach in our schema. In Section 2.10 we provided a general discussion on Gaussian process regression, that consider the bases of our hybrid IER power output stochastic prediction approach. Then, in Section 2.12 we discussed the main concepts of cooperative game theory that emerge in our work on collective advanced economic control.

Furthermore, in Section 2.13 we provided a general discussion on Markov decision processes, MDPs, and dynamic programming with respect to our work on IER efficient control. There we also outlined the value of dynamic programming to large MDPs (like the one considered in our work, as further justified in Chapter 5). In particular, formalizing solar tracking, ST, as a dynamic programming problem in the context of a MDP is a key contribution of our work (as discussed in Section 6.4). This choice is justified by our performance optimality requirement, along with the generality and applicability requirements (as discussed in Section 1.2.2). In more detail, current ST approaches although they generally meet the generality requirement (as detailed in Section 1.2.2), in the best case only partially meet the performance optimality requirement (and usually with a trade-off in their applicability). This fact motivates the development of PreST to meet all the respective ST requirements, as stated in Section 1.2.2 (i.e., optimality, applicability and generality), utilizing optimal control.

Chapter 3

AdaHeat: A General Adaptive Domestic Heating Automation System

In this chapter we describe our domestic heating automation system, AdaHeat. As discussed in Section 1.3.1, AdaHeat is able to account for: (i) simple heating control, (ii) simple economic control, as well as (iii) advanced economic control. In this context, AdaHeat is also able to exploit the coalition potential that arises in advanced economic control. To this end, here we also detail our proposed scheme for collective advanced economic control in the context of AdaHeat. In more detail, in Section 3.1 we discribe our domestic heating automation system with respect to simple heating control, simple economic control and advanced economic control in single houses. Subsequently, in Section 3.2 we detail our scheme for collective advanced economic control. Finally, Section 3.3 summarizes this chapter.

3.1 A General Adaptive Domestic Heating Automation System

Here we detail our domestic heating automation system with respect to single-house space heating control. In general, AdaHeat consists of the following components: (i) the *thermal comfort model*, (ii) the *thermal model* of the building, (iii) the *heating cost model*, and (iv) the *controller*, that utilizes the aforementioned components, as seen in Figure 3.1. We now proceed to describe each component in detail.

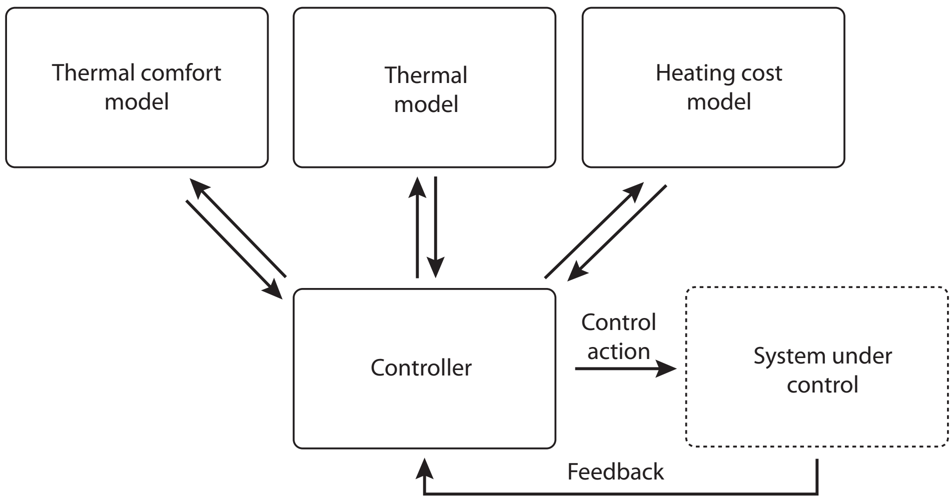

Figure 3.1: AdaHeat flow diagram

3.1.1 Thermal Comfort Model

In essence, thermal comfort is a complex response to several potentially interacting and less tangible physical, physiological, psychological, and other factors (e.g., differences in mood, activity, biology, clothing, air temperature, humidity, and air speed) (Djongyang et al. (2010)).[1] As such, based on different assumptions, a variety of metrics have been proposed to measure thermal discomfort (for a review on thermal discomfort metrics see Djongyang et al. (2010)). In this work, for reasons of simplicity, we assume discomfort to depend only on the inside air temperature, T^{IN}; and any discomfort experienced, at each instance that the house is occupied, is the absolute deviation of T^{IN} from the user-provided, set-point temperature, T^{SP} (as seen in Figure 3.2). As such, thermal discomfort during interval τ of length δ is calculated as:

$$\mathrm{Disc}\left(\cdot_{\tau}\right) = \begin{cases} \displaystyle\int_{t_0}^{t_0+\delta} \left|T^{SP} - T^{IN}(t)\right| \, \mathrm{d}t, & \text{occupied} \\ 0, & \text{otherwise} \end{cases} \tag{3.1}$$

where t_0 is the starting time of the interval. Now, assuming constant T^{IN} within the particular interval, Equation 3.1 becomes:

$$\mathrm{Disc}\left(\cdot_{\tau}\right) = |T^{SP} - T^{IN}_{t_0}|\delta \ \mathbb{1}_{\textbf{occupied}} \tag{3.2}$$

More advanced thermal comfort modeling (e.g, Auffenberg et al. (2015); Langevin et al. (2013)) can be directly incorporated in our approach by simply adjusting the above equation accordingly.[2]

[1] Optimal thermal comfort has been defined as "the condition of the mind in which satisfaction is expressed with the thermal environment" (according to the ANSI/ASHRAE Standard 55, "Thermal Environment Conditions for Human Occupancy", 2013).

[2] In contrast to heating automation systems that rely on complexity-specific (e.g., linear, quadratic and convex) programming approaches for planning (e.g., Halvgaard et al. (2012); Oldewurtel et al. (2010c)),

Now, as outlined above, any thermal discomfort within the house is experienced only when the house is occupied. As such, the occupancy schedule is essential for modeling and predicting thermal discomfort. However, the occupancy schedule is usually unknown in domestic settings and needs to be predicted. To this end, several approaches to predict the occupancy schedule have been proposed over time (see Section 2.4). In this work, we employ the schedule-based occupancy prediction approach proposed by Scott et al. (2011) due to the general low instrumentation needs of schedule-based approaches and its particular efficiency. That said, any other occupancy prediction approach can be used to provide the future occupancy estimates in our model. Now, this algorithm predicts the occupancy schedule on-line (in real time) and returns a vector that corresponds to the probability of occupancy in a 15-minute interval over the predicting horizon.[3] As such, assuming a constant T^{IN} during interval τ of length δ, the expected thermal discomfort is calculated as:

$$\mathbb{E}\left[\mathrm{Disc}(\cdot_\tau)\right] = O_{t_0}\, |T^{SP} - T^{IN}_{t_0}|\, \delta \tag{3.3}$$

where O_{t_0} is the occupancy probability during the particular interval and t_0 is the starting time of the interval.

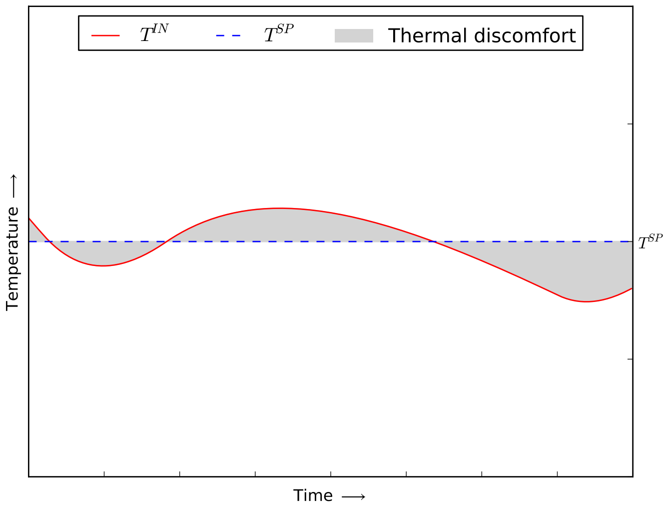

FIGURE 3.2: We measure thermal discomfort as the absolute deviation of T^{IN} from T^{SP} (over time) and, as such, thermal discomfort corresponds to the shaded area in-between T^{IN} and T^{SP} as illustrated above.

the dynamic programming planning approach of AdaHeat+ (further discussed in Section 3.1.4) does not generally raise constraints on the form of the components' modeling.

[3]For the needs of this work we interpolated any estimates where necessary.

3.1.2 Thermal Model

As discussed in Section 2.2, the thermal dynamics of domestic buildings are much harder to model accurately than their non-domestic counterparts, as they rely on additional factors which are hard to predict. Nevertheless, adaptive thermal modeling has been shown to be resilient and effective in such high dynamic settings where the partial observation can be interpreted as non-stationarity of the assumed thermal process. Moreover, gray-box modeling requires relatively small amounts of training data in order to demonstrate adequate performance and does not require explicit knowledge or on-site measurements of the thermal characteristics of a building (see Section 2.2). As such, in our system we employ adaptive gray-box thermal modeling where the equivalent thermal parameters (ETPs), are estimated on-line and are assumed to be time-varying.

Now, in general, a thermal model predicts the thermal response of a building based on: (i) the current thermal state vector of the building, $\mathbf{x}$; (ii) the vector of heating control actions to be executed, $\mathbf{a}$; and (iii) the vector of information variables regarding exogenous processes that affect the thermal process (e.g, incident solar radiation, outside or adjacent buildings' temperature), $\mathbf{i}$.[4] As such, in discrete-time form, a thermal model can generally be defined as:

$$\mathbf{x_{t+1}} = \mathcal{TM}(\mathbf{x_t}, \mathbf{a_t}, \mathbf{i_t}) \tag{3.4}$$

In more detail:

- **Thermal state vector, x:** The current thermal state of a building considers variables that influence the thermal response of the building but are also influenced by the heating control actions followed. For instance, depending on the thermal model used, the thermal state vector might only consider the indoor air temperature, T^{IN}, or also consider temperature values of intermediate nodes of the thermal process, such as the floor temperature, T^{FL}, or the building envelope temperature, T^{EN}. That said, the thermal state of a building can potentially be defined as only the current thermal condition values or as the current values along with historical values.

- **Heating control action vector, a:** In general, the possible heating control actions depend on the heating system employed. However, typical domestic heating systems usually allow for the end-user to define a set-point temperature (commonly a fixed integer value or a value of limited decimal accuracy; e.g., with intervals of 0.5 or 0.1) and to entirely switch on/off the heating system through the thermostat. In this context, a low-level control is utilized to retain the indoor air temperature, T^{IN}, close to the user-defined set-point temperature, T^{SP}. This low-level control

[4] In this thesis we use the typical notation of denoting vectors and matrices with bold lower-case and bold upper-case letters respectively. Moreover, when not stated otherwise, a vector is assumed to be a column vector.

is typically a simple ON/OFF control that switches off heating if $T^{IN} > T^{SP}$ and turns it on otherwise. In this context, setting T^{SP} higher (or lower) than T^{IN} within a low-level control cycle is equivalent to switching on (or off respectively) heating directly. As such, the available actions for a DHAS depend only on the different modes of operation of the heating system employed. For instance, for a typical system with only one operating mode, the heating actions can be denoted by a single binary scalar a populated with either 0 or 1 to indicate heating on and off respectively. In contrast, if the heating system employed has multiple operating modes (such as an electric radiative space heater which can operate with either 2 or 4 elements, i.e. resistances) the heating control action vector (or, potentially, the possible populations of a single scalar) should be extended accordingly.

- **Vector of exogenous processes variables, i:** The external stochastic processes variables can range from simple observations (or estimations) to historical averages and/or predicted values. The exact variables utilized depend on the thermal model used and the availability of these quantities.

As such, any thermal model, $\mathcal{TM}(\mathbf{x}, \mathbf{a}, \mathbf{i})$, is essentially a function over the inputs ($\mathbf{x}$, $\mathbf{a}$, and $\mathbf{i}$) that aims to predict the thermal response of a building with high accuracy. In gray-box modeling the parameters of this function consider the ETPs that need to be estimated and can be assumed to be either time-varying or fixed. This choice depends on the dynamics of the underlying process and the model's structural predictive abilities (i.e the form of the equation and the information that it accounts for).

Example 3.1. *For instance, a linear (with respect to* $\mathbf{x}$*,* $\mathbf{a}$*, and* $\mathbf{i}$*), discrete-time thermal model can be generally defined as:*

$$\mathbf{x_{t+1}} = \mathbf{C_a x_t} + \mathbf{C_b a_t} + \mathbf{C_c i_t} \tag{3.5}$$

Now, considering the thermal state vector $\mathbf{x} = \begin{pmatrix} T^{IN} \\ T^{FL} \end{pmatrix}$*, which could be held suitable for underfloor heating systems, and only one heating control binary variable,* a*, (of switching on/off heating) the above set of linear equations will be:*

$$T^{IN}_{t+1} = (\mathbf{C_a})_{1,1}\ T^{IN}_t + (\mathbf{C_a})_{1,2}\ T^{FL}_t + (\mathbf{c_b})_1\ a + (\mathbf{C_c})_\mathbf{1}\ \mathbf{i} \tag{3.6}$$

$$T^{FL}_{t+1} = (\mathbf{C_a})_{2,1}\ T^{IN}_t + (\mathbf{C_a})_{2,2}\ T^{FL}_t + (\mathbf{c_b})_2\ a + (\mathbf{C_c})_\mathbf{2}\ \mathbf{i} \tag{3.7}$$

where, all matrices and vectors denoted with $\mathbf{C}$ *and* $\mathbf{c}$ *respectively correspond to time varying or fixed ETPs that need to be estimated. Nevertheless, some of the ETPs can be set manually if information is available (and/or to maintain a physical meaning). For instance, the coefficient* $(\mathbf{c_b})_1$ *can be manually set to zero as the control actions do not directly affect* T^{IN} *for underfloor heating systems.*

In this context, the vectors $\mathbf{x}$, $\mathbf{a}$, and $\mathbf{i}$ are designed to provide sufficient information for a thermal model to be able to predict the building thermal response. Taking this concept to its furthest limit, incorporating all the necessary information in order for the next system state to depend *only* on the information provided by the vectors $\mathbf{x}$, $\mathbf{a}$ and $\mathbf{i}$, would theoretically allow us to define a thermal model with potentially absolute accuracy. Moreover, given that all the necessary information is being accounted for by our thermal model, an inadequate modeling of the relationship between these variables would lead to prediction inaccuracies that are fixed in time (i.e., the deviation of the predicted thermal state from the real one would be the same for every time a possible combination of $\mathbf{x}$, $\mathbf{a}$ and $\mathbf{i}$ populations is encountered). However, in real settings, it is often the case that not all the necessary information is accounted for by a thermal model (see Section 1.1.2) In such cases, the deviation of a predicted state and the real one might change for a possible combination of the variables in $\mathbf{x}$, $\mathbf{a}$ and $\mathbf{i}$ over time, as the evolution of the system depends on additional effects that are not considered (non-modeled dynamics). In such cases, the variation observed could be stationary or non-stationary. For instance, if the non accounted for information at each time instance is just an i.i.d. random variable, then the observed deviation will follow a stationary distribution. On the other hand, if the non-accounted for information considers occupant activity (as is the case in domestic settings), such as opening a window or cooking, the deviation observed might demonstrate non-stationarity.

Now, as noted in Section 2.2, assuming time-varying ETPs, in the context of gray-box thermal modeling, has been shown to be effective in the highly dynamic domestic settings. In this context, depending on the complexity of the thermal model (linear or non-linear), different real-time parameter identification methods can be employed such as recursive least square with forgetting factor, Kalman filters or extended Kalman filters (as discussed in Section 2.2). That said, the state vector can be partially observable as well. For instance, readings of the floor temperature in the thermal model provided in the above example might not be available or might be very inaccurate. As such, estimating state variables along with the parameters might be needed. Even in the case that a linear thermal model is used, the simultaneous estimation of parameters and state variables introduces a non-linear problem in general (Haykin (2001)). Now, model selection is an essential part of gray-box thermal modeling. Although several methodologies are proposed for thermal model selection (e.g., Prívara et al. (2012); Bacher and Madsen (2011); Andersen et al. (2000); Kristensen et al. (2004)), identifying the most suitable model depends highly on the thermal process being modeled and the application requirements. As such, it is typically undertaken by the designer. Nevertheless, our system is able to handle both linear and non-linear thermal models due to our general control approach based on dynamic programming (as further discussed in Section 3.1.4). As such, it does not raise respective restrictions in identifying the most suitable thermal model. Specific thermal modeling instantiations are provided for our case study systems in Chapter 4.

3.1.3 Heating Cost Model

The heating systems employed in domestic settings are diverse in technology and type. However, all of them consume an amount of energy over their operation. Strictly speaking, this amount of energy is the energy provided to the space over the efficiency of the system. Now, the input power of a heating system, at any time instance t, is:

$$\mathrm{Cons}\left(\cdot_t\right) = \frac{\mathrm{Pwr}\left(\mathbf{x_t}, \mathbf{a_t}\right)}{C^{Eff}} \tag{3.8}$$

where $\mathrm{Pwr}(\mathbf{x}, \mathbf{a})$ stands for the output power of the space heating system according to the heating control action vector and the building thermal state; and C^{Eff} stands for the heating system efficiency.[5] Considering import and export tariffs as well as domestic intermittent energy generation capacity (as per the simple and advanced economic control requirements, see Section 1.1), the energy consumption cost of a heating system during interval τ of length δ (where all variables remain unchanged) is:

$$\begin{aligned} \mathrm{Cost}(\cdot_\tau) &= \Big(\max\left(0, \mathrm{Cons}(\cdot_{t_0}) - R_{t_0}\right) P_{t_0}^{Buy} \\ &+ \min\left(0, \mathrm{Cons}(\cdot_{t_0}) - R_{t_0}\right) P_{t_0}^{Sell}\Big)\ \delta \end{aligned} \tag{3.9}$$

where $P_{t_0}^{Buy}$, $P_{t_0}^{Sell}$ and R_{t_0} respectively stand for the import tariff, the export tariff and the power output of the intermittent energy resources (IERs) at the starting time of the interval, t_0. Note that in the special case where no domestic IER generation capacity is available, Equation 3.9 becomes:

$$\mathrm{Cost}(\cdot_\tau) = \mathrm{Cons}(\cdot_{t_0})\ P_{t_0}^{Buy}\ \delta \tag{3.10}$$

Now, in contrast to other approaches where fixed equivalence formulas or multiple user-provided parameters are used to balance heating cost and discomfort (e.g., Mozer et al. (1997); Gao and Keshav (2013b)), in our approach this balancing is adaptive to the user preferences (enhancing the usability of our approach as discussed in Sections 1.1.3 and 2.7). In particular, this adaptation is carried out in real-time, based on a single Boolean feedback from the user which progressively adjusts a *single* weighting parameter, as we further discuss in Section 3.1.4. For this reason, in the cases of simple heating control and simple economic control, we are not generally interested in the particular consumption of the heating system over a fixed time of operation. In those cases, the particular heating system efficiency, as well as the energy provided to the space for a particular heating action, can be set to arbitrary values without any loss in performance

[5] C^{Eff} can be a simple scalar or a function of the temperature difference between the heat source and the destination to also account for the coefficient of performance variability when heat pump technology is considered (see Section 2.1.2).

(as long as any needed ratios are retained).[6] Nevertheless, in the special case of advanced economic control the particular energy consumption of the heating system is needed and, hence, so are the particular $\mathrm{Pwr}(\cdot)$ and C^{Eff} values (or adequate estimates).

In addition, in the special case of advanced economic control, adequate estimates of the future IER power output are also needed in order for any heating cost model to have a practical value. To account for the respective uncertainty, a stochastic approach can be considered as discussed in Section 2.9. In particular, the power output of an IER can be modeled and predicted as a stochastic process R valued within the range $[0, r^{\max}]$, where $r^{\max}$ is the maximum power output that can be achieved, at any time, by the particular IER. In this context, a variety of stochastic modeling approaches, including Kalman filters (Louka et al. (2008)) and Gaussian processes (GPs) (Jiang et al. (2010)), can be used to model the IER power output in order to provide probabilistic estimates of the future power output, henceforth denoted R' (Zamo et al. (2014)) (as discussed in Section 2.9). In this context, the expected heating cost (Equation 3.11) is subject to the particular IER power output modeling:

$$\mathbb{E}\left[\mathrm{Cost}(\cdot_{\mathcal{T}})\right] = \mathbb{E}\left[\max\left(0, \mathrm{Cons}(\cdot_{t_0}) - R'_{t_0}\right) P^{Buy}_{t_0} + \min\left(0, \mathrm{Cons}(\cdot_{t_0}) - R'_{t_0}\right) P^{Sell}_{t_0}\right] \delta \tag{3.11}$$

In Chapter 4 we provide an instantiation of IER stochastic prediction based on adaptive GP modeling for our case study of wind turbine generators. There we also provide a closed form of Equation 3.11 in accordance with our GP-based IER modeling. Now, we note here, that although our work is motivated on import and export tariffs that are known in advance, our approach can also work in conjunction with tariffs that are not generally known in advance (Jenner et al. (2013)), as long as an appropriate tariff prediction approach is considered in the context of our heating cost modeling.[7]

3.1.4 Controller

The aim of our DHAS is to be effective for a variety of heating systems, such as heating systems with considerable thermal lags and heating systems with variable heating cost. As such, we regard heating control in an infinite horizon optimization manner. To this end, we employ model predictive control (MPC) which also allows us to directly and effectively incorporate on-line adaptation of the thermal model to account for the dynamic domestic thermal characteristics. Now, as discussed in Section 2.6, the thermal process of buildings is typically slow and stable, thus enabling us to design an MPC controller generally based on performance criteria alone. As such, in this work we employ

[6] For instance, for a simple fixed-efficiency electric radiative space heater which can operate with either 2 or 4 elements (i.e. resistances) of the same nominal power, the corresponding energy provision values for each heating action, a (i.e., number of resistances in operation, i.e. 0,1,2), should correspond to the ratio 0:1:2.

[7] In this context, $P^{Buy}_{t_0}$ and $P^{Sell}_{t_0}$ in Equation 3.11 will consider, potentially stochastic, estimates of $P^{Buy}_{t_0}$ and $P^{Sell}_{t_0}$ derived in accordance to the tariff prediction approach utilized.

an adaptive MPC approach (see Section 2.6) that works as follows. Every δ amount of time, the controller plans a heating schedule over a fixed horizon into the future (utilizing the above models and relevant predictions) and executes the first action of the planned schedule. Then, new estimates of O, R , $\mathbf{i}$, (potentially, P^{Buy} and P^{Sell}) and the thermal model parameters are acquired and the procedure is repeated shifting the planning horizon into the future. We now proceed to describe our planning approach, detailing our planning objective formalization and our optimization approach.

3.1.4.1 Planning (Objective formalization)

As discussed in Section 1.1.1, the objective of a heating automation system is to balance the discomfort experienced by the occupants and the cost of heating according to a predefined condition. As there is a conflict between discomfort and cost, defining the optimal heating schedule is a two-objective optimization problem. In order to tackle the respective complexity, in this work we combine the objective functions to form a single scalarized function. For the single scalarized function we use the common *weighted sum* (Marler and Arora (2004)) which is a sufficient but not necessary formalization for Pareto optimality (Zionts (1989); Zadeh (1963)). Hence, we ensure that all the derived heating schedules fall in the Pareto optimal set. However, it is not guaranteed that our method is able to capture *all* the optimal schedules, as long as the Pareto optimal hypersurface is not convex (Marler and Arora (2010)). That being said, other scalarization methods exist that are both necessary and sufficient conditions for Pareto optimality even for the case of non-convex problems (for a comprehensive review of these methods see Marler and Arora (2004)). However, we use the weighted sum due to its simplicity and good observed performance (as further illustrated in Chapter 4). Moreover, using the weighted sum allows our system to adapt to the user's preferences through a simple linear feedback procedure as it depends on only one weighting factor that can be learned on-line to reflect the user preferences. For instance, if the user feels that the system is consuming a lot of energy and he/she is willing to experience some thermal discomfort, he/she can simply progressively reduce the weighting factor by a constant value until his/her preferences are met.

In more detail, in our approach we plan for the MPC horizon, of length Δ, by breaking it down into a set of non overlapping intervals of length δ. As such, ensuring that Δ is an integer multiple of δ, the planning horizon corresponds to a set of intervals, noted H, where $|H| = \Delta/\delta$. During each interval, identified by a unique id number, τ, all environmental conditions are assumed constant. Now, our planning objective is to assign at each interval a heating control action vector, $\mathbf{a}_{\boldsymbol{\tau}}$, in order to minimize a unifying scalarized function of heating cost and expected thermal discomfort. More formally, the objective of our planning is to find the sequence of action vectors, $\mathbf{a}_{\boldsymbol{\tau}}$, that minimizes

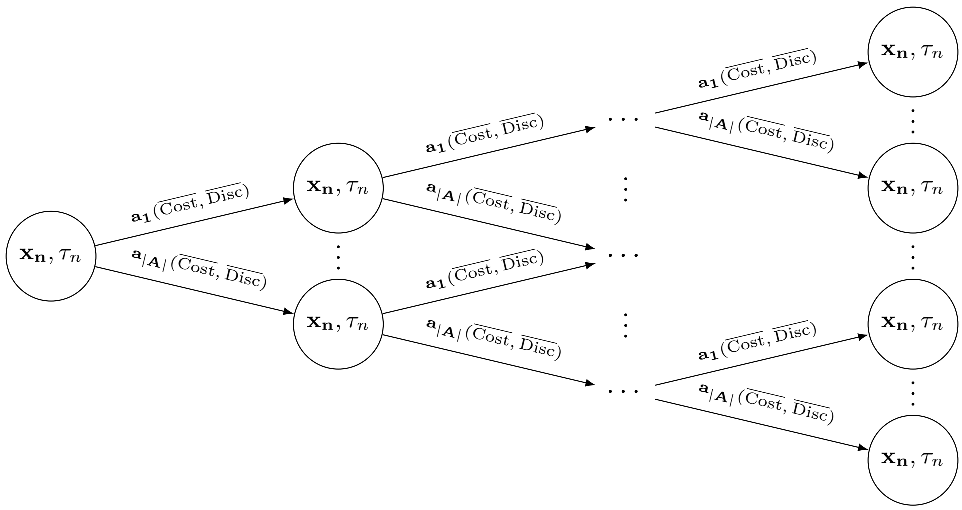

FIGURE 3.3: For each interval, τ, all environmental conditions are assumed constant. The edges of the DAG correspond to the possible action vectors, $\mathbf{a_1}, \ldots, \mathbf{a_{|A|}} \in A$, with their weights, noted $\overline{\text{Cost}}$ and $\overline{\text{Disc}}$, accounting for the corresponding expected heating cost and expected thermal discomfort during the particular interval, respectively. Each node is identified by a thermal state vector, $\mathbf{x}$, and the interval id number, τ.

the expected unifying cost, $\bar{J}$, over the planning horizon, as defined below:

$$\begin{aligned} \underset{\mathbf{a_1},\ldots,\mathbf{a_{|H|}}}{\text{minimize}} \quad & \bar{J}(\mathbf{a_1}, \ldots, \mathbf{a_{|H|}}) = \sum_{\tau=1}^{|H|} \lambda \, \mathbb{E}\left[\text{Disc}(\cdot_\tau)\right] + (1-\lambda) \, \mathbb{E}\left[\text{Cost}(\cdot_\tau)\right] \\ \text{subject to} \quad & \mathbf{a_1}, \ldots, \mathbf{a_{|H|}} \in A \end{aligned}$$

Here, $\mathbb{E}\left[\text{Disc}(\cdot_\tau)\right]$ and $\mathbb{E}\left[\text{Cost}(\cdot_\tau)\right]$ return the expected discomfort and expected heating cost during each interval, τ, respectively, λ stands for the weighting factor between cost and discomfort that is learned on-line, and A is the set of all feasible heating control action vectors (used here to intuitively denote all possible constraints on heating actions).[8] In general, λ values should be in the range $(0, 1)$, without considering the limits 0 and 1, in order to ensure strict Pareto optimality (Marler and Arora (2004))—as for these limit values, minimizing only one of the objectives is considered (i.e., either only heating cost or thermal discomfort). That being said, as there exists *only one* unique heating strategy for minimizing cost, i.e., of not doing any heating, the derived heating schedule is Pareto optimal even with $\lambda = 0$. However, searching for this *already known* trivial strategy is inefficient. As such, for the extreme case of eliminating heating cost with the minimum discomfort, heating should be off. On the other hand, for eliminating discomfort with the minimum heating cost a λ value very close to 1 can be used. Note that normalized values for cost and discomfort can be used. We now proceed to describe our planning algorithm.

[8]Note that the absence of any output variable constraints in our optimization objective ensures that we will not face any feasibility issues in planning (as discussed in Section 2.6).

3.1.4.2 Planning (Optimization)

The slow nature of the thermal process of buildings enables us to tackle the optimization problem defined in Section 3.1.4.1 in a general, dynamic programming manner, as the real-time computation constraints are not typically strict (i.e., the amount of time between two consecutive heating control actions, δ, is typically set to several minutes, allowing the heating control actions to take effect—hence, providing us with a sufficient amount of time for planning). Moreover, the utilization of dynamic programming enables us to generally choose the objective function and the system model based on performance criteria alone. As such, we reduce planning to finding the shortest path in a directed acyclic graph (DAG) and we provide a dynamic programming planning algorithm that exploits the property of topological ordering of a DAG through depth first search (DFS) in order to find the shortest path in linear time (Cormen et al. (2001)). It is worth noting here that, although in this work we do not deal with constraint optimization approaches, these appropriately correspond to the respective constraint shortest path finding problems.[9]

In more detail, each node, n, of the graph G corresponds to a distinct tuple that contains all the necessary information to predict the next state, n', based on the heating control action vector to be followed, $\mathbf{a}$, through the appropriate transition model, $n' = T(n, \mathbf{a})$. Let us note here that the vector of exogenous processes, $\mathbf{i}$, at each instance can be inferred by the time step id, τ, alone. Hence, the n tuple will be $\langle \mathbf{x_n}, \tau_n \rangle$. As such, the transition model $n' = T(n, \mathbf{a})$ considers the thermal model and a simple incrementer to appropriately provide $\mathbf{x'_n}$ and τ'_n respectively. Now, the edges of the graph correspond to the transitions in time due to the heating action. In this context, each of the edges, noted by e, has two weights corresponding to the cost of heating and the expected discomfort during the respective interval. As such, each edge corresponds to a tuple that contains the initial node, the successor node and the aforementioned weights. Given the above, the corresponding graph will be directed and acyclic as illustrated in Figure 3.3. In general, dynamic programming needs a discretization of the node variables to work effectively (Cormen et al. (2001)). Now, given the limited predicting ability of any thermal model used in practice, a discretization of the thermal state vector comes naturally and would not introduce any additional uncertainty,[10] while the time step id, τ, is appropriately already discrete.

In Algorithm 3 we provide a dynamic programming planning algorithm based on the DFS recursion. In particular, we extend the DFS recursion with constant time expressions, thus the time complexity is retained at $O(|V_G| + |E_G|)$ where V_G and E_G stand for the set of edges and vertices of the graph respectively. Specifically, the algorithm follows the DFS recursion to create the DAG and populates the dictionaries $\mathrm{Min\bar{J}}\{n\}$

[9] The *Constraint Shortest Path Problem (CSPP)* is well known to be NP-complete making long-term planning inefficient (Garey and Johnson (1979))

[10] In contrast to the claim in Mozer et al. (1997).

Algorithm 3 General Heating Planning Algorithm

1: **procedure** HEATINGPLANNING(G, n)
2:   **for** every $\mathbf{a} \in A$ **do**
3:     $\overline{\text{Cost}} \leftarrow \mathbb{E}\left[\text{Cost}(\cdot_{\tau_n})\right]$ ▷ Obtain succesor node and edge weights
4:     $\overline{\text{Disc}} \leftarrow \mathbb{E}\left[\text{Disc}(\cdot_{\tau_n})\right]$
5:     $n' \leftarrow T(n, \mathbf{a})$
6:     $e \leftarrow < n \mapsto n', \overline{\text{Cost}}, \overline{\text{Disc}} >$
7:     add e to E_G
8:     **if** $n' \notin V_G$ **then**
9:       add n' to V_G
10:       **if** $\tau_{n'} < |H|$ **then**
11:         $G \leftarrow$ HEATINGPLANNING(G, n')
12:       **else**
13:         $\text{Min}\bar{\text{J}}\{n'\} \leftarrow 0$
14:     $\text{Tmp} \leftarrow \text{Min}\bar{\text{J}}\{n'\} + \lambda\, \overline{\text{Disc}} + (1-\lambda)\, \overline{\text{Cost}}$ ▷ Populate dictionaries with minimum additional expected cost and best action
15:     **if** $\text{Min}\bar{\text{J}}\{n\} = \text{NaN}$ **or** $\text{Min}\bar{\text{J}}\{n\} > \text{Tmp}$ **then**
16:       $\text{Min}\bar{\text{J}}\{n\} \leftarrow \text{Tmp}$
17:       $\text{BestAction}\{n\} \leftarrow \mathbf{a}$
18:   **return** G

and BestAction$\{n\}$ with the minimum additional expected unifying cost, $\bar{J}$, and the best action vector for each node respectively. As such, after the algorithm terminates, the dictionary BestAction$\{n\}$ holds the optimal heating action vector for each node. The arguments at the initial call of the recursion consider an empty graph and the root node. The recursion then creates the DAG in a pre-order (Cormen et al. (2001)) manner, lines 2-13, and when a final node is reached the additional minimum $\bar{J}$ is populated with 0, line 13. As the recursion folds back (i.e. the recursive calls return) the minimum additional expected unifying cost of each node is populated along with the best action vector for this node, lines 14-17.

3.2 Collective Advanced Economic Control

As discussed in Section 1.1.3, in contrast to simple heating and simple economic control, the benefits of advanced economic control can be further amplified in domestic coalitions. In such settings, a number of houses share their energy generation and shift their heating consumption in order to further minimize the energy imported from the grid. As such, in this work we also propose a practical scheme for collective advanced economic control. In particular, we propose the formation of coalitions of houses that coordinate their heating system operation, ahead of time, so as to efficiently use shared grid-connected IERs. Notably, these houses are not required to be geographically adjacent (as discussed in Section 1.1.3). To this end, we propose a heuristic heating schedule planning approach in the context of our MPC approach to account for collective advanced economic control. Our heuristic approach ensures the practical applicability of AdaHeat in contrast to optimal planning, as further discussed in the following paragraphs. Finally, AdaHeat also incorporates a cost allocation mechanism to share the coalition gains in such settings, that respects individual rationality. We now proceed to detail the proposed scheme.

3.2.1 Optimal Planning

To account for collective advanced economic control, we adjust the planning objective, in the context of our MPC approach. To do so, we first model the expected collective thermal discomfort and expected collective heating cost of the coalition. The former, appropriately considers a simple extension of Equation 3.3 to account for the total comfort of the coalition. More formally, assuming all variables remain constant during interval τ of length δ, the expected thermal discomfort of the coalition C, is the sum:

$$\mathbb{E}\left[\mathrm{Disc}_C(\cdot_\tau)\right] = \sum_{i=1}^{|\,\mathrm{C}\,|} \mathbb{E}\left[\mathrm{Disc}_i(\cdot_\tau)\right], \quad \text{where } \mathbb{E}\left[\mathrm{Disc}_i(\cdot_\tau)\right] = O_{i,t_0} \left|T_i^{SP} - T_{i,t_0}^{IN}\right| \tag{3.12}$$

where O_{i,t_0} and T_{i,t_0}^{IN} are the occupancy probability and inside temperature at time t_0 (the starting time of interval τ) for each house $i \in$ C, respectively (assumed constant during interval τ as noted above).[11] In addition, the aggregate input power of the heating systems of C, at any time instance t, considers a simple extension of Equation 5.7, in a similar manner:

$$Cons_C(\cdot_t) = \sum_{i=1}^{|\,\mathrm{C}\,|} Cons_i(\cdot_t), \quad \text{where } Cons_i(\cdot_t) = \frac{\mathrm{Pwr}_i\left(\mathbf{x_{i,t}}, \mathbf{a_{i,t}}\right)}{C_i^{Eff}} \tag{3.13}$$

Nevertheless, given the fact that the members of the coalition share their IER power output, the collective heating cost is not a simple sum of the heating cost of each member. In particular, during interval τ of length δ and assuming all variables remain constant to their value at time t_0 (the starting time of interval τ), the expected heating cost of coalition C will be:

$$\begin{aligned} \mathbb{E}\left[\mathrm{Cost}_C(\cdot_\tau)\right] &= \mathbb{E}\Big[\max\left(0, Cons_C(\cdot_{t_0}) - R'_{C,t_0}\right) P_{t_0}^{Buy} \\ &+ \min\left(0, Cons_C(\cdot_{t_0}) - R'_{C,t_0}\right) P_{t_0}^{Buy}\Big]\delta \end{aligned} \tag{3.14}$$

where R'_{C,t_0} stands for the stochastic estimate of the cumulative future IER power output at time t_0. In particular, letting I stand for the set of the shared IERs owned by the coalition C, the cumulative power output of the shared IERs can be modeled and predicted as a stochastic process R_C. In essence, R_C considers the sum of |I| stochastic processes R_i where $i \in \{1, 2, \ldots, |\mathrm{I}|\}$. Each one of these sub-processes considers the power output of each one of the shared IERs. Now, letting t stand for time, R_C will be

[11] Note that T_i^{SP} and O_i are also defined according to house i, as different houses generally have different set-point preferences and occupancy schedules.

a collection of random variables $R_{C,t}$ with the cumulative distribution function (CDF) $F_C(r;t) = \mathbb{P}(R_{C,t} \leq r)$, which has support on a subset of $\mathbb{R}_+^N$. The corresponding probability density function is denoted by $f_C(r;t)$. Furthermore, the corresponding quantile function is $F_C^{-1} : [0,1] \leftarrow [0, r_C^{\max}]$, where $r_C^{\max}$ is the maximum power output that can be achieved in total, and at any time, by the shared IERs. We note here that a variety of stochastic modeling approaches can be used to model the power output of each one of the shared IERs (as discussed in Section 3.1.3) or directly the cumulative IER power output in order to provide probabilistic estimates of the cumulative future power output R'_C. As discussed in Section 3.1.3, in Chapter 4 we provide an instantiation of IER stochastic prediction based on adaptive GP modeling for our case study of wind turbine generators. There we also provide a closed form of Equation 3.14 in accordance with our GP-based IER modeling.

Now, given the above modeling, we are able to define our optimization objective that corresponds to the optimal plan for collective advanced economic control in the context of our MPC approach. In particular, the planning optimization objective is to find the action matrix, $\mathbf{A}$, that is the sequence of heating control actions for each house, that minimizes the expected unifying cost, $\bar{J}_C$, over the planning horizon, H:

$$\begin{aligned} \underset{\mathbf{A}}{\text{minimize}} \quad & \bar{J}_C(\cdot) = \sum_{\tau=1}^{|\mathrm{H}|} \left(\mathbb{E}\left[\mathrm{Cost}_C(\cdot_\tau)\right] + \sum_{i=1}^{|\mathrm{C}|} \lambda_i \; \mathbb{E}\left[\mathrm{Disc}_i(\cdot_\tau)\right] \right) \\ \text{subject to} \quad & \mathbf{A}_{i,\tau} \in \mathrm{A}_i \, \forall (i \in \mathrm{C}, \tau \in \mathrm{H}) \end{aligned}$$

where A_i stands for the set of feasible heating control actions for each house $i \in \mathrm{C}$ and λ_i is a scaling parameter that defines the balancing between thermal discomfort and cost for each one of the coalition members. Having a different scaling parameter for each house enables us to capture diverse household preferences in balancing cost and discomfort which is vital in real settings.

Solving the above objective corresponds to a shortest path finding problem, in the emerging directed acyclic graph, G. In this context, the shortest path can be found by exploiting the topological ordering property of G (Cormen et al. (2001)), in linear $\mathrm{O}(|V_G| + |E_G|)$ time, with respect to the number of edges, $|E_G|$, and vertices, $|V_G|$, as discussed in Section 3.1.4. However, even with this low complexity (with respect to the size of the graph), the size of the problem leads to poor scaling performance, as the size of the graph scales exponentially with the size of the coalition. Note that the branching factor of the graph, corresponds to the size of the set of all possible heating action combinations of all houses in the coalition, i.e., $\prod_{i=1}^{|\mathrm{C}|} |\mathrm{A}_i|$. As such, our experiments show that only two-house coalitions are feasible to solve in real time. Moreover, the collective control of all heating systems via a common objective raises further usability issues with respect to real settings. In particular, although different scaling parameters exist for each household, the scalar adjustment of one household affects the cost-discomfort balancing in all the others (through the common objective). This fact leads to a complex and non-linear outcome of the λ population that stops their adjustment through an adaptive boolean feedback procedure (as achieved in the case of single house control, see Section 3.1.4). This, in turn, undermines the applicability of optimal planning in realistic settings.

3.2.2 Heuristic Planning

Given the shortcomings of optimal planning discussed in Section 3.2.1, we propose a heuristic planning approach for collective advanced economic control that enhances the practical applicability of AdaHeat. In particular, we divide equally all available IER energy among the members and plan their heating schedule, given respective IER predictions, independently. More formally, the planning objective for each house $i \in \mathrm{C}$ is:

$$\begin{aligned} \underset{\mathbf{A}_i}{\text{minimize}} \quad & \bar{J}_i(\cdot) = \sum_{\tau=1}^{|\mathrm{H}|} \mathbb{E}\left[\mathrm{Cost}_i^*(\cdot_\tau)\right] + \lambda_i \, \mathbb{E}\left[\mathrm{Disc}_i(\cdot_\tau)\right] \\ \text{subject to} \quad & \mathbf{A}_{i,\tau} \in \mathrm{A}_i \, \forall \tau \in \mathrm{H} \end{aligned}$$

where $\mathbb{E}\left[\mathrm{Cost}_{i,\tau}^*(\cdot)\right]$ is Equation 3.11 adjusted to consider an even share of the total IER energy availability. In particular, during interval τ of length δ and assuming all variables remain constant to their value at time t_0 (the starting time of interval τ), the expected heating cost is:

$$\begin{aligned} \mathbb{E}\left[\mathrm{Cost}_i^*(\cdot_\tau)\right] \quad = \quad & \mathbb{E}\Bigg[\max\left(0, Cons_i(\cdot_{t_0}) - \frac{R'_{C,t_0}}{C}\right) P_{t_0}^{Buy} \\ + \quad & \min\left(0, Cons_i(\cdot_{t_0}) - \frac{R'_{C,t_0}}{C}\right) P_{t_0}^{Buy}\Bigg]\delta \end{aligned} \tag{3.15}$$

The above optimization can be solved, independently for each house, as a shortest path finding problem. As such, the complexity of our approach scales in a linear and parallelizable manner with the size of the coalition, that is $\mathrm{O}(|\mathrm{C}|)$, enabling it to be executed in separate units in every house/member. Note that although an even share of the total IER energy availability is allocated to each house, some houses might end-up consuming less energy than originally allocated to them, while others might require more. Nevertheless, in the context of our collective advanced economic control scheme, planning generally dictates the executed heating actions and not the energy exchange among the members. As such, any excess of energy is shared among the members before buying from the grid. We note here that a more sophisticated allocation of the IER energy among the members, in accordance to their next-day expected needs, could be of value when the respective differences are large. Nevertheless, our evaluation results, considering a wide range of typical houses in the UK (with distinct occupancy patterns and thermal characteristics), illustrate that our approach has high efficiency (and even a near-optimal one) in such a genuine domestic scenario (further discussed in Section 4.1.1).[12]

Apart from its efficiency (further demonstrated in Chapter 4) this approach enables us to populate the λ parameter of each house through a simple boolean feedback procedure. In particular, the expected cost (based on the planned heating schedule) is a worst case scenario with respect to the realized cost—since some houses might consume less energy than originally allocated to them, enabling others to fulfill their needs with lower cost than expected. This worst-case cost is not affected by the λ parameter population of other houses, and considers an upper bound limit that an occupant is willing to pay for the discomfort experienced. As such, our heuristic approach enables the households to balance thermal discomfort and worst-case heating cost, independently and in an adaptive manner for each house, further illustrated in Chapter 4.

3.2.3 Gain Allocating Mechanism

Our work incorporates a cost allocation mechanism to share the realized coalition gains that respects individual rationality (see Section 2.12), motivating, as such, the households to join a coalition. Appropriately, the outside option that we compare to is independent advanced economic control, where each household optimizes its heating control process (considering both export tariffs and its own IER generation capacity) independently.[13] Now, in order to identify the gains of collective advanced economic control compared to independent advanced economic control (and subsequently allocate them), we follow the successive two-step procedure: Initially we perform appropriate simulations to identify the λ population for each house where independent advanced economic control leads to the same discomfort for the house as collective advanced economic control. Subsequently, the difference between the corresponding estimated heating cost of independent advanced economic control (for all members, in total) and the cost of collective advanced economic control is calculated. This difference appropriately corresponds to the gains of forming a coalition. Now, in order to guide our initial λ population search, and provide a computationally practical mechanism, we exploit the fact that in the context of our weighted sum formulation for independent advanced economic control, λ is in a monotonous relationship with discomfort (as also further discussed in Chapter 4). Hence, the specific λ

[12] In this context, considering coalitions where industrial or commercial entities are also able to participate (and, hence, the energy requirement differences are expected to be significant among the members) is a possible future work direction with respect to broader advanced economic control schemes considering a variety of thermostatically controlled loads (e.g, refrigerators, water heaters and air conditions).

[13] Independent advanced economic control considers AdaHeat+ for single-house coalitions and, hence, it is proposed for the first time in this work.

value is identified by progressively increasing (or decreasing) its value (in a hill climbing manner) until the desired discomfort (or a fair approximation) is reached.[14]

Given this initial identification of the gains of collective advanced economic control compared to independent advanced economic control, we proceed with the respective allocation. Now, any allocation of the these gains where each member receives a positive share respects individual rationality (Osborne and Rubinstein (1994)). In this context, we allocate the gains of the coalition, at the end of each day, proportionally to the normalized ratio of produced over consumed energy of each member $i \in \mathrm{C}$, over the day (as realized following collective advanced economic control). More formally, we use the ratio:

$$\frac{\left(\int_{t=0}^{t=D} R_{i,t}\ dt + \theta_0\right) / \left(\int_{t=0}^{t=D} Cons_{i,t}(\cdot)\ dt + \theta_1\right)}{\sum_{i=1}^{|\mathrm{C}|} \left[\left(\int_{t=0}^{t=D} R_{i,t}\ dt + \theta_0\right) / \left(\int_{t=0}^{t=D} Cons_{i,t}(\cdot)\ dt + \theta_1\right)\right]} \tag{3.16}$$

where D is the length of the day, and θ_0 and θ_1 are arbitrarily set small constants to ensure that even the members with nonexistent IER generation still receive a share of the collective gains (motivating them, as such, to join a coalition),[15] and avoid division by zero issues, respectively. This ratio considers an intuitive ranking index of each member that aims to represent its contribution to the collective gains. Notably, using this normalized ranking ensures that all the collective gains are allocated. As such, we provide a computationally practical allocation mechanism that provides imputations, satisfying individual rationality and allocation efficiency (see Section 2.12).

3.3 Summary

In this chapter we described our MPC-based DHAS, AdaHeat. In particular, we proposed an MPC approach that enables our DHAS to work in conjunction with a diverse range of heating systems typically employed in domestic settings, such as: (i) heating systems with considerable thermal lags, (ii) heating systems with variable overall efficiency, and/or (iii) heating systems that exhibit a variability of the heating cost over time (both direct, i.e., variable energy prices, or indirect, i.e., due to the utilization of house-integrated intermittent energy resources), as per the respective generality and economic control requirements (see Section 1.1.3). To this end, we provided a new algorithm for

[14] In practice, these simulations can be carried out independently in separate units in each house/member, at the day's end. In the special case where the houses fall in the same region (which is expected if a micro-grid facilitates our approach) and have the same IER capacity, there is no need for such simulations. In that case, our heuristic collective advanced economic control reduces to independent advanced economic control, with the only difference being that the members share their energy before buying from the grid. As such, the cost of independent advanced economic control can be projected based on the realized consumption and IER production of each member.

[15] This choice is supported by the fact that the participation of such a member in a coalition that has spare generation (i.e., exports to the grid) leads to additional gains for the coalition. The risk of an opposite setting (i.e. there is no spare generation) can be managed through appropriately populating the θ_0 and θ_1 parameters and/or by choosing whether to accept such a member in the first place.

heating planning in the context of our MPC approach that fully exploits the probabilistic occupancy estimates (in fulfilling also the respective occupancy uncertainty requirement, see Section 1.1.3). Our algorithm is based on dynamic programming, enabling our system to work in conjunction with both linear and non-linear system models and arbitrarily complex optimization objectives (in contrast to other optimization methods typically used in the context of MPC for planning, such as quadratic and linear programming), enhancing the respective generality requirement (see Section 1.1.3). Moreover, regarding the particular optimization objective used in planning, the reliance on a single parameter for balancing heating cost and thermal discomfort enables our system to match to the user preferences through an effective adaptive procedure (as per the requirements: (i) minimal user-input, (ii) matching the user preferences and (iii) Pareto efficiency, see Section 1.1.3). In the context of our thermal comfort modeling approach, we utilized a cost-effective occupancy prediction algorithm developed by Scott et al. (2011) that limits the cost of our approach (as per the applicability requirement, see Section 1.1.3). Finally, our gray-box adaptive thermal modeling approach makes our system resilient and effective for employment in the highly dynamic thermal settings of houses, without facing the shortcomings of black-box and white-box approaches (as detailed in Section 2.2). As such, fulfilling the reliable thermal modeling requirement while respecting the applicability and generality requirements (see Section 1.1.3).

Furthermore, to account for the coalition potential in the case of advanced economic control, we detailed a scheme for collective advanced economic control (as per the respective coalition potential requirement, see Section 1.1.3). In this context, we formulated the optimization objective for optimal heating planning in the context of collective economic control and detailed its limitations in terms of deployment in real settings. Subsequently, we detailed our heuristic planning approach for collective advanced economic control that aims to overcome these limitations (as per the applicability requirement, see Section 1.1.3). Finally, we detailed our allocation mechanism to share the realized gains of the coalition that respects individual rationality and allocation efficiency (as per the respective cost allocation requirement, see Section 1.1.3). To conclude, in essence AdaHeat is a core framework where different system models can be incorporated to capture specific characteristics. In Chapter 4, along with our evaluation results, we provide a specific instantiation of our DHAS for our case study systems.

Chapter 4

Evaluating AdaHeat

In this chapter we provide a thorough evaluation of our domestic heating automation system (DHAS) approach, AdaHeat. In particular, in Section 4.1 we evaluate AdaHeat with respect to simple heating and simple economic control, while in Section 4.2 we evaluate AdaHeat with respect to advanced economic control. In both sections we: (i) describe the case study of the evaluation and how we collected the necessary data; (ii) describe the specific instantiation of AdaHeat with respect to the case study (iii) discuss our evaluation set-up and the instantiation of the benchmark systems; and (iv) report the evaluation results. Section 4.3 summarizes this chapter.

4.1 Simple Heating Control and Simple Economic Control

In this section we evaluate AdaHeat with respect to simple heating and simple economic control (see Section 1.1). We now proceed to describe the case study of our evaluation and how we collected the necessery data.

4.1.1 Case Study and Data Collection

For the case study of our evaluation of simple heating and simple economic control, we consider the living room of a family house in Cambridge, UK, (as seen in Figure 4.1) utilizing data coming from the original PreHeat deployment (Scott et al. (2011)). In particular, the house has both radiators (4 of 10 independently heatable rooms) and underfloor heating (8 of 10 rooms) and is equipped with custom hardware using .NET Gadgeteer to control the heating system and record data. In more detail, heating is controlled on a per-room basis through a *room unit* that controls wireless radiator valves, and a *control unit* that controls the underfloor heating valves, as shown in Figure 4.2(a)

FIGURE 4.1: Case study location (powered by Google maps).

(a) Room unit

(b) Control unit

FIGURE 4.2: Case study custom hardware (Scott et al. (2011)).

and Figure 4.2(b), respectively. In addition, the room unit has (i) an indoor air temperature sensor (Sensiron SHT15)[1] and (ii) a passive infra-red motion sensor for detecting occupancy (Panasonic PIR-AMN34111J),[2] while the control unit has no sensors but a relay for controlling the underfloor heating valves. Both units have 802.15.4 radio modules to establish a wireless mesh communication with a central server, where the data (further utilized in this work) are recorded. See Scott et al. (2011) for more details on this set-up and the original employment.

We choose the living room for our study as: (i) it is largely in use when the house is occupied and (ii) its thermal dynamics are particularly challenging due to its physical properties and household activity. In particular, it has two doors and three windows, as

[1] `www.farnell.com/datasheets/1563786.pdf`
[2] `www.panasonic-electric-works.com/eu/ds_61804_en_pir_motion_sensor.pdf`

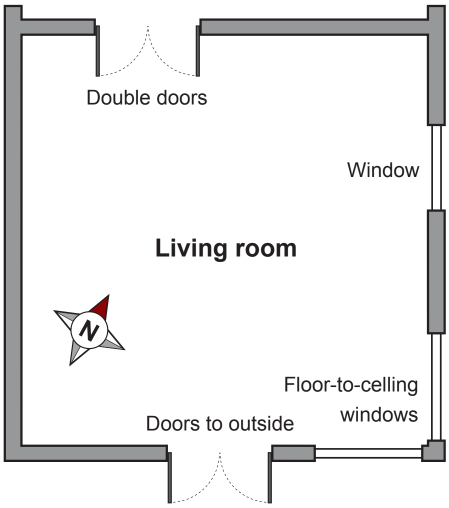

Figure 4.3: Layout of our case study living room.

seen in Figure 4.3, and it is equipped with an underfloor heating system (no radiators) and an auxiliary fan heater (that is occasionally used). Underfloor heating involves multiple heat transfer processes whereby heat is transferred from the source (i.e., the boiler) to an intermediate thermal mass (i.e., underfloor system), which then slowly leaks to its surroundings (e.g., air, ground, and house envelope). This process introduces considerable thermal lags, thus making this room interesting from a heating control perspective. In addition, the per-room-based heating in the adjacent rooms and the weather conditions can affect the indoor and outdoor thermal leakage. Furthermore, the occupants' activities have a substantial effect on the thermal dynamics of the room, specific examples of such events in our case include opening a window or operating the auxiliary fan heater. Taken together, these factors make this room a challenging testbed on the generality and efficiency of our approach both in terms of thermal modeling and control.

For the purpose of our research, we collected inside air temperature readings, T^{IN}, and occupancy events from November, 2011 to March, 2012 (150 days in total) via the .NET Gadgeteer hardware as discussed above. For the outside temperature, T^{OUT}, we use of publicly available data from the Cambridge computer laboratory (30 minute intervals).[3] Finally, to estimate solar radiation, G_T, we use the dataset from the EU joint research commission[4], which consists of estimated solar irradiance data on a typical day for a given month (note that this means that all days in any given month have the same solar irradiance). This data consists of average solar irradiance for 15 minute intervals and we assume this to be constant within this interval.

[3] We use linear interpolation where needed.

[4] `www.cl.cam.ac.uk/research/dtg/weather` and `http://re.jrc.ec.europa.eu/pvgis/apps4/pvest.php`

4.1.2 Instantiating AdaHeat

In this section we describe in detail the instantiation of our DHAS, AdaHeat, for controlling our case study underfloor heating system. In particular, in the following paragraphs we provide the instantiation of each of the principle components of our system, as detailed in Section 3.1 (i.e., the thermal comfort model, the thermal model, the heating cost model, and the controller).

- **Thermal Comfort Model:** As discussed in Section 3.1.1, any thermal discomfort experienced, at each instance that the house is occupied, considers the absolute deviation of T^{IN} from the user-provided, set-point temperature T^{SP}. In our case study the preferred set-point temperature is set to 22°C and, hence, T^{SP} in Equation 3.3 is set to 22°C. Moreover, for the occupancy schedule we utilized the occupancy data collected via the .NET Gadgeteer, as discussed in Section 4.1.1

- **Thermal Model:** Thermal modeling is highly dependent on the thermal process that is being modeled (as discussed in Section 3.1.2). Here, we identify the most suitable thermal model for our case study underfloor heating system by starting with the simplest feasible model and iteratively refining it into a more complex one. By doing so, we derive a thermal model where the transfer of heat from the heat source to the indoor air is assumed to be via an intermediate thermal mass (underfloor system), and the transfer of heat to the outside is via the house envelope. Thus, our thermal model captures the thermal lags of this system (in contrast to the simple thermal models originally employed by both PreHeat and SPOT+, see Section 2.7). Moreover, our thermal model also accounts for the effects of solar radiation on the indoor air temperature and the house envelope temperature. In more detail, the derived thermal model is:

$$T_{t+1}^{FL} = T_t^{FL} + r^h a + \phi_a(T_t^{IN} - T_t^{FL}) \tag{4.1}$$
$$T_{t+1}^{IN} = T_t^{IN} + r_a^s G_T + \phi_a(T^{FL} - T_t^{IN}) + \phi_b(T_t^{EN} - T_t^{IN}) \tag{4.2}$$
$$T_{t+1}^{EN} = T_t^{EN} + r_e^s G_T + \phi_b(T_t^{IN} - T_t^{EN}) + \phi_c(T^{OUT} - T_t^{EN}) \tag{4.3}$$

  where t stands for time, and T^{FL}, T^{IN} and T^{EN} stand for the floor-mass temperature, the inside air temperature and the house envelope temperature, respectively, and consider the thermal state of the system, i.e., $\mathbf{x} = \begin{pmatrix} T^{FL} \\ T^{IN} \\ T^{EN} \end{pmatrix}$. Furthermore, G_T and T^{OUT} represent the global solar irradiance and the outside temperature, respectively, and consider the vector of exogenous processes variables, i.e., $\mathbf{i} = \begin{pmatrix} T^{OUT} \\ G_T \end{pmatrix}$. In addition, ϕ_a, ϕ_b, and ϕ_c represent the leakage rates (between the floor mass and the indoor air, the indoor air and the house envelope, and the house envelope and the outside environment, respectively),[5] and r^h, r_a^s and r_e^s represent additional coefficients that aim to capture the effect of the heating system output to the inside air temperature, the effect of solar radiation on the inside air temperature, and the effect of solar radiation on the envelope temperature, respectively. Now, the leakage rates, ϕ_a, ϕ_b, and ϕ_c, along with the coefficients r^h, r_a^s and r_e^s, consider the equivalent thermal parameters that are assumed to be time-varying and need to be estimated. Finally, $a \in \{1, 0\}$ (on/off) is the heating control action and trivially considers $\mathbf{a}$.

Now, T^{FL} and T^{EN} consider non-observable thermal state variables that need to be estimated along with the equivalent thermal parameters. Note here, that, in essence, the above defined thermal model is affine with respect to the thermal state vector, $\mathbf{x}$ (or linear with respect to an appropriately extended thermal state vector). However, the simultaneous estimation of both parameters and state variables yields a non-linear problem in general (as discussed in Section 3.1.2). To this end, as is common practice, we use an extended Kalman filter (EKF) for the joint estimation of state and parameter variables (Grewal and Andrews (2011))—for details on joint state and parameter estimation with an extended Kalman filter see Fux et al. (2014). We evaluated our procedure over the 150 days dataset (collected as described in Section 4.1.1) to achieve the 95th percentile of the absolute prediction error to be 0.95°C and 1.37°C for 2 and 4 hours predictions, respectively.

- **Heating Cost Model:** Regarding the heating cost model, the particular heating system efficiency, as well as the energy provided to the space for a particular heating action, can generally be set to arbitrary values in the case of simple heating and simple economic control, as long as any needed ratios are retained. This is due to our adaptive approach in meeting the user preferences, as discussed in Section 3.1.3. As such, the instantiation of our heating cost model considers Equation 3.10 where we have set $C_{eff} = 1$ and $\mathrm{Pwr}(a) = a$ (i.e., $\mathrm{Pwr}(0) = 0$ and $\mathrm{Pwr}(1) = 1$).[6]

- **Controller:** Regarding the controller, the planning horizon and the planning interval length of AdaHeat were set to 1 hour ahead and 5 minutes, respectively (i.e., $\delta = 5$ min and $|H| = 12$). In particular, those MPC design characteristics have been found to be effective for efficient heating control, after experimenting with various design characteristics.

4.1.3 Experimental Setup

In this work, we evaluate AdaHeat *with* and *without* adaptive thermal modeling. We do so, to identify the benefits of such modeling in our DHAS. Moreover, we compare

[5] With respect to a RC-network representation (Deng et al. (2010)), a leakage rate considers a cumulative representation of thermal capacitance, C_{th}, and thermal resistance, R_{th} (i.e., $\phi = \frac{1}{C_{th} R_{th}}$).

[6] These particular parameters are the same for all systems evaluated in this work and, as such, their value does not alter the conclusions of this evaluation.

against the well-known SPOT+ and PreHeat which, essentially, employ MPC along with heuristic planning (see Section 2.7). As such, our evaluation can provide significant insights about the trade-off between heuristic and optimal planning, in the context of MPC. Now, although these DHASs employ simple fixed thermal modeling, we also evaluate them with a more advanced *fixed* model (that captures the thermal lags of the case study system) and with our *adaptive* model. In addition, we use the same occupancy prediction algorithm (i.e., Scott et al. (2011)); cost and discomfort metrics (i.e., Equation 3.10 and 3.3 respectively); and planning horizon and interval length (i.e., $\delta = 5$ min and $|H| = 12$) for all systems. We do so for two reasons: (i) to identify the benefits of adaptive modeling in various DHASs; and (ii) to compare various DHASs without being affected by any model and design differences. Moreover, we evaluate all DHASs *with* and *without* considering variable energy cost in order to characterize them in different settings. That said, our case study of an underfloor heating system with variable energy cost is a worst case scenario system and its efficient control can confirm (or disprove) the intended generality of AdaHeat. For completeness, we also evaluate the performance of three simple heating strategies: (i) always-on, which retains T^{IN} at T^{SP} throughout the whole day, (ii) always-off, in which heating is always off, and (iii) reactive, in which heating responds to occupancy (this is equivalent to a strategy where heating is manually switched on and off, when the occupants leave and return to the house, respectively). In more detail, the aims of this evaluation are:

- To identify the benefits of incorporating adaptive gray-box thermal modeling in different DHAS approaches.
- To identify the trade-off between heuristic planning and a planning approach that fully exploits the probabilistic occupancy estimates in the context of MPC (without being affected by any modeling and design differences of the DHASs considered).
- To provide a comprehensive comparison of different DHASs in different operational settings (also without being affected by any modeling and design differences).

In more detail, we evaluate all DHASs for a typical *winter* day (of February 2011), ensuring (via an iterative procedure) that the initial and final thermal state, $\mathbf{x}$, at the beginning and at the end of the day respectively, are the same in all our experiments, as seen in Figure 4.4. As such, our evaluation results consider *long-term average performance evaluation*, assuming that the same day repeats over time (i.e., same occupancy schedule, environmental conditions and predictions). We followed this procedure to provide *long-term performance estimates* for *various* DHAS parameter settings within feasible computational time. In particular, by doing so, we were able to evaluate all DHASs for a wide range of parameters and identify their performance in meeting the user preferences. In more detail, we evaluate SPOT+ for all combinations of a weighting factor within (0,1) with a step of 0.01, and a threshold value within (0,1) with step 0.1. In addition,

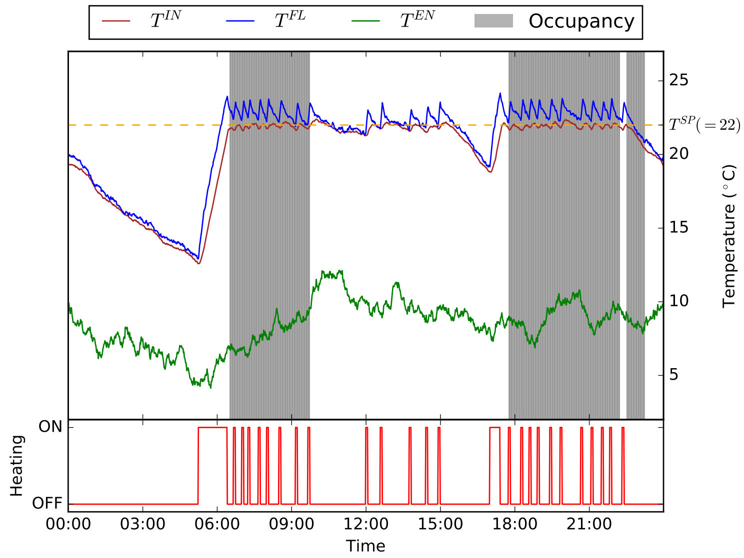

(a) AdaHeat heating control, $\lambda = 0.99$

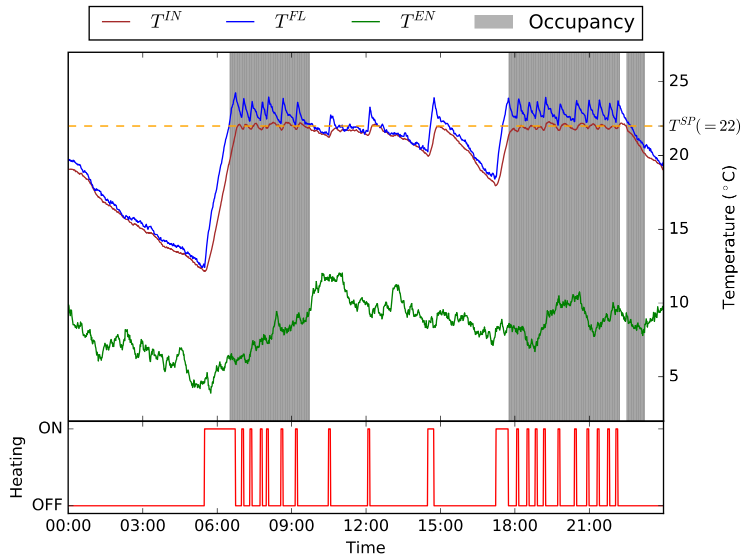

(b) AdaHeat heating control, $\lambda = 0.96$

FIGURE 4.4: AdaHeat heating control instances.

we evaluate PreHeat and AdaHeat for the same threshold and weighting factor range, respectively. We note here that the SPOT+ objective has been normalized to work with a weighting parameter in the range (0,1) *without any performance loss* to reinforce our comparison. In more detail, the objective of SPOT+, as discussed in Section 2.7, is given by Equation 2.2, reported here again (for concreteness):

$$J(\cdot) = \sum_{\tau=1}^{|H|} \mathrm{Cost}(\cdot_\tau) + O_\tau \lambda' \, \mathrm{Disc}(\cdot_\tau)$$

where $\lambda' \in [0, \infty)$. That being said, by setting $\lambda' = \frac{\lambda}{1-\lambda}$, λ' remains in $[0, \infty)$ as long as $\lambda \in [0, 1)$. Moreover, in general, minimizing $J(\cdot)$ is equivalent to minimizing $cJ(\cdot)$ for any $c > 0$. As such, and given that $0 \leq \lambda < 1 \Rightarrow 1 \geq (1 - \lambda) > 0$, the objective of SPOT+ can be transformed, without any performance loss, to:

$$J(\cdot) = \sum_{\tau=1}^{|H|} (1 - \lambda)\,\mathrm{Cost}(\cdot_\tau) + O_\tau \lambda\,\mathrm{Disc}(\cdot_\tau) \tag{4.4}$$

where $\lambda \in [0, 1)$. As such, we have derived an equivalent objective function based on a double-bounded parameter.[7]

Now, we chose a week-day in winter due to the heating needs of the particular season and to avoid any week-end day peculiar features.[8] To this end, we used the collected data for the ground truth of the occupancy schedule (and derived respective occupancy predictions for this day based on historical data according to Scott et al. (2011)—see Section 2.4) and the weather conditions (see Section 4.1.1).[9] Furthermore, in order to simulate our thermal model inaccuracies, we simulated the underlying thermal process by sampling $\mathbf{x}$, at each instance, from the respective EKF derived distributions. As such, the thermal model is not completely accurate with respect to our simulation, making our experiments more realistic.

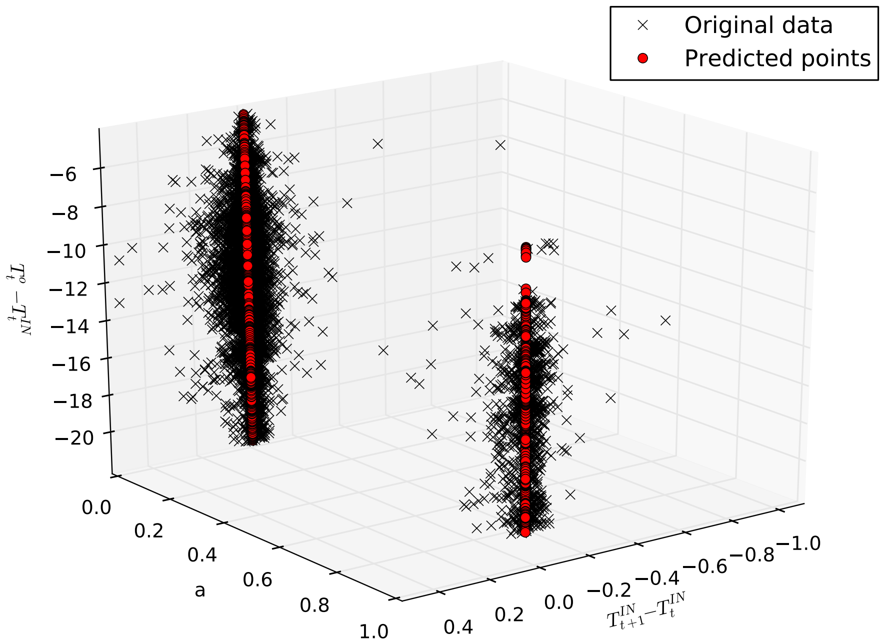

Figure 4.5: Least squares fitting of SPOT+ thermal model.

[7]Note that normalized cost and discomfort values can be used.
[8]For instance if the house is unoccupied during a week-end day, due to a trip, there will be zero potential savings.
[9]We linear interpolate whenever needed.

As outlined above, we evaluate SPOT+ and PreHeat also with their original fixed thermal models. Thus, as proposed in the respective publications, we estimated SPOT+ model via least squares regression and PreHeat's heat-rate as a historical average. In particular, we estimated both models based on the two first months of the 150-days dataset—thereafter, the estimated equivalent thermal parameters are fixed. In more detail, the SPOT+ thermal model (Equation 2.1) is equivalent to:

$$
\begin{aligned}
T_{t+1}^{IN} &= T_t^{IN} + \frac{e}{C}P_t^{hvac} - \frac{k}{C}(T_t^{IN} - T_t^{OUT}) \Rightarrow \\
(T_{t+1}^{IN} - T_t^{IN}) &= +\frac{e}{C}P_t^{hvac} + \frac{k}{C}(T_t^{OUT} - T_t^{IN})
\end{aligned}
\tag{4.5}
$$

where e/C and k/C are the linear regression coefficients. Note that in our case study $C^{Eff} = 1$ and $\mathrm{Pwr}(a) = a$ (see Section 4.1.2) and hence, $P^{hvac} = a$. As such, the dependent variable is $(T_{t+1}^{IN} - T_t^{IN})$ and the independent ones are a and $(T_t^o - T_t^{IN})$. Now, the derived ETPs are $e/C \simeq 0.043$ and $k/C \simeq 0.00043$ for SPOT+ and a heat-rate of ~ 0.441°C/hr for PreHeat. For concreteness, Figure 4.5 illustrates the predicted thermal response according to the thermal model of SPOT+ against the corresponding training dataset. Note, that the data points have a considerable variability over the $(T_{t+1}^{IN} - T_t^{IN})$ axis which is mainly due to the thermal lags of the system. The simple SPOT+ thermal model is not able to consider the thermal lags of the case study heating system which is illustrated by its inability to capture this variability. Now, in order to evaluate the DHASs with a more advanced fixed model, we use the equivalent thermal parameters of our adaptive model as derived exactly 30 days before the evaluation day (as such, the last model "calibration" is done, approximately, one month ago).[10] Finally, we note that in our evaluation, discomfort and cost are estimated with a numerical evaluation of an *one-min* interval.

4.1.4 Evaluation Results

In the context of this work, we benchmark AdaHeat against both PreHeat and SPOT+ using their original thermal modeling approaches. As we further detail in the rest of this section, PreHeat is not very sensitive to the accuracy of the thermal model used (due to the simple heating control strategy utilized), in contrast to both AdaHeat and SPOT+. That being said, the original simple thermal model of SPOT+ is not able to capture the thermal lags of the case study underfloor heating system. As such, when a planning horizon of one hour ahead is used, SPOT+ is not able to execute any heating schedule other than the Always-off heating strategy. Hence, only for this experiment we consider SPOT+ with an extended planning horizon of *two* hours ahead (in contrast to AdaHeat and PreHeat where *one*-hour planning horizon is used). Figure 4.6 illustrates the corresponding evaluation results.

[10] Although this simple technique is used to "approximate" a fixed thermal model, estimation techniques for fixed parameters can potentially demonstrate higher accuracy (Keesman (2011)).

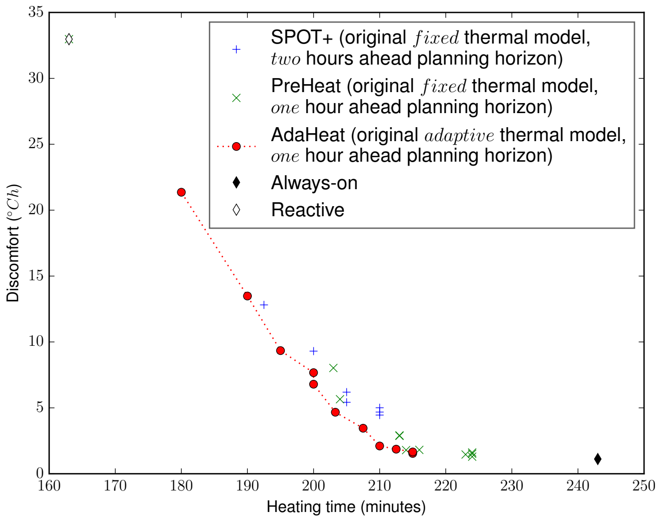

FIGURE 4.6: Initial evaluation results (Always-off wields ∼ 154°Ch discomfort).

From this we can see that AdaHeat has a better performance, in terms of Pareto efficiency, compared to both SPOT+ (with two hours ahead planning horizon) and PreHeat systems (while SPOT+ and PreHeat have a comparable efficiency). In particular, the balancing points captured by AdaHeat fall closer to the origin compared to SPOT+ and PreHeat. Moreover, AdaHeat demonstrates a better distribution of the obtained solutions compared to both SPOT+ and PreHeat as it is able to capture a wider and more evenly distributed solution set. However, this experiment is not very informative on whether this is due to the different thermal models employed by these systems, or due to the different control and/or planning approaches utilized. It is worth noting though, that none of the systems' solutions are dominated by the simple heating strategies Always-on, Always-off and Reactive and, hence, all systems can improve heating system efficiency compared to these simple strategies.

Given these initial observations, we proceed with a more comprehensive evaluation of the above systems. In particular, we first evaluate all systems with our thermal modeling approach, both fixed and adaptive (as discussed in Section 4.1.3), without considering energy cost variability (see Figure 4.7). As expected, adaptive thermal modeling significantly improves the efficiency of DHASs. In particular, as seen in Figure 4.8, both AdaHeat and SPOT+ are highly dependent on the accuracy of the thermal model employed and their performance improves significantly when adaptive thermal modeling is used. This is especially the case when low thermal discomfort values are intended. In particular, it can be seen in Figures 4.8(b) and 4.8(c) that the solutions captured with adaptive modeling fall closer to the origin compared to the fixed modeling ones. That said, in essence, SPOT+ is a non-domestic system developed for office buildings and relies on a simple fixed thermal model estimated through linear least squares fitting. However, the thermal characteristics of houses are much more dynamic than those in

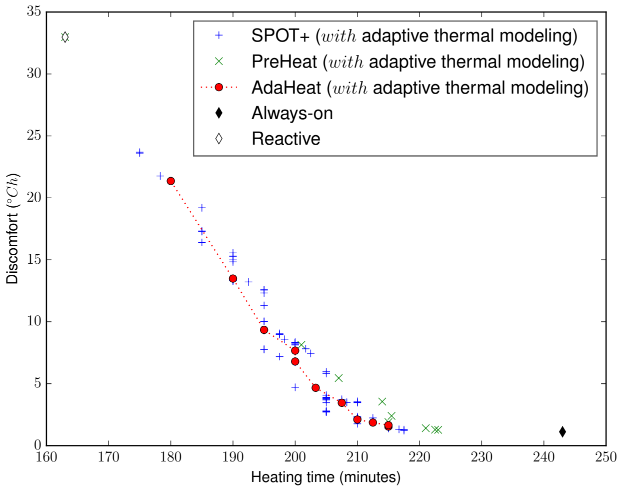

(a) *with* adaptive thermal modeling

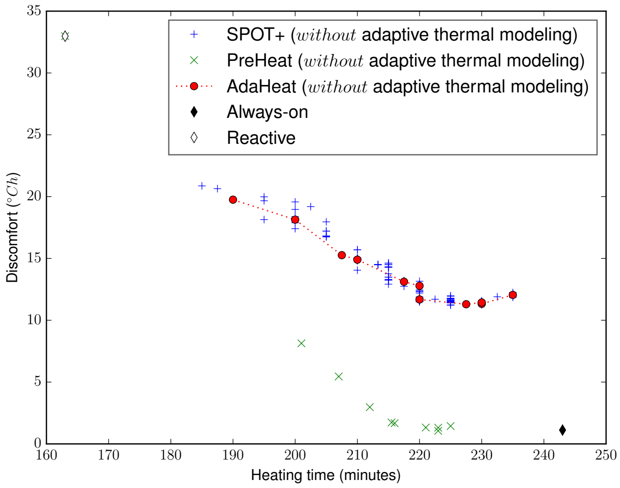

(b) *without* adaptive thermal modeling

FIGURE 4.7: Comprehensive evaluation results (without energy cost variability). "Always-off" wields ∼ 154°Ch discomfort (and, appropriately, 0 cost).

non-domestic buildings and are highly affected by several less tangible factors, such as the occupants' activity and the temperature in adjacent rooms (or buildings, see Section 2.2). Hence, the efficiency of SPOT+ (and AdaHeat) significantly improves when adaptive thermal modeling is considered. On the other hand, as seen in Figure 4.8(a), PreHeat is less sensitive to the accuracy of the thermal model employed due to its simple heating control strategy. However, this simple strategy deteriorates in terms of flexibility and efficiency (as discussed below). In general though, none of the systems' solutions are dominated by the simple heating strategies: Always-on, Always-off and Reactive, and, hence, all systems can potentially improve heating system efficiency compared to these strategies, even when fixed thermal modeling is incorporated into SPOT+ or AdaHeat.

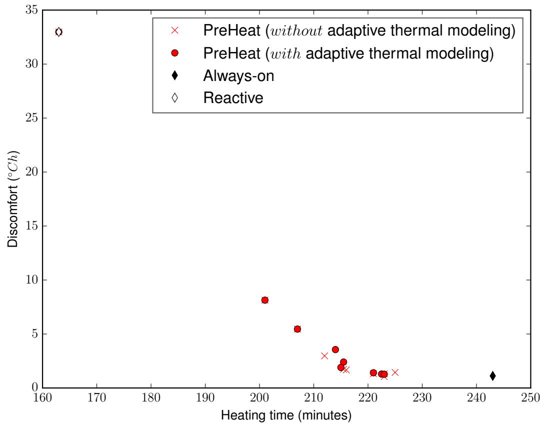

(a) PreHeat

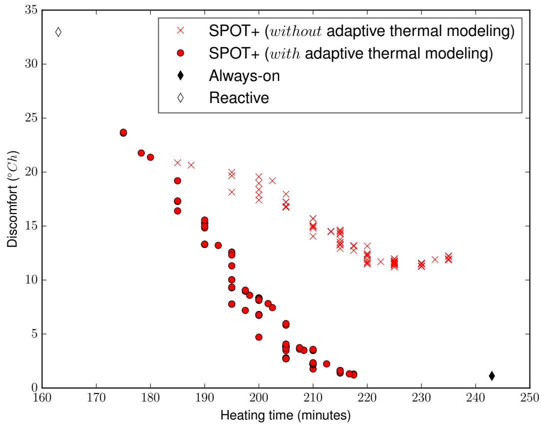

(b) SPOT+

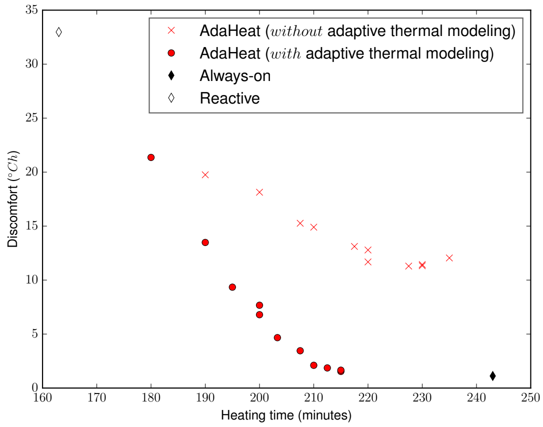

(c) AdaHeat

FIGURE 4.8: System by system comprehensive evaluation results (without energy cost variability). "Always-off" wields ∼ 154°Ch disc. (and 0 cost).

As far as the evaluated heating automation systems are considered with adaptive thermal modeling (Figure 4.7(a)), the evaluation results suggest that SPOT+ and AdaHeat have comparable Pareto efficiency. In general though, PreHeat demonstrates a slightly worse efficiency than SPOT+ and AdaHeat. This is due to its simple heating control strategy which is not generally able to capture heating systems with considerable thermal lags, such as the underfloor heating system considered in our case study, in a maximally efficient manner (see Sections 2.1 and 2.7). Moreover, SPOT+ demonstrates a less stable performance, in terms of Pareto efficiency compared to AdaHeat. This is, the solutions captured by SPOT+ are sometimes dominated by AdaHeat and vice-versa. On further investigation, SPOT+ is observed to occasionally *plan* a suboptimal heating schedule as seen in Figure 4.9.[11] However, the suboptimal planning of SPOT+ occasionally leads to higher or lower Pareto efficiency, as the MPC is not an optimal control approach (see Section 2.6).

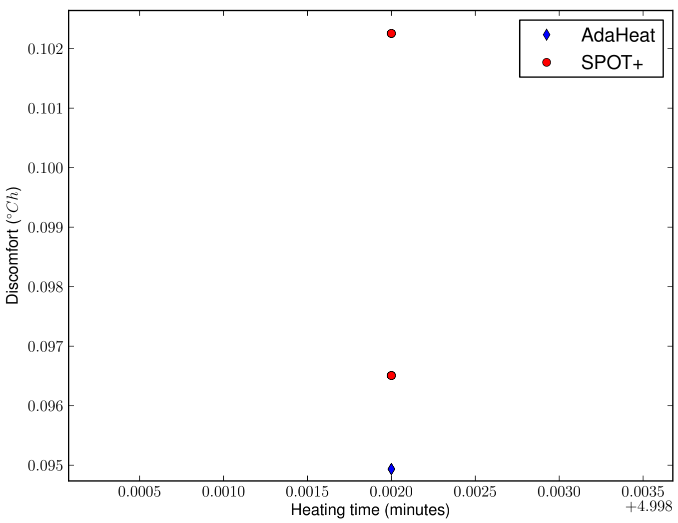

FIGURE 4.9: Planning instance. Note that SPOT+ solutions are dominated by AdaHeat solution.

Now, matching the, potentially time-varying, occupant preferences in balancing discomfort and cost is crucial in the context of DHASs (as discussed in Sections 2.7). To this end, SPOT+ relies on two user-provided parameters, i.e., the weighting factor and the threshold over the probabilistic occupancy estimates (see Section 2.7). However, in general, mathematical relationships between heating cost and quantifications of thermal discomfort are hard to comprehend for the users. As such, the usability of SPOT+ in domestic settings is questionable due to the complicated relationship between the threshold and the weighting parameter (see Figure 4.10). In more detail, many SPOT+ solutions (for different weighting and threshold parameters) are dominated by other solutions that

[11]This fact suggests that the non-closed form formalization of SPOT+'s planning objective (see Gao and Keshav (2013a)) is not a sufficient condition for Pareto optimality over cost and expected discomfort. However, we cannot conclude, whether it is a necessary condition just from these observations.

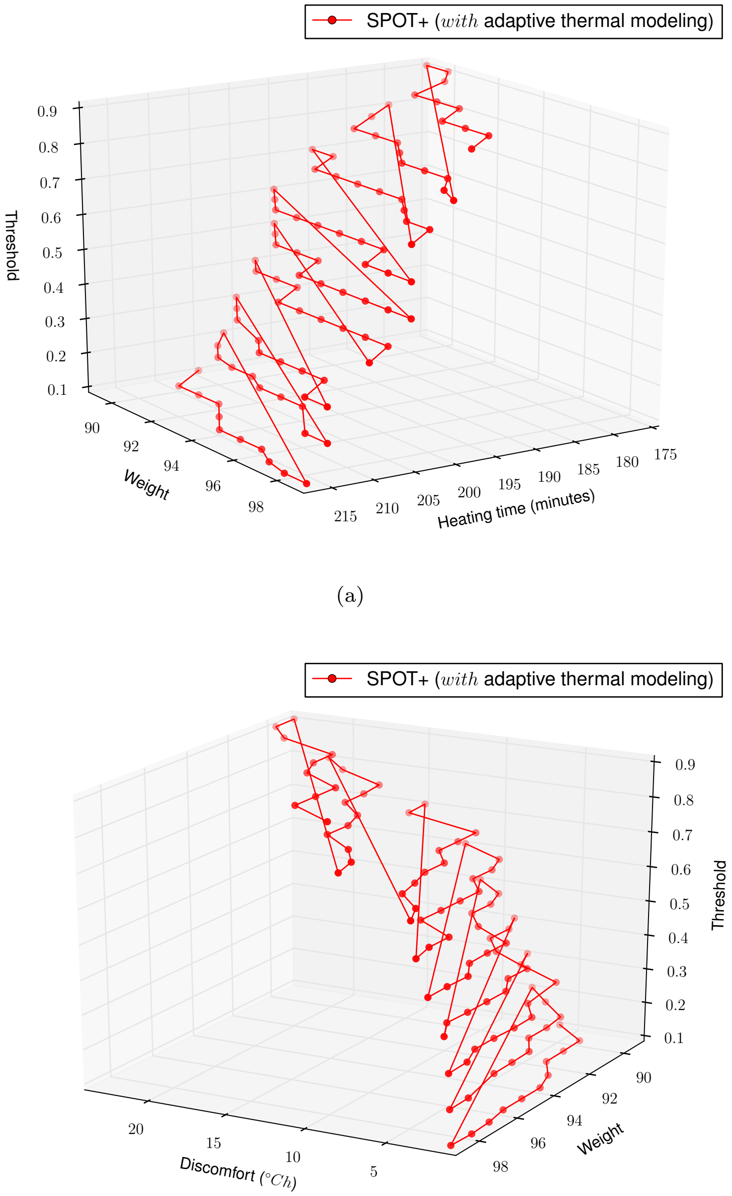

(a)

(b)

FIGURE 4.10: SPOT+ (*with* adaptive thermal modeling), balancing heating cost and thermal discomfort.

SPOT+ captures with different parameter choices. However, the exact performance of SPOT+ cannot be known in advance and we are not able to find any algorithm to appropriately populate the weight and the threshold parameter that can demonstrate a monotonic relationship with either the discomfort or the cost. For instance, one such algorithm could be to increase weight and threshold iteratively, starting from a particular weight for each threshold choice. This fact makes the parameter choice tricky as the user cannot know what to expect from different parameter value combinations.

On the other hand, both AdaHeat and PreHeat rely on only a single parameter for balancing heating cost and thermal discomfort. Moreover, the adjustable parameter of both AdaHeat and PreHeat demonstrates a monotonic correlation to thermal discomfort (and to heating cost for AdaHeat), when adaptive thermal modeling is considered, as seen in Figures 4.11 and 4.12. This fact enables the adjustment of these variables through a simple, real-time, adaptive procedure, based on a single Boolean feedback from the user, as discussed in Section 3.1.4. As discussed above though, when adaptive thermal modeling is considered (Figure 4.7), PreHeat illustrates a slightly lower Pareto efficiency than SPOT+ and AdaHeat. Moreover, in general, PreHeat is not able to capture a wide range of balancing points between cost and discomfort that allows a variety of user preference schemes to be captured—in contrast to AdaHeat. In particular, PreHeat operates on only a small region in balancing cost and discomfort which is not sufficient for appropriate heating control in domestic settings. As such, the occupants need to also adjust the origin of the discomfort metric (i.e., the set-point temperature), along with

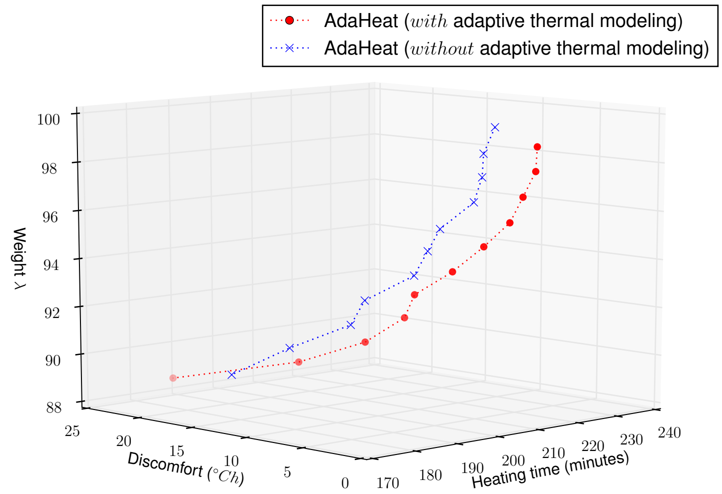

(a) AdaHeat balancing plot

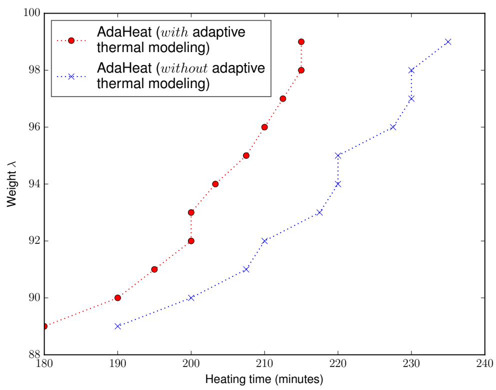

(b) Weight *versus* heating time (minutes)

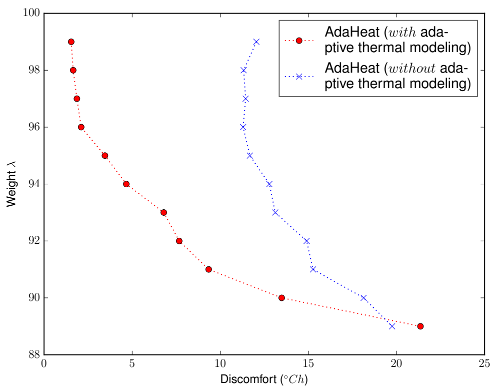

(c) Weight *versus* thermal discomfort (°C)

FIGURE 4.11: AdaHeat, balancing heating cost and thermal discomfort.

the threshold parameter, in order to meet their preferences. Thus AdaHeat is the only system that works sufficiently based on a single weighting parameter that can be learned on-line.

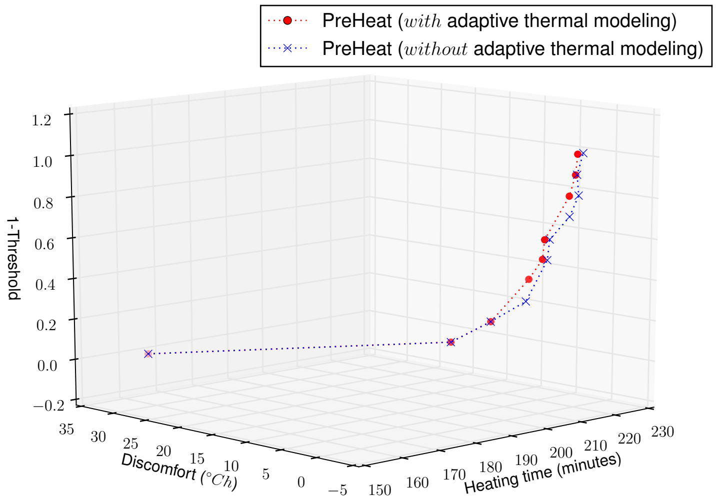

(a) PreHeat balancing plot

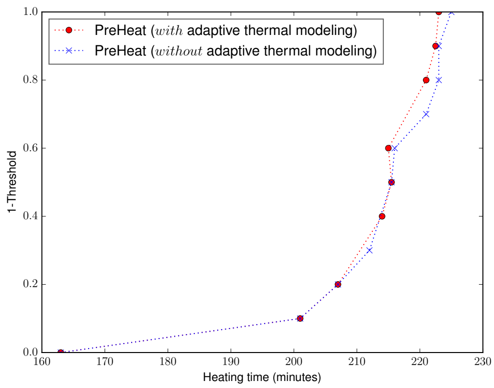

(b) Threshold *versus* heating time (minutes)

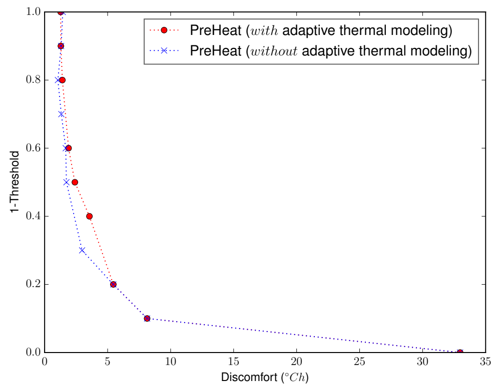

(c) Threshold *versus* thermal discomfort (°C)

FIGURE 4.12: PreHeat, balancing heating cost and thermal discomfort.

Lastly, the simple heating control strategy of PreHeat (i.e., heating for the minimum time required right before an occupancy event) does not allow this system to efficiently work in conjunction with heating systems that exhibit a variability of heating energy cost over time, time-varying overall efficiency or considerable thermal lags (as illustrated above). To further illustrate this we have conducted an additional experiment where arbitrarily variable energy prices have been assumed through the day. In particular, the energy prices have been designed to change every 5 minutes with their value being sampled from a uniform distribution within the range [1,10]. As can be seen in Figure 4.13, PreHeat's performance deteriorates significantly in this settings (both in terms of Pareto efficiency, and distribution and range of balancing points that it captures). Specifically, certain PreHeat solutions are dominated even by the Always-on strategy. Moreover, SPOT+ demonstrates significant variability over its performance for different parameter choices in this setting, as it captures many self-dominated solutions. In contrast, AdaHeat is generally stable in terms of Pareto efficiency, and generally smooth in terms of the distribution and range of solutions captured (Figure 4.13).

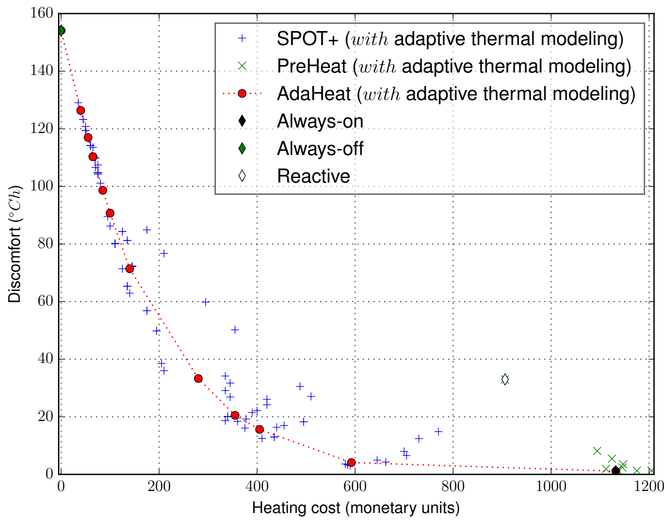

FIGURE 4.13: Evaluation results with energy cost variability

4.2 Advanced Economic Control

Here, we provide a thorough evaluation of AdaHeat with respect to advanced economic control. We now proceed to describe the evaluation case study and how we collected the necessary data.

4.2.1 Case Study and Data Collection

For our evaluation case study of advanced economic control, we modeled 30 houses in Mablethorpe, Lincolnshire, UK, (as seen in Figure 4.14) equipped (for simplicity) with the same wind turbine generator nominal capacity WTG^{nom} of 6kW and heating system nominal power HS^{nom} of 4kW. We choose the UK for our evaluation, since its weather climate leads to a generally heavy usage of space heating systems (while the specific UK location is chosen randomly). Due to the absence of concrete data for our particular region of interest, we modeled these houses by combining real data coming from different datasets (considering the UK), i.e., our case study considers fictional houses based on real data. In particular, we utilized a month of occupancy data (3/2012) from the dataset of the original PreHeat deployment (Scott et al. (2011)) considering a family house in Cambridge, UK. Furthermore, we utilized the dataset of the original MyJoulo (Rogers

FIGURE 4.14: Case study location (powered by Google maps).

et al. (2013)) deployment which considers one-week indoor and outdoor temperature readings of hundreds of houses across the UK, randomly choosing 30 for our research. Finally, we acquired 50 days (starting from 15/2/2011) worth of archival weather readings and forecasting reports from on-line providers,[12] for the location of our interest (i.e., Mablethorpe, UK).

4.2.2 Instantiating our Approach

Here, we detail the instantiation of AdaHeat for our case study. In particular, in the following paragraphs we provide the instantiation of each of the principle components of our system, as detailed in Section 3.1 (i.e., the thermal comfort model, the thermal model, the heating cost model, and the controller).

- **Thermal Comfort Model:** Regarding thermal comfort modeling, we set the set-point temperature, T_i^{SP}, of each house $i \in$ C to the actual one as captured in the MyJoulo dataset for all 30 houses (see Section 4.2.1). Moreover, for the occupancy schedule we utilized the one-month's worth of data of the PreHeat original deployment (see Section 4.2.1). In particular, we used each different day's realized occupancy and corresponding predictions to model the "ground truth" occupancy schedule and the occupancy predictions for each one of the 30 case study houses, respectively.

- **Thermal Model:** Since our case study heating systems do not experience considerable thermal lags, we employ a simple standard thermal modeling formulation,

[12] `www.uk.weather.com` and `www.wunderground.com`

where heat leaks from the house at a rate that is proportional to the temperature difference between the house, $T_{i,t}^{IN}$, and the outside environment, $T_{i,t}^{OUT}$ (Rogers et al. (2013); Bacher and Madsen (2011)). More formally, the thermal model employed for each house, i in coalition C, is:

$$T_{i,t+1}^{IN} = T_{i,t}^{IN} + \left(a_{i,t} r_i^{in} - r_i^{out}\left(T_{i,t}^{IN} - T_{i,t}^{OUT}\right) + c_i\right)\ \delta \tag{4.6}$$

where t is time, r_i^{in} stands for the heat provided by the heater; r_i^{out} for the leakage rate; $a_{i,t} \in \{1, 0\}$ for the heating control action (and trivially considers $\mathbf{a}$); and c_i is a heat bias to capture additional non-modeled heat transfer sources (e.g., incident solar radiation, occupant activity or adjacent buildings' temperature). In essence r_i^{in}, r_i^{out} and c_i consider time-varying parameters that need to be estimated. As discussed in Section 3.1.2, to this end, we utilize moving training window least-squares estimation (Soderstrom and Stoica (1989)). We trained our model using a moving training window least squares fitting on the 30 houses of the MyJoulo dataset (see Section 4.2.1). Our thermal modeling approach (Equation 4.6), can be easily transformed in (least squares form) as

$$\left(T_{i,t+1}^{IN} - T_{i,t}^{IN}\right) / \ \delta = a_{i,t} r_i^{in} + r_i^{out}\left(T_i^{OUT} - T_{i,\tau}^{IN}\right) + c_i \tag{4.7}$$

where $\frac{(T_{i,t+1}^{IN} - T_{i,t}^{IN})}{\delta}$ is the dependent variable/quantity, $a_{i,t}$ and $\left(T_i^{OUT} - T_{i,\tau}^{IN}\right)$ the independent ones, and r_i^{in}, r_i^{out} and c_i the time-varying coefficients to be estimated. For concreteness, we performed a simple evaluation of our approach, using the data for all 30 houses in our dataset. In particular, we used the last day for our evaluation with a moving training window of 6 days for all houses, and calculated the RMSE with respect to the predictive horizon. The results are reported in Figure 4.15(a). It can be seen that this simple approach provides adequate predictive accuracy for our MPC approach which is in alignment to the literature (Coley and Penman (1992)). Now, in order to capture the thermal modeling inaccuracies in our simulation experiments, we used the standard deviation of our fitted thermal models, as derived in the evaluation dataset for each house, to model a zero-mean Gaussian distribution. A sample from this distribution was added to the thermal model predictions, at each instance, to simulate the "ground truth" thermal response.

- **Heating Cost Model:** As discussed in Section 3.1.3, in the special case of advanced economic control, adequate estimates of the future intermittent energy resource (IER) power output are needed in the context of the respective heating cost modeling. Now, the shared IERs of our case study are wind turbine generators. To this end, we used the weather data considering the location of our interest (Section 4.2.1) and the RENES[13] simulation platform (Panagopoulos et al. (2012);

[13] www.intelligence.tuc.gr/renes

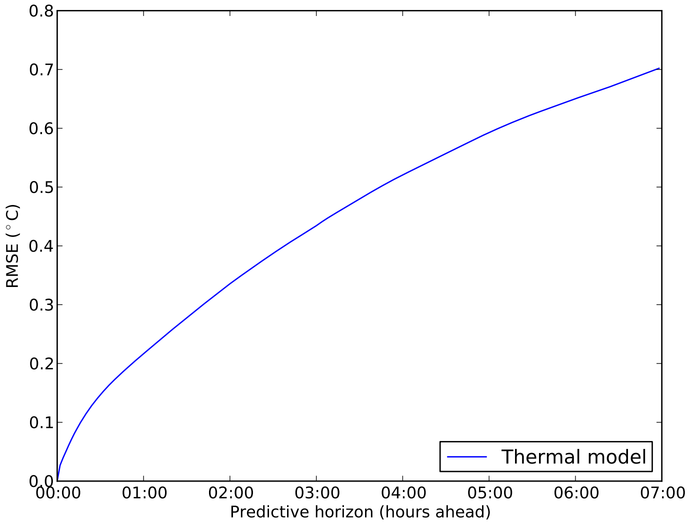

(a) Thermal model.

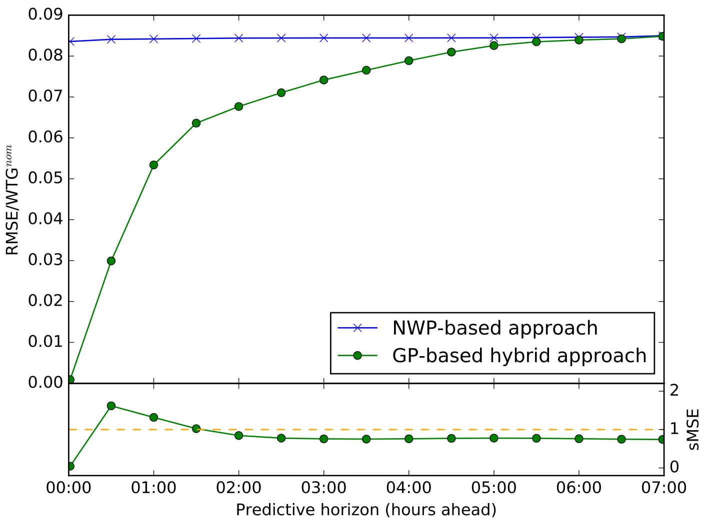

(b) Wind turbine generator power output prediction.

FIGURE 4.15: Instantiation evaluation results.

Panagopoulos (2013)) to create synthetic power output data. In more detail, we used the archival weather readings and the forecasting reports to derive power output "readings" and predictions, respectively. In essence, RENES uses a well-known sigmoid function to transform wind speed into power output (see Panagopoulos et al. (2012)). The sigmoid parameters that we use are the ones suggested by RENES, and are consistent with a typical commercial system. It should be noted here that since all houses are assumed to lie in the same region (and have the same WTG$^{\text{nom}}$), we use the same ground truth and derived predictions based on physical numerical weather prediction (NWP) (see Section 2.9), assuming the same sigmoid

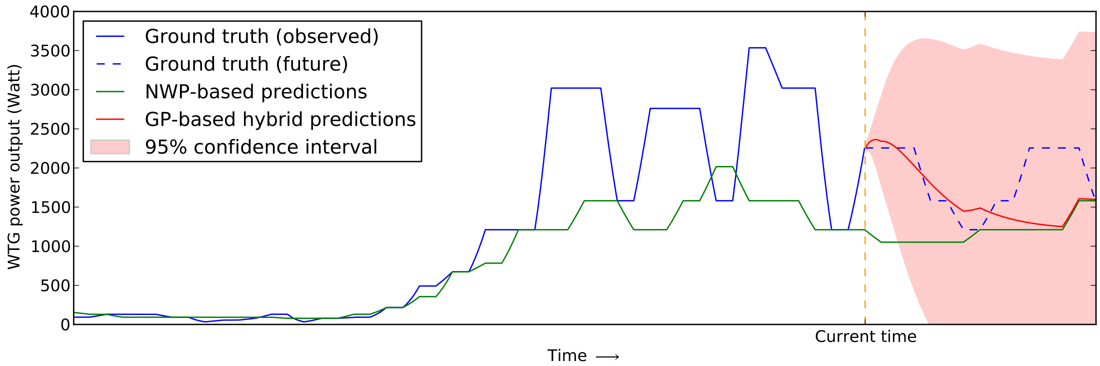

FIGURE 4.16: Instance of our GP-based hybrid wind turbine generator power output prediction approach.

parameters (for simplicity) for all IERs.

Now, as discussed in Section 2.9, hybrid approaches can improve the (short term) accuracy of NWP-based approaches due to their site-specific calibration, and GP regression is effective in this context, providing also stochastic estimates. As such, in this work we combine the derived NWP-based power output predictions with historical observations, based on adaptive GP regression. We utilize an adaptive approach to maximally account for up-to-date NWP-based predictions and historical observations. In particular, we fix the mean of the GP to follow the NWP-based predictions and learn the fluctuations of the NWP-based predictions compared to the ground truth, based on the historical observations, in real time. More formally, given a set of observations $(x_1, y_1), \ldots, (x_n, y_n)$ and NWP-based prediction points $(x_1, g_1), \ldots, (x_n, g_n), \ldots (x_m, g_m)$ (where m is usually greater than n), we take the difference between the two outputs, i.e., $z_i = y_i - g_i$, to obtain a new time series $(z_1, x_1), \ldots (z_n, x_n)$. This new series, defined over the same period as the historical observations $i < n$, essentially considers the fluctuation of the NWP-based predictions compared to the observations. Now, we further extrapolate these points using gradient descent in the context of a GP with a squared exponential kernel, i.e., $K_{\mathrm{SE}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left((\mathbf{x} - \mathbf{x}')/2l^2\right)$, which is suitable for modeling our smoothly changing NWP-based time series. Finally, in order to derive power output predictions, we add to the GP projected series the NWP-based prediction. This procedure is repeated in time, as new observations and new NWP-based predictions are acquired, shifting the training window (which considers a day's worth of data) and predictive horizon into the future. As such, we derive probabilistic estimates of the cumulative future power output, R'_C, at each instance, which follow a Gaussian distribution; $R'_C \sim N(\mu, \sigma)$. However, the power output of a wind turbine generator is bounded within $[0, r_C^{\max}]$. As such, we derive the corresponding truncated Gaussian distribution by bounding R'_C within this range.

Following the above, we are able to increase the accuracy of the NWP-based predictions through a simple and effective adaptive site-specific calibration technique. In particular, our approach significantly improves the prediction accuracy for a short-term horizon, while for longer horizons it progressively reduces to the NWP-based approach, as can be seen in Figure 4.15(b), where an instance is illustrated for 7 hours ahead predictive horizon. In order to evaluate our approach in a concrete manner we calculated the RMSE of our predictions with respect to the predictive horizon for the 50 days in our dataset (employing our approach every half an hour) modeling one of our case study wind turbine generators. In this context, we also calculated a standardized mean square error (sMSE) which considers the MSE over the GP predicted variance. This metric captures the accuracy of our GP variance modeling. In particular, values closer to 1.0 indicate that the modeled variance is consistent with the realized one, while values significantly over or under 1.0 mean that our approach's estimates are under or over confident, respectively. As can be seen in Figure 4.15(b), the results indicate that our hybrid GP approach significantly increases the accuracy of the NWP-based approach (especially for short term horizons) while the modeled variance closely follows the realized one—in Section 4.2.4 we also evaluate the significance of our hybrid approach within AdaHeat. It is worth noting that the performance of our simple hybrid GP-based approach is in alignment to the literature (Chen et al. (2014)) and also outperforms far more complex GP-based approaches (in terms of RMSE throughout the predictive horizon examined) evaluated in the proximity of our region of interest (and, in particular, in Ireland) (Yan et al. (2016, 2014)).

Now, since we consider advanced economic control, we used Equation 3.11 for our heating cost modeling. In this context, we assume that the case study electricity-based space heating systems do not consider heating pump technology. As such, we set $C_i^{Eff} = 1$ and $\mathrm{Pwr}_i\left(u_i, T_{i,t}^{IN}\right) = u_i \mathrm{HS}^{\mathrm{nom}}$ (i.e., $\mathrm{Pwr}_i\left(0, T_{i,t}^{IN}\right) = 0$, and $\mathrm{Pwr}_i\left(1, T_{i,t}^{IN}\right) = \mathrm{HS}^{\mathrm{nom}}$), for each house $i \in \mathrm{C}$.[14] Furthermore, following the P^{Buy} and P^{Sell} rates in the UK market in 2015, we set $P^{Buy} = 11.25$ p/kWh and $P^{Sell} = 4.5$ p/kWh (where p is British penny) which correspond to a P^{Buy}/P^{Sell} ratio of 2.5. That said, given $R_C' \sim N(\mu, \sigma)$, $R_C \in [0, r_C^{\max}]$, and the straightforward conditions that $r_C^{\max} > 0$, and $Cons_t(\cdot) \geq 0$, Equation 3.11 can be transformed in closed form (Appendix A).

- **Control and Planning:** Regarding control and planning, the slow nature of building thermal dynamics (Široký et al. (2011)) dictates a relatively long planning horizon and enable us to use a long planning interval. As such, we use a planning horizon of 3.5 hours with a 10-min interval. In particular, those MPC design characteristics have been found to be effective for efficient advanced economic control, after experimenting with various design characteristics.

[14] We note that setting $C_i^{Eff} = 1$ ensures that we do not favor our approach via modeling inefficient heating systems.

4.2.3 Experimental Setup

In our evaluation procedure, we compare our heuristic collective advanced economic control approach, in a contemporary market reality with flat import and export tariffs, against two benchmark approaches: (i) independent advanced economic control, where the houses optimize space heating independently with respect to their own IER generation, and (ii) simple economic/heating control where the houses optimize space heating (independently) disregarding their IER generation. As such, we evaluate the benefits of collective advanced economic control, compared to independent advanced economic control, as well as the benefits of independent advanced economic control compared to (independent) simple economic/heating control. Note here that in the case of flat import tariffs, simple economic control, that merely considers import tariffs (see Section 1.1), reduces to simple heating control, as there is no variation in the energy cost over time. Now, we evaluate all approaches for a typical winter day of February 2011 (randomly chosen from our weather dataset), ensuring (via an iterative procedure) that the initial and final thermal state, at the beginning and at the end of the day respectively, are the same for each house in all our experiments (as in Section 4.1.3). As such, our results consider long-term average performance evaluation, assuming that the same day repeats over time (i.e., same occupancy schedule, environmental conditions and predictions). By doing so, we are able to provide long term average evaluation results considering a wide range of coalition sizes (1-30 houses) and a a wide range of cost-discomfort balancing preferences for each house (sampling 100-10,000 different combinations, depending on the coalition size). In order to identify the potential improvement margin of our heuristic approach, we also evaluated collective advanced economic control considering optimal planning, for the case of a two-house coalition (where the respective optimization is feasible, as discussed in Section 3.2.1). In addition, in order to identify the benefits of incorporating our hybrid wind turbine generator predictive approach in AdaHeat, we also evaluated AdaHeat, for the case of a single house coalition, using the initial NWP-based predictions.[15]

4.2.4 Evaluation Results

Here, we report our evaluation results. In all figures reported in this section we use the following abbreviations:

- **AEC**: Advanced economic control
- **SEC**: Simple economic control
- **SHC**: Simple heating control

[15]Note here, that for single house coalitions, optimal and heuristic collective advanced economic control, as well as independent advanced economic control, become the same. Hence, we evaluate the benefits of our hybrid approach without being affected by any heuristic planning losses.

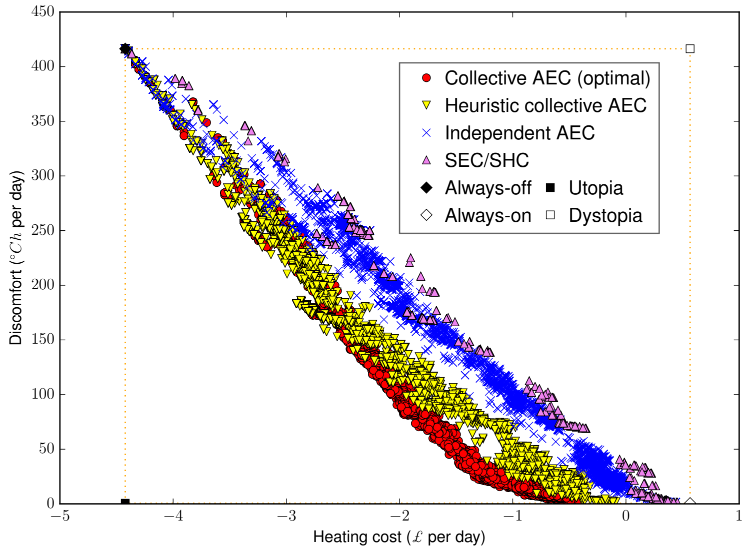

(a) Balancing points.

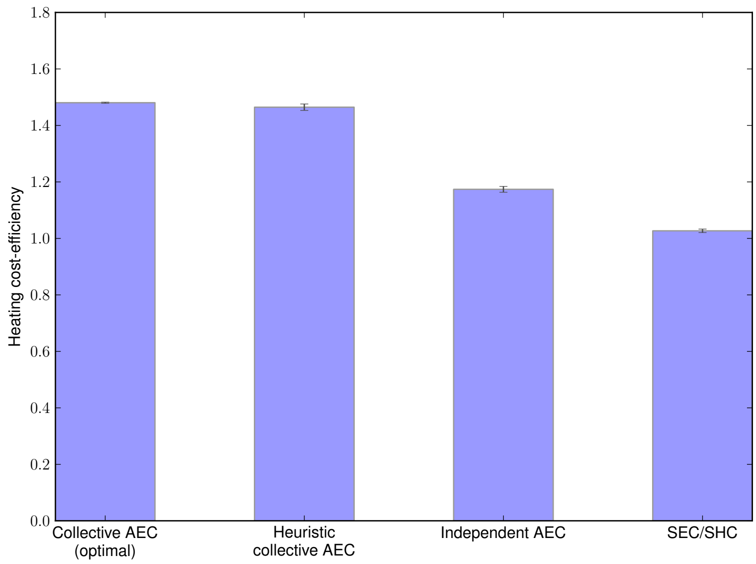

(b) Heating cost-efficiency of each approach.

FIGURE 4.17: Initial evaluation results (two-house coalition)

Figure 4.17(a) illustrates our evaluation results for a two-house coalition (where optimal planning can feasibly be used, as discussed in Section 3.2.1), considering the collective discomfort and cost of the coalition per day (negative "cost" indicates profit). We note here, that since this is a two-objective optimization, points closer to the origin indicate higher Pareto efficiency. Within this context, Figure 4.17(a) illustrates that our heuristic collective advanced economic control approach has considerably higher Pareto efficiency when compared to independent advanced economic control and even higher when compared to simple economic/heating control (denoted in all figures as SEC/SHC). Furthermore, independent advanced economic control demonstrates a consistently better Pareto efficiency compared to simple economic/heating control (even though not by a wide margin). These results are not surprising since further information is considered as we proceed to a more advanced system. In addition, our heuristic collective advanced economic control approach demonstrates a Pareto efficiency that is very close to the optimal one, which illustrates that there is not a big margin of potential improvement (at least in the case of a two-house coalition). That said, in order to compare the evaluated approaches in a more concrete manner we estimate the heating cost-efficiency of each approach. To do so, we evaluate two simple static timer heating policies; namely *always-on* and *always-off*, illustrated in Figure 4.17(a) (for concreteness). In essence, always-off is the heating policy where there is no heating at all, while always-on is the policy where the indoor temperature of each house is maintained as close as possible to the respective set-point temperature throughout the day (as discussed in Section 4.1.3). Intuitively, always-off considers a possible worst-case scenario in terms of discomfort, $\mathrm{Disc}^{\max}$, while it considers the best-case scenario in terms of heating cost, $\mathrm{Cost}^{\min}$. In contrast, always-on considers a possible worst-case scenario in terms of cost, $\mathrm{Cost}^{\max}$, and the best-case scenario in terms of discomfort, $\mathrm{Disc}^{\min}$. Now, we utilize these best-case and worst-case cost and discomfort values to estimate a utopia and a dystopia point of our approaches, respectively. These points, along with the points captured by the always-on and always-off policies create a 4-point imaginary window that frames all solution points of our approaches (as seen in Figure 4.17(a)). Given this window, we normalize the points captured by the evaluated approaches within the emerging ranges. Finally, using the supplementary fraction of the normalized discomfort (that essentially considers comfort) we calculate the mean comfort to cost ratio for each one of our approaches, which considers the heating cost-efficiency. More formally, given the set of all the solution points of an evaluated approach $\mathrm{A} := \{\mathrm{P}_1, \mathrm{P}_2, ..., \mathrm{P}_{|A|}\}$, where the coordinates of any given point P_x is a tuple $(\mathrm{Cost}_{\mathrm{P}_x}, \mathrm{Disc}_{\mathrm{P}_x})$ considering the respective cost and discomfort values, the heating cost-efficiency of the approach is calculated as $\frac{1}{|\mathrm{A}|}\sum_{i=1}^{|\mathrm{A}|} \frac{1-(\mathrm{Disc}_{\mathrm{P}_i} - \mathrm{Disc}^{\min})/(\mathrm{Disc}^{\max} - \mathrm{Disc}^{\min}))}{(\mathrm{Cost}_{\mathrm{P}_i} - \mathrm{Cost}^{\min})/(\mathrm{Cost}^{\max} - \mathrm{Cost}^{\min})}$. Figure 4.17(b) illustrates the heating cost-efficiency of each approach. It can be seen that the results are consistent with the observations in Figure 4.17(a) which renders heating cost-efficiency a valid comparison metric. The bars correspond to standard mean error capturing the diverse cost-discomfort preferences of the houses and indicate a statistically significant improvement in the heating cost-efficiency when moving from simple economic control to independent advanced economic control and, further on, to heuristic collective advanced economic control. Importantly for the efficiency of AdaHeat+, further t-test evaluation shows no statistically significant difference in the heating cost-efficiency between optimal and heuristic collective advanced economic control at the 0.05 significance level (with a p-value of $\sim$0.17).

Now, the above results indicate a significant improvement in the heating cost-efficiency of collective advanced economic control compared to independent advanced economic control (and even more when compared to simple economic/heating control) which can

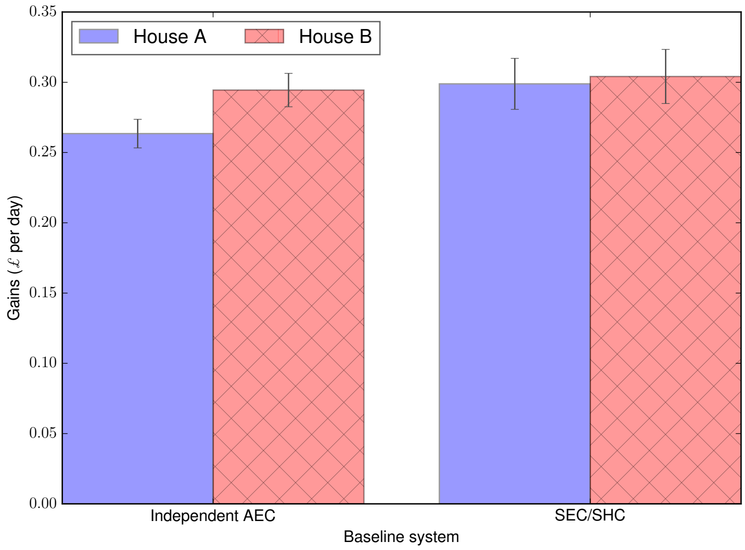

FIGURE 4.18: Two-house coalition monetary gains.

lead to considerable monetary gains for the coalition members. Figure 4.18 illustrates the monetary gains of using collective advanced economic control, for each one of the houses in the coalition (after the cost allocation), compared to: (i) independent advanced economic control and (ii) simple economic/heating control. As can be seen, house A and B gain over £0.20 and ∼£0.30 (respectively) on average, *each day*, by using collective advanced economic control compared to independent advanced economic control (and even more when compared to simple economic/heating control). These gains correspond to £5 − £10 per month, which is a considerable saving, if one considers that they come only from efficient heating control. The bars correspond to standard mean error capturing the diverse share of gains allocated to each house, as well as the diverse cost-discomfort preferences of the houses. We note here that these preferences vary from maximum interest in heating cost (with minimum interest in discomfort) to maximum interest in discomfort (with minimum interest in cost). As such, the realized absolute gains of advanced economic control fluctuate considerably along with the heating cost (since, for instance,

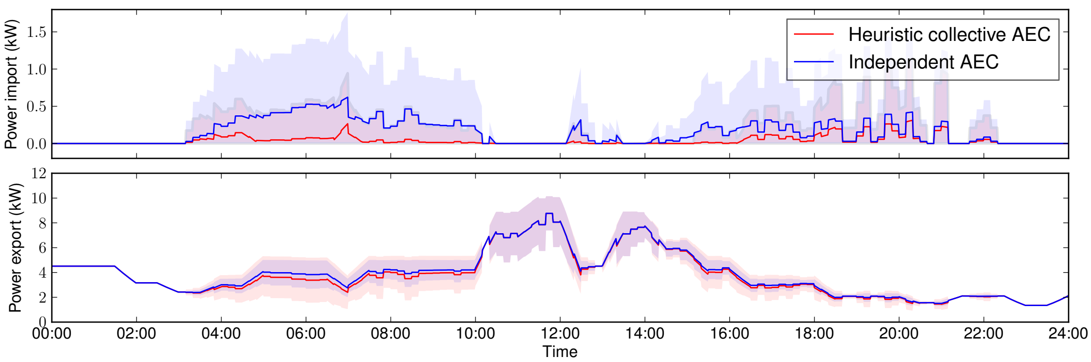

FIGURE 4.19: Two-house coalition energy exchange.

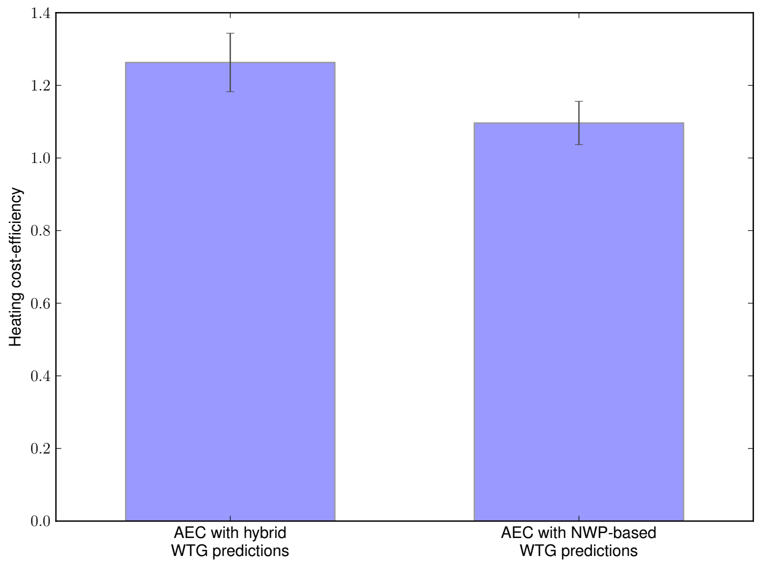

FIGURE 4.20: Hybrid predictions contribution.

there cannot be any gains if heating is switched off). These gains come along with providing respective regulation services to the grid (i.e., minimize the energy import as per the pricing motivation). Figure 4.19 illustrates the mean energy import and export (from and to the grid) for the two-house coalition—the shaded region considers (unbiased estimation) of standard deviation capturing the diverse cost-discomfort balancing preferences. As seen, the energy consumption flattens significantly (towards 0) even for a small coalition of only 2 houses, compared to independent advanced economic control. Now, as discussed in Section 4.2.2 here we also identify the benefits of incorporating our hybrid wind turbine generator predictive approach to AdaHeat. We do so, by evaluating our approach, for a single house coalition, both with NWP-based and our hybrid approach (Figure 4.20). Not surprisingly, we see that our approach increases the cost-efficiency of AdaHeat (although with statistical significance only in the 0.1 significance level, i.e., p-value of $\sim$0.1) as more accurate predictions are being considered.

As discussed in Section 3.2.2, AdaHeat also enables the coalition members to balance worst-case heating cost and thermal discomfort through a simple boolean feedback procedure. To demonstrate this, Figure. 4.21 illustrates worst-case cost and thermal discomfort against the balancing parameter λ for the two houses in our two-house coalition case study. As can be seen, both worst-case heating cost and thermal discomfort are in a monotonous relationship with λ (irrespectively of the λ population of the other house). This, in turn, suggests that a household can progressively adjust the single balancing parameter λ through a simple feedback procedure until the respective preferences are met. Nevertheless, as discussed in Sections 3.2.1 and 3.2.2, the realized heating cost depends also on the λ population of the other coalition members. To demonstrate this dependence, Figure 4.22 illustrates the realized cost and discomfort balancing of each house given a fixed $\lambda = 90$ and for various populations of the λ parameter of the other

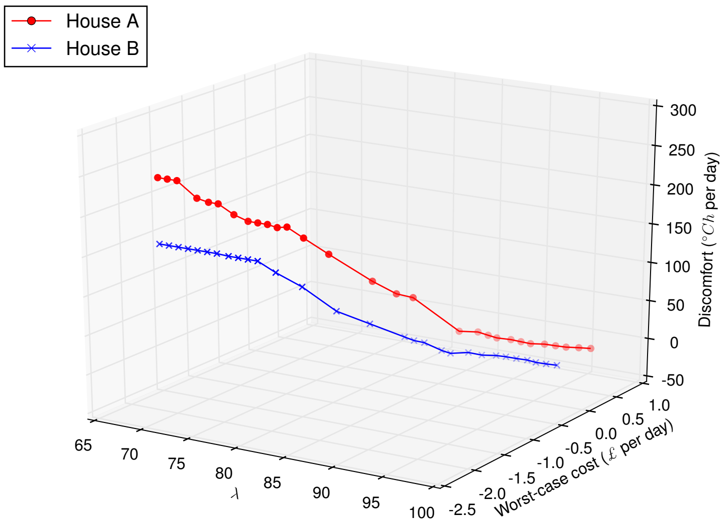

FIGURE 4.21: Balancing worst-case cost and discomfort.

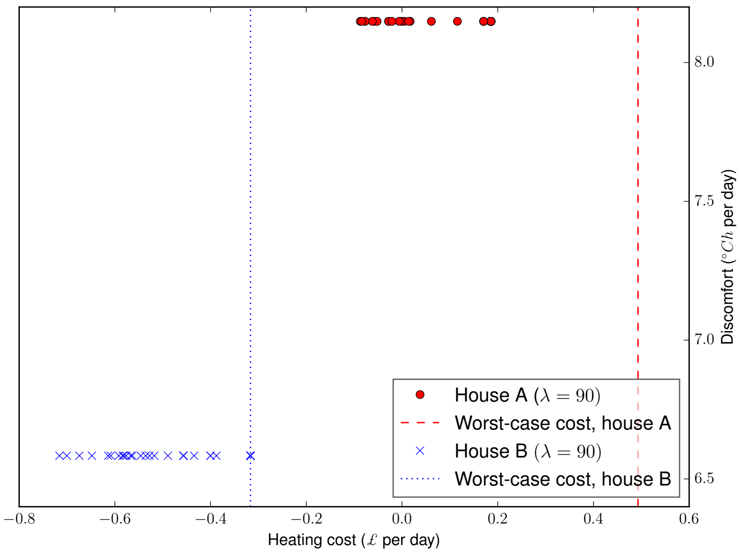

FIGURE 4.22: Fixed λ evaluation results.

house.[16] As can be seen, the realized cost depends on the λ population of the other house. Nevertheless, it is always lower (or equal) to the worst-case cost enabling the user to take an informed decision. Further evaluation of the user behavior with respect to this balancing requires a real world trial (that considers a future work direction, as further discussed in Chapter 7).

Regarding scalability, we evaluate how our approach scales with the size of the coalition. Figure 4.23 illustrates how the heating cost-efficiency scales with the size of the coalition up to 30 houses—the bars consider standard mean error corresponding to statistically significant differences, in the 0.05 significance level, among all methods evaluated for all

[16]For the effective range of λ for each house consult Fig. 4.21.

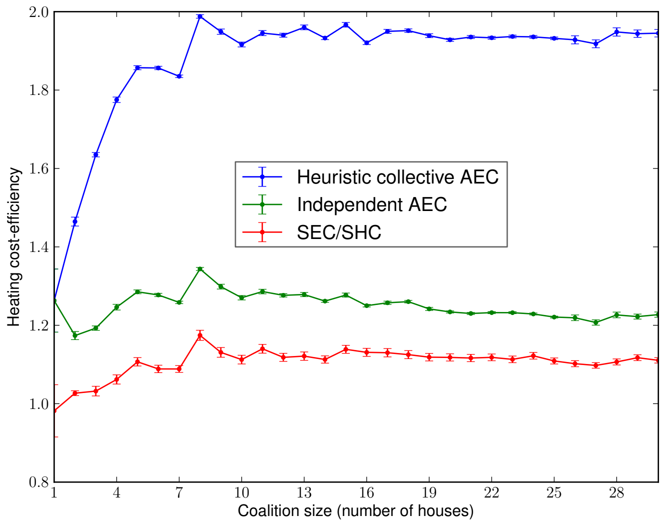

FIGURE 4.23: Heating cost-efficiency vs coalition size.

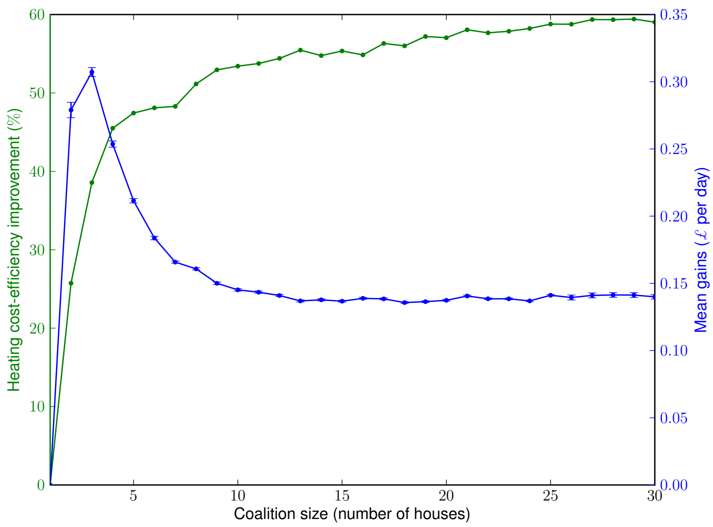

FIGURE 4.24: Collective compared to independent AEC.

coalition sizes considered (apart between heuristic collective and independent advanced economic control for the single-house coalition case where both methods are essentially the same, as discussed in Section 4.2.3). As can be seen the cost-efficiency of simple economic/heating control and independent advanced economic control remain relatively constant with the size of the coalition, as there is no energy exchange within the coalition. In contrast, the collective advanced economic control cost-efficiency increases rapidly with the size of the coalition until ∼9 houses and then remains generally constant. However, a careful investigation of the relative improvement of collective advanced economic control over independent advanced economic control in Figure 4.24 illustrates that the efficiency improvement flattens at around 27 houses. As can be seen in Figure 4.25, this is the size

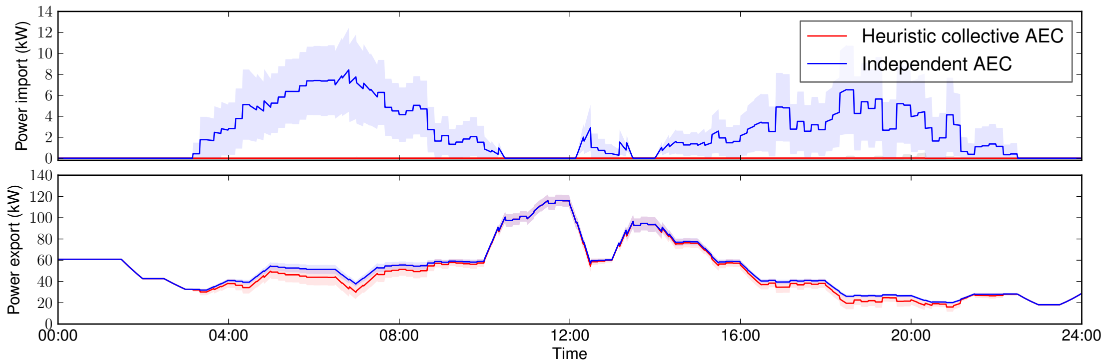

FIGURE 4.25: 27-house coalition energy exchange.

where the coalition is able, for the first time, to meet its heating consumption from its own IER capacity with negligible energy import (the shaded region considers standard deviation as per Figure 4.19). Figure 4.24 also illustrates the mean monetary gains of the houses, compared to independent advanced economic control (the bars consider standard mean error). We note here that the gains of every coalition member are positive in all cases (by definition) in accordance to our cost allocation mechanism that respects individual rationality (Section 3.2.3). As can be seen, the mean house gains flatten relatively fast with the size of the coalition while they peak at a coalition of 3 houses. The above observations are not surprising. In particular, the benefits of collective advanced economic control are already great for relatively small coalitions where the energy import is minimized considerably. Hence, as we move forward larger coalitions it is progressively only the energy import of a new member that is further minimized and, hence, it is progressively only the corresponding gains that are further allocated among an ever-growing number of coalition members. Now, the above illustrate that a coalition of 27

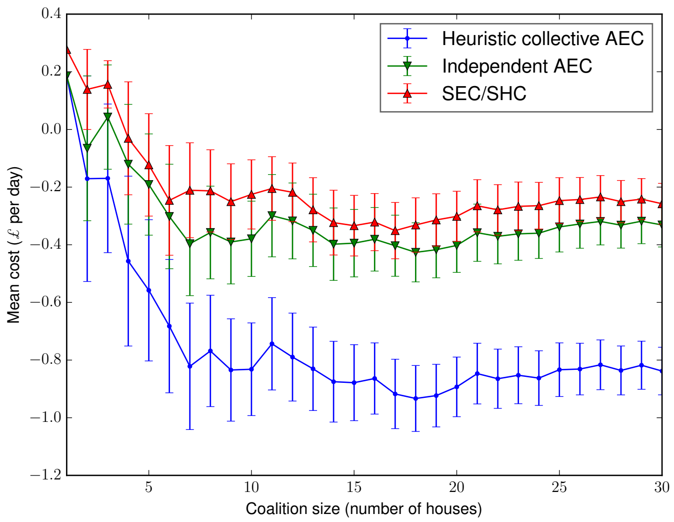

FIGURE 4.26: Cost with fixed mean discomfort at ∼12°Ch.

houses is already optimal in terms of flattening the demand, while the mean monetary gains of the members are higher for smaller coalitions. Although these exact results are specific to the peculiarities of our evaluation case study, they provide an indication of the required size of such a coalition in a typical scenario (Section 4.2.1). Finally, in order to further identify the financial gains of AdaHeat, we evaluate a specific instance of our approach given fixed discomfort-cost preferences. Figure 4.26 illustrates how the mean cost of the houses scales with the size of the coalition given a fixed discomfort preference of $\sim$12°Ch for all houses. According to the houses' occupancy schedule this discomfort could corresponds to a 1-2°C deviation (of the indoor temperature from the set-point) per occupied hour which is a reasonable real-setting preference. It can be seen that collective advanced economic control leads to considerable savings (as a fraction of the daily heating cost) in this realistic scenario—the bars correspond to standard mean error capturing the diverse share of gains allocated to each house according to the allocation index. Interestingly, there is no statistically significant difference (at a 0.05 significance level) between the mean house cost in independent advanced economic control and the cost in simple economic/heating control in this discomfort scenario, for all coalition sizes examined. Nevertheless, as we move towards bigger coalitions, the mean house cost of heuristic collective advanced economic control is progressively more-and-more statistically significantly lower than the one of any of the other two methods in this discomfort scenario. As a final note, our evaluation results suggest that, due to their shifting potential (as discussed in detail in Section 1.1.3), heating loads can effectively support collective advanced economic control even if similar occupancy schedules and/or household preferences are considered, the IER power output does not generally follow the energy requirements of the households, even when the IER power output does not generally follow the energy requirements of the households, and when a wide range of typical UK houses (that translate to a wide range of thermal characteristics) are considered.

4.3 Summary

In this chapter we provided a thorough evaluation of our domestic heating automation system (DHAS) approach, AdaHeat. In particular, we independently evaluated AdaHeat with respect to simple heating and simple economic control, as well as advanced economic control. In both cases: (i) we discussed how we choose the case study of our empirical evaluation and how we collected the necessary data, (ii) we described the specific instantiation of our DHAS for the respective case study system(s), (iii) we discussed our evaluation set-up and the instantiation of the benchmark systems, and (iv) we reported the respective evaluation results.

In more detail, regarding simple heating and simple economic control we provided a thorough respective evaluation, along with a comprehensive comparison of state-of-the-art heating automation systems. To this end we choose the living room of a family house in Cambridge, UK for our case study. We showed that this case study is a challenging testbed on the generality and efficiency of our approach (both in terms of thermal modeling and control) due to its specific thermal characteristics. In this context, our evaluation results suggest that AdaHeat deals with the thermal dynamics of houses in a more effective manner compared to the benchmark systems, due to its reliance on adaptive thermal modeling. In more detail, our evaluation results suggest that adaptive thermal modeling can significantly improve the efficiency of domestic heating automation systems, especially when advanced heating automation systems are considered (i.e., SPOT+, AdaHeat). Furthermore, our results suggest that AdaHeat also deals with the inherent uncertainty of the occupancy schedule in domestic settings in a more effective manner compared to the benchmark systems. In particular, it is suggested that the optimal exploitation of the occupancy probabilistic estimates (in the context of AdaHeat heating planning) leads to a more stable performance, in terms of Pareto efficiency, compared to the benchmark planning approaches. In addition, it is suggested that AdaHeat improves the usability and effectiveness of state-of-the-art DHAS approaches in meeting the user preferences. In particular, the evaluation results suggest that reliance on a simple parameter in balancing heating cost and thermal discomfort (as is the case in AdaHeat) is sufficient for efficient DHAS performance, and ensures the usability of the systems in domestic settings. In this context, we showed that AdaHeat is the only system that works sufficiently based on a single weighting parameter which can be learned on-line. Lastly, the evaluation results also suggest that AdaHeat is able to work in conjunction with a diverse range of heating systems in various operational settings. This is shown due to its ability to effectively and efficiently consider both simple heating control and simple economic control in our challenging case study system.

Now, regarding advanced economic control we provided a thorough evaluation considering 30 houses in Mablethorpe, Lincolnshire, UK for our case study. We showed that this case study enables us to evaluate both collective and independent advanced economic control. In this context, our evaluation results suggest that Adaheat is able to efficiently consider the economic aspects that arise in controlling electricity-based space heating systems with respect to the electricity market, and also exploit the respective coalition potential. In more detail, we showed that collective advanced economic control can significantly improve heating cost-efficiency compared to independent advanced economic control, and even more when compared to independent simple economic/heating control. In addition, we showed that this coalition potential can be exploited without any loss in the usability and generality of our approach.

To sum up, our evaluation results suggest that our requirements as stated in Section 1.1.2 have been met. In particular, AdaHeat is general enough to consider a wide range of space heating systems typically employed in domestic settings and effectively accounts for the dynamic domestic thermal characteristics as well as for the occupancy uncertainty that arises in domestic settings. In addition, AdaHeat relies to the minimum extent on user-input and is able to meet the, potentially time-varying, user preferences through a simple boolean feedback procedure. In this context, AdaHeat demonstrates an adequate performance in terms of Pareto efficiency and distribution of captured solutions compared to state-of-the-art benchmark systems, enabling also a wide range of user preferences to be met. AdaHeat is able to efficiently and effectively consider simple heating control, simple economic control, as well as advanced economic control, exploiting also the coalition potential that arises in the later. To this end, AdaHeat also comes complete with a practical gain allocation mechanism to share the realized collective gains among the coalition members in the case of collective advanced economic control. Finally, AdaHeat has low computational complexity and efficiency that allows it to be applicable in real settings with limited computational resources, minimum instrumentation, and operating time constrains.

Chapter 5

PreST: A Dynamic Programming Predictive Solar Tracking Approach

In this chapter we detail our dynamic-programming-based predictive solar tracking (ST) approach, PreST. As discussed in Section 6.4, PreST considers a low-cost (i.e., does not make use of expensive equipment or sensors) and generic (i.e., applicable to a wide range of commonly employed ST architectures) optimal control ST approach that aims to meet the requirement of performance optimality (Section 1.2.2) respecting at the same time the applicability requirement (as detailed in Section 1.2.2). The backbone of PreST is the estimation of the optimal trajectories a day before, based on weather forecasts that can come from online providers for free. To this end, in Section 5.1 we, first, outline the necessary astronomical background with respect to ST, revealing the key concepts in maximizing incident solar radiation. Then, in Section 5.2 we provide a detailed discussion on the popular azimuth-altitude dual axis tracking (AADAT) and vertical single axis tracking (VSAT) systems, that are the main focus of our ST work. Subsequently, in Section 5.3 we provide our formalization of ST as a dynamic programming problem. In particular, in Section 5.3.1 we define the corresponding Markov decision process (MDP), in Section 5.3.2 a general ST consumption model to appropriately model the process dynamics, and, in Section 5.3.3, we provide a discussion on optimally solving the respective MDP and the challenges that arise. Then, in Section 5.4 we describe our approach for estimating the next day ST policy. In particular, in Section 5.4.1 we describe our policy iteration approximation technique, solar tracking policy iteration (STPI), that we devised to compute a beneficial ST policy, in Section 5.4.2 we describe our myopic method used to enhance the effectiveness of STPI and in Section 5.4.3 we describe our approach in calculating the next-day-optimal fixed photovoltaic system (PVS) orientation, suitable for fixed (yet readjustable) PVSs. Finally, Section 5.5 summarizes this chapter.

5.1 Astronomical Aspects of Solar Tracking

As discussed in Section 1.2, ST can be used to increase the power output of a photovoltaic system (PVS), by orienting the system towards the greatest possible levels of incoming solar irradiance. Now, the total irradiance G_T falling on an arbitrarily oriented surface, consists of the beam G_B, sky-diffuse G_D and ground-reflected G_R components (Luque and Hegedus (2011)) as seen in Equation 5.1 below (and illustrated in Figure 5.1):

$$G_T = G_B + G_D + G_R \tag{5.1}$$

Usually, the cosine effect is used to model the variations of the G_B component, as seen in Equation 5.2 below:

$$G_B = G_B^{max} \cos\theta_s \tag{5.2}$$

where θ_s is the angle between the normal to the surface and the direction to the sun (as seen in Figure 5.2) and G_B^{max} is the incident beam irradiance when the surface is oriented normally to the incoming radiation (i.e., $\theta_s = 0°$). G_B^{max} is the maximum beam irradiance that the PV module can orient to, and depends on weather conditions and solar position.

The G_D component varies according to Equation 5.3 which assumes that every point of the celestial sphere emits light with equal radiance (Liu and Jordan (1961)):

$$G_D = G_D^{max}(1 + \cos\beta)/2 \tag{5.3}$$

where β is the inclination angle of the surface and G_D^{max} is the incident diffuse irradiance for $\beta = 0°$. G_D^{max} is the maximum diffuse irradiance that the PV module can orient to, and depends on weather conditions and solar position.

Finally, the G_R component is modeled by Equation 5.4 which assumes that the ground is horizontal, of infinite extent, and reflects uniformly in all directions (Luque and Hegedus (2011)):

$$G_R = G_R^{max}(1 - \cos\beta) \tag{5.4}$$

where G_R^{max} is the maximum reflected irradiance that the PV module can orient to (for $\beta = 90°$), and depends on weather conditions and solar position.

5.2 Solar Tracking Architectures

As discussed in Section 1.2.1, many ST architectures are commonly employed in practice. However, for the reasons outlined in Section 1.2.1, we focus on both azimuth-altitude dual axis tracking (AADAT) and vertical single axis tracking (VSAT) systems. The AADAT has two degrees of freedom, rotating over a slope (elevation) and an azimuth

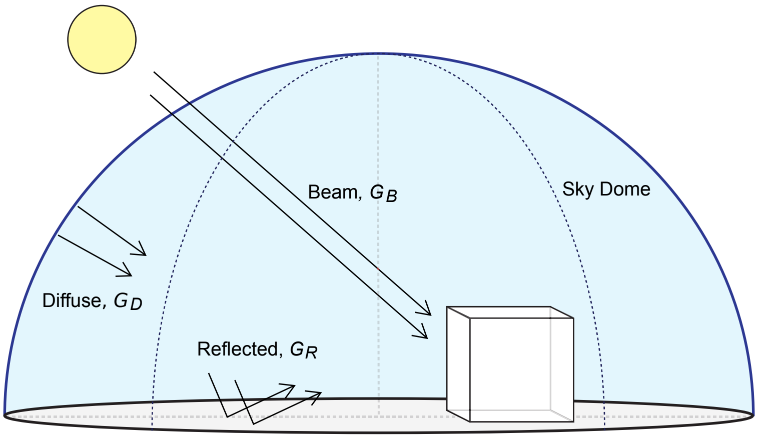

FIGURE 5.1: Solar irradiance components.

axis. An abstract AADAT is illustrated in Figure 5.2. The VSAT rotates only over the azimuthal axis, while its slope angle is kept fixed.

Typically, the movement allowed in tracking systems is constrained within a certain range in both the azimuthal and elevation axis. We henceforth denote the allowed azimuthal and elevation angular range by r_{Az} and r_{Sl} respectively.

The possible slope and azimuth orientations of a dual axis system consist of a discrete number of possible positions within the allowed range at each rotation axis, depending on the tracker step size. Now, a misalignment of $\pm 1°$ causes only a minor drop of $\sim$0.015% in the incident beam irradiance G_B (cf. Equation 5.2). Thus, small misalignments are not a concern for typical commercial systems.

The controller step-size (i.e., the system's minimum angular displacement) θ gives rise to two sets of distinct possible orientation positions for the PVS (one such set per axis of movement). We denote these by K, the set of azimuth orientation positions, and by Λ, the set of possible positions on the elevation axis. In particular, we have $|K| = \lfloor r_{Az}/\theta \rfloor + 1$ and $|\Lambda| = \lfloor r_{Sl}/\theta \rfloor + 1$. The time required for a minimum displacement θ to occur is denoted by δ; its value is assumed constant in our model, in order to maintain a fixed mean angular velocity for every minimum displacement.

Now, the controller requires some time to interact with the PVS, and it takes the system some time to execute the controller commands. For simplicity, the controller in our model is synchronous, meaning that any two consecutive controller-system interactions are separated by a fixed-length time interval Δ. A natural choice is to pick a Δ length that is sufficient to move the PV panel at any orientation starting from an arbitrary position (i.e., $\Delta \geq \delta \cdot max(|K| - 1, |\Lambda| - 1)$), and small enough so as the environmental conditions do not change abruptly during this interval.

Hence, the operation time of a PVS is naturally divided into a number of equal time intervals of length Δ each. Letting Δ_{Day} stand for the day-length (i.e., the time between

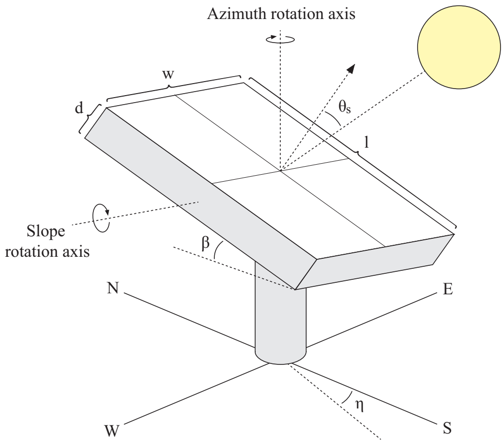

FIGURE 5.2: Abstract AADAT (in VSAT β is fixed)

sunrise and sunset), the required controller-system interactions are contained in a set I, with $|I| = \lceil \Delta_{Day}/\Delta \rceil$. Note that each interaction corresponds to a specific *time-step* or *interaction id* τ; and that these interactions are sufficient to orient the PV panel to any orientation right after the sunrise, and to move it back to the initial position after the sunset.[1]

5.3 A Dynamic Programming Approach

The aim of our work is to calculate optimal ST trajectories for the day ahead, based on available weather forecasts—which can actually come from online providers for free.[2] To this end, we formulate ST as a Markov decision process (MDP). Due to the size of the corresponding MDP we employ dynamic programming for solving it (see Section 2.13) and, in particular, an intuitive *policy iteration* technique (and variants). The proposed optimal control approach within the context of PreST aims to meet the requirement of performance optimality (as stated in Section 1.2.2) with low cost, low complexity and effectiveness that allows our system to be widely applicable in real settings (as per the applicability requirement, detailed in Section 1.2.2). Moreover, our general and parameterizable ST formulation, along with the generic consumption model proposed, aims to enable our ST approach to work in conjunction with the diverse range of ST architectures typically employed in practice (meeting the respective generality requirement, Section 1.2.2).

[1]In particular, the first interaction will be right before (or exactly at) sunrise, and the last one right before (or exactly at) sunset.

[2]We note here, that all our methods are able to also consider on-line updates of the weather forecast reports, if available.

5.3.1 Defining the MDP

The problem is naturally modeled as a fully observable, finite-horizon, discrete-time Markov decision process (MDP) corresponding to an $\langle S, A, P, R\rangle$ tuple as follows:

First, S is a finite set of *states*, where each state $s \in S$ corresponds to a tuple $\langle \kappa_s, \lambda_s, \mathbf{w}_s\rangle$ with $\kappa_s \in [1, |K|]$ denoting the azimuth orientation position, $\lambda_s \in [1, |\Lambda|]$ denoting the slope orientation position, and $\mathbf{w_s}$ standing for the vector of *stochastic weather condition variables* which are required to calculate the MDP reward dynamics (i.e., prevailing wind speed and direction, relative humidity, temperature, and sky conditions). As such, $|S| \geq |K \times \Lambda|$. The value of each state depends on the $\tau \in [1, |I|]$ time-stamp at which the state is visited—that is, $|I|$ represents the horizon for our problem. Note that the time-stamp τ at which s is visited is sufficient to extract all necessary information regarding the *non-stochastic environmental conditions* (i.e., the solar position angles) relevant to s and τ. Therefore, these do not have to be explicitly included in the state representation.

Then, A is a finite set of *actions*, with each action $a \in A$ positioning the PVS to some specific orientation. Thus, a corresponds to tuple $\langle \kappa_a, \lambda_a\rangle$, with $\kappa_a \in [1, |K|]$ and $\lambda_a \in [1, |\Lambda|]$. Hence, we have $|A| = |K \times \Lambda|$.

The *transition model* P defines the $P(s, a, s')$ probability that taking action $a = \langle \kappa_a, \lambda_a\rangle$ in state s will lead to s'. Thus, given a particular action a at a state s, for the successive state s' we will have: $\kappa_{s'} = \kappa_a$, $\lambda_{s'} = \lambda_a$; while the transition probabilities will depend entirely on the $P(\mathbf{w'_s}|\mathbf{w_s})$ ones. Note that in the case that $\mathbf{w'_s}$ is independent of $\mathbf{w_s}$ we will not have a *probabilistic transition model*, but rather a *probabilistic reward model*. Hence, in that case, $\mathbf{w_s}$ can be omitted from the state representation and expected reward values can be extracted directly from the controller interaction id, τ, at which s is visited. The same holds if only non-probabilistic weather forecasts are available.

Finally, R is a *reward model* determining the $R_a(s, s')$ reward received for a transition from state s to s' after taking action a. This reward is the energy produced during the time between two consecutive controller interactions, minus the energy consumed due to the movement of the tracker throughout this interval. Thus:

$$R_a(s, s') = Prod(s, s') - Cons(s, s') \tag{5.5}$$

where $Prod(s, s')$ and $Cons(s, s')$ are functions estimating the energy produced and consumed as a result of the PVS system moving from s to s' (after some action a taken at s).

Calculating $Prod(s, s')$ is straightforward, assuming that the PVS power output is steady throughout a time interval Δ between two consecutive controller interactions:

$$Prod(s, s') = \frac{(Pwr(s) + Pwr(s'))}{2}\Delta \tag{5.6}$$

where $Pwr(s)$ stands for the PVS power output at state s. In our work, the $Pwr(s)$ estimates are provided by RENES, given the PVS orientation (i.e., κ_s,λ_s), the particular time of day (derived based on τ and used to estimate the solar position angles), the (fixed for a given system)) PV characteristics, as well as the stochastic weather conditions in $\mathbf{w_s}$.

5.3.2 Consumption model

A distinct contribution of our work is the construction of a generic and parameterizable solar tracker consumption model. This is the first time that such a model is proposed which is necessary for modeling the ST dynamics in the context of an optimal control approach. We derive an appropriate model by using a white-box modeling approach based on well-known physical principles and our own mathematical derivations. In more detail, with an arbitrary displacement corresponding to an aggregation of minimum angular displacements on each one of the rotation axes, we calculate the consumption of an arbitrary displacement as the (efficiency-weighted) sum of the consumptions corresponding to these minimum angular displacements. Now, in order to maintain a fixed mean angular velocity for every minimum angular displacement θ, every such θ is assumed to follow a simple trapezoid motion profile with three motion phases (of equal time duration): *(1)* an *angular acceleration phase*, *(2)* a *constant angular velocity phase*, and *(3)* an *angular deceleration phase*. As such, the consumption for θ is calculated as the sum of the consumption for all three motion types in sequence. Then, the system consumption $Cons(s, s')$ is:

$$Cons(s,s') = \frac{1}{c_{eff}}\Bigg(\sum_{1}^{|\kappa_s-\kappa_{s'}|} Cons_\theta^{az} + \sum_{1}^{|\lambda_s-\lambda_{s'}|} Cons_\theta^{sl}\Bigg) \tag{5.7}$$

where c_{eff} stands for the efficiency factor of the tracking system. This corresponds to the mean efficiency of the motors, multiplied by the mean efficiency of the gears, and is further reduced to best fit all other secondary losses of the system during a displacement. $Cons_\theta^{az}$ and $Cons_\theta^{sl}$ represent the consumption of every minimum angular displacement θ over the azimuth and slope (elevation) axis respectively. Their values are calculated by the following equations:

$$Cons_\theta^{az} = \sum_{\mu=1}^{3}(\alpha_\mu \mathcal{I}_{A_{(\theta,\mu)}} - T^w_{A_{(\theta,\mu)}})\theta_\mu \tag{5.8}$$

$$Cons_\theta^{sl} = \sum_{\mu=1}^{3}(\alpha_\mu \mathcal{I}_{S} - T^w_{S_{(\theta,\mu)}})\theta_\mu \tag{5.9}$$

where θ_μ and α_μ stand for the angular displacement and acceleration for each one of the motion phases, and can be computed as $\theta_1 = \theta_3 = \theta_2/2 = \theta/4$, and $\alpha_1 = -\alpha_3 = 9\theta/2\delta^2$

and $\alpha_2 = 0$. Now, $\mathcal{I}_{A_{(\theta,\mu)}}$, $\mathcal{I}_S$, $T^w_{A_{(\theta,\mu)}}$ and $T^w_{S_{(\theta,\mu)}}$ stand for the *moment of inertia* and wind *torque*, for the azimuth and slope axis respectively.

For the slope rotation, the moment of inertia is independent of the azimuthal orientation, and, assuming the panel is a cuboid, can be given by $\mathcal{I}_S = \frac{m}{12}(l^2 + d^2)$ (Myers (1962)), where m stands for the mass of the panel, and l and d for the length and thickness of the panel as seen in Figure 5.2. Note, however, that the azimuthal motion occurs simultaneously with the slope one. Hence, T^w_S, T^w_A and $\mathcal{I}_A$ are not constant during the motion. Nevertheless, due to the very small displacement corresponding to each motion phase, these quantities are assumed constant and equal to their value at the beginning of each motion phase.

For the azimuthal rotation, the moment of inertia depends on the slope orientation. Assuming that the panel is a cuboid, the moment of inertia for the azimuthal rotation given a particular slope angle β can be computed as follows (a proof can be found in Appendix B):[3]

$$\mathcal{I}_A = \frac{m}{12}\left(l^2\cos^2(\beta) + d^2\sin^2(\beta) + w^2\right) \tag{5.10}$$

where w stands for the *width* of the panel.

We also modeled the PVS aerodynamics, estimating the torque on the rotation axes due to the incident wind as $T^w_X = \frac{1}{2}\rho w l^2 V^2 c_X$, where $X \in \{A, S\}$, ρ denotes the *air density*, V the prevailing *wind speed*, and c_A and c_S denote the non-dimensionalized slope and azimuth *moment coefficients*, respectively. These coefficients depend on the orientation of the system. The air density was estimated based on the *local pressure*, the *relative humidity*, and the *temperature*, based on standard meteorological equations (Picard et al. (2008)). The moment coefficients are calculated based on *wind direction* and *system orientation*, as in Roos (2012) which provides correlations on the well known Peterka dataset (Peterka et al. (1986)).

5.3.3 Optimal Solar Tracking

With the above MDP at hand, the *optimal ST policy* can be derived by solving the corresponding Bellman optimality equation (Equation 2.12, Sutton and Barto (1998)). However, due to the size of the state and action spaces (typically $|I| \cdot |S| \cdot |A| > 4$Bn, without even considering $\mathbf{w_s}$),[4] the optimal tracking policy for the day-ahead cannot be computed exactly in realistic settings applications (Sutton and Barto (1998); Puterman (2014)). Rather, it can only be approximated. To this end, we have devised several approximation methods, which we now proceed to describe.

[3] Also presented in Panagopoulos and Chalkiadakis (2015).

[4] In more detail, for a day with 12 daylight hours (i.e., $\Delta_{Day} = 12$ *hours*); a typical system (like the one considered in our evaluation) with $r_{Az} = 270°$, $r_{Sl} = 63°$ and $\theta = 1.8°$ (at each axis); and a control interval of 5 *min* (i.e., $\Delta = 5$ *min*), there will be: $|I| = \lceil \Delta_{Day}/\Delta \rceil = 144$, $|A| = |K| \cdot |\mathit{\Lambda}| = \lfloor r_{Az}/\theta \rfloor + 1 \cdot \lfloor r_{Sl}/\theta \rfloor + 1 = 5{,}436$, $|S| \geq |K| \cdot |\Lambda| \Rightarrow |S| \geq 5{,}436$, and, hence, $|I| \cdot |S| \cdot |A| \geq 4{,}255{,}213{,}824$.

5.4 Approximation Methods

In this work we propose three methods to approximate the optimal ST policy: (i) Solar Tracking Policy Iteration (STPI), (ii) Myopic and (iii) next-day-optimal fixed PVS orientation. The STPI method considers a policy iteration schema that aims to approximate the optimal policy by improving on the ST policy as derived by the Myopic method. In particular, the Myopic method aims to approximate the optimal policy by maximizing power generation alone (disregarding of the tracking consumption cost). In this context, STPI improves this policy by taking into account the tracking consumption cost. Finally, the *next-day-optimal fixed PVS orientation* defines the next day optimal fixed-orientation (and has minimum computational requirements). We now describe the approximation techniques we developed in order to compute effective ST policies.

5.4.1 Solar Tracking Policy Iteration method (STPI)

We devised a policy iteration (PI) approximation technique to compute a beneficial ST policy. The technique interweaves two distinct PI procedures, which are used in an alternating fashion. The first PI procedure, *SlopePI*, considers an arbitrary input policy for the above MDP, e.g., a myopic one. It then attempts to improve that policy, in the usual PI fashion, nevertheless assuming a *fixed azimuthal policy*, π_κ. Given this fixed π_κ policy, it computes the respective optimal slope-positioning policy, π_λ. The output policy is then fed in a second PI algorithm, which estimates an optimal (given π_λ) azimuth-positioning policy, π_κ. The process repeats until convergence, or until some computational or time limit is reached. By combining the derived policies computed for each axis, we can derive a ST policy. The same PI algorithm can be readily employed for single axis tracking, with the action selection process for the static axis—the slope one, in the case of vertical single axis tracking (VSAT)—considering only a set of fixed possible orientations for the whole motion (so as to estimate the best possible fixed slope angle for VSAT tracking during the next day). The overall PI technique is shown in Algorithm 4, while Algorithm 5 describes the PI process to derive a slope policy (the PI for deriving an azimuthal policy is entirely similar). Note that STPI effectively alternates between solving MDPs with state-action spaces which are orders of magnitude smaller than that required by the original problem formulation. Moreover, STPI is expected to converge to a fixed point. In particular, each iteration of STPI improves on the input policy (as per the typical policy iteration monotonicity analyses, see Puterman (2014)), and, hence STPI experiences monotonicity. In addition, by definition there is an upper value bound considering the policies (i.e. the value of the optimal policy) (as per the MDP formalization, see Puterman (2014)). As such, in the case of finite state-action spaces, STPI is expected to converge to a fixed point in a finite number of steps in accordance to the monotone convergence theorem (see Puterman (2014)). We note here again that this is the first time that an approach that alternatively optimizes over MDP action sub-spaces is proposed for optimal policy approximation.

Algorithm 4 "Alternating" Policy Iteration for ST

procedure STPI(π)
  Initialize π_λ and π_κ based on π
  while π_λ and π_κ are not stable **do**
    $\pi_\lambda \leftarrow$ SlopePI(π_λ, π_κ)
    $\pi_\kappa \leftarrow$ AzimuthPI(π_κ, π_λ)
  Derive π' by combining π_κ and π_λ
  return π'

Algorithm 5 Slope Policy Iteration

procedure SlopePI(π_λ, π_κ)
  while π_λ is not stable **do**
    for all $\tau \in I$ in descending order **do**
      for all $s \in S$ that can emerge based on π_κ at τ **do**
        $a \leftarrow \langle \kappa_a = \pi_\kappa(s,\tau), \lambda_a = \pi_\lambda(s,\tau) \rangle$
        $V_\tau(s) \leftarrow \sum_{s'} P(s,a,s')\,(R_a(s,s') + V_{\tau+1}(s'))$
    for all $\tau \in I$ (in any order) **do**
      for all $s \in S$ that can emerge based on π_κ at τ **do**
        $\pi_\lambda(s,\tau) \leftarrow \operatorname{argmax}_\lambda \sum_{s'} P(s,a,s')(R_a(s,s')+$
          $V_{\tau+1}(s'))$, where $a = \langle \kappa_a = \pi_\kappa(s,\tau), \lambda_a = \lambda \rangle$
  return π_λ

That said, the choice of the initial policy used is crucial for the efficiency of any policy iteration algorithm (as discussed in Section 2.13). The initial policy we use is a *myopic* one, which maximizes power output alone (disregarding any associated repositioning costs). Now, the tracking consumption of a PVS is a very small fraction of its production (typically less than 1% (Mousazadeh et al. (2009))). As such, the *myopic policy* is essentially a *near-optimal* one, achieving over 99% of the optimal performance. Moreover, any gains achieved by a method that lowers consumption is typically accompanied by costs due to lower production. Hence, the near-optimality guarantees of the myopic policy are even stronger; and they are *inherited by STPI*, as any derived improved policy cannot be worse than the initial one. We now describe how to derive the myopic policy.

5.4.2 Myopic method

We define the *myopic policy* in this domain as a method that maximizes power generation alone. As such, we can modify the MDP reward function to account for the PVS power output alone: $R_a(s,s') = Prod(s,s')$. Given the fact that all possible PVS orientations are accessible from any state, it is clear that the optimal policy in this modified MDP is equivalent to the one that chooses the action that gives the maximum expected reward for the next time interval.

Now, the power output of a PVS increases proportionally to the incident irradiance. In order to maximize the incident G_T we need to maximize the sum $G_B + G_D + G_R$ as seen in Equation 5.1. Moreover, as seen in Equations. 5.3 and 5.4, G_D and G_R components vary from their maximum based only on the slope angle of the PV module. On the other hand, Equation 5.2 illustrates that the G_B component varies from its maximum based on the incident angle, which for any given sun position depends on the slope and azimuth angle of the PV module. In particular, G_B reaches its maximum as the incident angle reaches zero (i.e., when the azimuth and slope angle of the PV module are the same as the azimuth and slope angle of the sun). As such, fixing the azimuth angle to follow the sun azimuth, ensures that we are always able to track the maximum G_T (for any weather conditions). The only thing that we need to do is to optimize the PVS slope angle at every time step, so that we balance the G_T components and get the maximum expected G_T. For (vertical) single axis tracking, the problem is further simplified into following the sun over the azimuth (and just defining the best next-day fixed slope orientation).

5.4.3 Next-Day-Optimal Fixed-Orientation

In the context of this work, we also propose and calculate the *next-day-optimal fixed PVS orientation*, by simply evaluating the whole space of possible orientations, given the next-day weather prediction. The derived orientation is suitable for any fixed-orientation (yet re-adjustable) PVS operating within the geographical region of a given weather station. Moreover, next-day-optimal fixed positioning can also be used by trackers in the case of scheduled power cuts. We note that this method can also be extended for weekly-optimal or some hours-optimal positioning, as needed.

5.5 Summary

In this chapter we described our dynamic-programming-based predictive solar tracking (ST) approach, PreST. In particular, we first provided a general discussion over key astronomical concepts with respect to solar tracking, ST, revealing the key concepts in maximizing incident solar radiation. In particular, maximizing incident solar radiation is a key element within the performance optimality requirement of effective and efficient ST (as discussed in Section 1.2.2) which is lacking in current open-loop ST approaches (as discussed in Section 1.2.3). Furthermore, we provided a general discussion of ST architectures with a focus on the common implementations of vertical single axis tracking (VSAT) and azimuth-altitude dual axis tracking (AADAT) (that consider also our evaluation systems—as discussed in Section 6.4).

Subsequently, we provided an MDP formulation of ST, which allows it to be tackled as a dynamic programming problem, along with a general ST consumption model to appropriately consider the dynamics of ST. Our general and parameterizable ST formulation, along with the generic consumption model proposed, aims to enable our ST approach to work in conjunction with the diverse range of ST architectures typically employed in practice (meeting the generality requirement as discussed in Section 1.2.2). Subsequently, we provided a discussion on optimally solving the corresponding MDP pointing out the challenges that arise due to the great dimension of the problem. To this end, we provided our approximation solutions. Importantly, we described our policy iteration approximation algorithm; STPI, which is suitable for large MDPs, like the one considered in this work. We also provided a myopic method that is used to enchance the effectiveness of STPI, and a next-day-optimal method that is suitable for any fixed-orientation (yet re-adjustable) PVS operating within the geographical region of a given weather station, enabling efficient manual tracking. All our methods come with optimality or near-optimality guarantees, and our next-day policy comes complete with an expected PVS power output estimation, which is crucial for the smooth integration of PVSs into the electrical grid (as discussed in Section 6.4).

The proposed optimal control approach aims to meet the requirement of performance optimality (Section 1.2.2) respecting at the same time the applicability requirement (as detailed in Section 1.2.2). In the following chapter, we provide an evaluation of our ST approach, showing that it outperforms all commonly employed ST benchmark techniques (i.e., chronological, sensor-based and/or fixed-orientation), which can lead to significant monetary gains.

Chapter 6

Evaluating PreST

In this work we provide a detailed evaluation of our solar tracking (ST) approach, PreST based on real data. In the following sections we first describe our evaluation setup, considering also the ST techniques used as benchmark (Section 6.2). Then, in Section 6.1 we describe the case study of our evaluation and how we collected the necessary data. Then, in Section 6.3 we report the evaluation results. Finally, Section 6.4 summarizes this chapter.

6.1 Case Study and Data Collection

For the case study of our evaluation we consider a photovoltaic system (PVS) located at Chania, Crete, Greece (as seen in Figure 4.14). We choose Crete for our evaluation due to the great degree of PVS penetration on this sunny Greek island.[1] Moreover, this choice ensures that RENES provides accurate PVS power output predictions (see Panagopoulos et al. (2012)). Now, our case study PVS is a typical $72m^2$ system (i.e., $w = 6.0m$, $l = 12.0m$, $d = 0.20m$), weighting $\sim 2500kg$, with 270° of azimuthal motion range, and 63° of elevation motion range. The system has a *step-size* of $\theta = 1.8°$ at each axis, which can lead to a maximum misalignment of $\arccos(\cos^2(\theta))/2 \simeq 1.27°$, corresponding to a G_B drop of $\sim$0.025%. In this context, the time δ required for a minimum displacement θ to occur, was set to $1s$, and the interval between two consecutive controller interactions was set to $\Delta = 5min$. Finally, as the efficiency of the motors and gears depends on the speed and load at all times (Burt et al. (2008)), we used a mean efficiency of 30% for both[2]. We note that this choice leads to a tracking consumption that is close to the reported practical value for such systems (Mousazadeh et al. (2009)).

[1] $\sim 60MW$ of installed PV power, corresponding to $\sim$7% of Crete's energy production and $\sim$70% of the total Greek islands' installed PV power (Public Power Corporation of Greece).

[2] Based on data provided at `www.acosolar.com`; and `users.ece.utexas.edu/~valvano/Datasheets`.

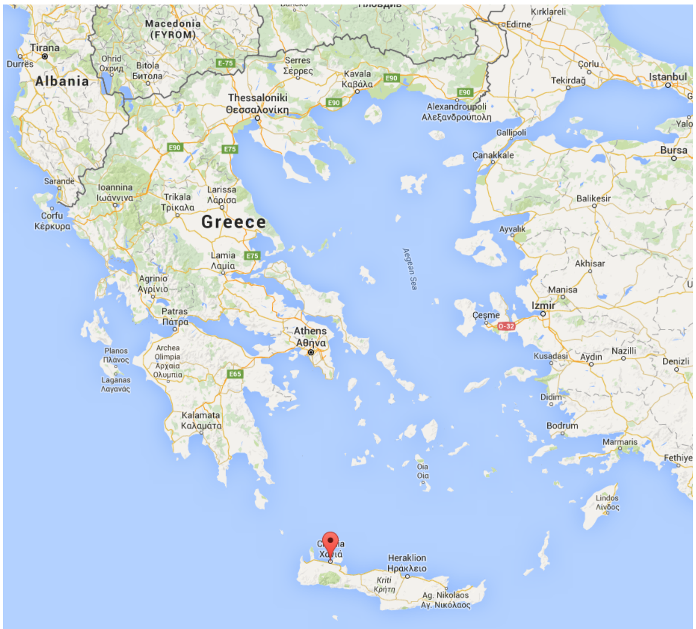

FIGURE 6.1: Case study location (powered by Google maps).

For the purposes of our research, archival weather data was accumulated from weather underground[3] considering distinct limit-case days within the $2008 - 2012$ period, for our location of interest. In particular, the acquired meteorological variables are: *relative humidity*, *temperature*, *wind speed*, *wind direction*, and *qualitative cloud coverage observations* (appropriately transformed to quantitative values, as in Panagopoulos et al. (2012)). As there is a 30-minute gap between consecutive archival weather data, we use linear interpolation to meet the model's time interval requirements. Furthermore, as we are interested in the prevailing conditions within a $\Delta = 5min$ interval, all variables are assumed constant and equal to their mean value within that interval.

6.2 Experimental Setup

In order to evaluate PreST we consider a photovoltaic system (PVS) located at Chania, Crete (as discussed in Section 6.1) and estimate its output energy gain from employing each one of our methods. We note here that considering locations with lower sunshine and greater wind speeds for our evaluation, would only favor our methods. This is supported by the fact that PVSs in such locations would exhibit a higher consumption over production ratio and, hence, a greater output energy gain from using ST methods that also consider the tracking consumption (further illustrated by our evaluation with "fictional" weather data in Section 6.3).

Now, our evaluations are performed for 8 different daily weather patterns; 4 of them corresponding to actual, real days, and 4 of them fictional. Specifically, we utilized

[3] `www.wunderground.com`

archival weather data for the 20/03/2011 equinox, the 22/09/2012 equinox, the solstice of 21/06/2012, and the solstice of 21/12/2008 (collected as detailed in Section 6.1). These days are noted from now on as day 1, 2, 3 and 4 respectively. The general weather patterns of the days considered is as follows:

- **Day 1:** The general weather pattern for *Day 1* consists of several transitions in the cloud coverage levels (moving from mostly sunny to scattered clouds, then to mostly cloudy, and back to mostly sunny).
- **Day 2:** The general weather pattern for *Day 2* consists of a simple transition in the cloud coverage levels, from mostly cloudy to clear sky.
- **Day 3:** There was a clear sky throughout *Day 3*.
- **Day 4:** There was full cloud coverage throughout *Day 4*.

As such, though only four, these days exhibit weather patterns that are quite distinct from each other. Moreover, these days also consider distinct cases in terms of the day length. Hence, they enable us to evaluate PreST under a wide range of meteorological conditions at our location of interest. Nevertheless, throughout these four days, the prevailing wind speed was quite low. Given the fact that power consumption grows with *high* prevailing wind speed, we decided to evaluate PreST with fictional wind data. We thus created four additional, "fictional" days, with exactly the same weather conditions as their real counterparts—apart from the wind, whose speed was set to *60 km/h*. That value corresponds to a typical maximum wind speed that a PVS can withstand without a need to orient itself to a safe position (Peterka and Derickson (1992)).

In this context, we compared our methods against three additional baseline methods we implemented for this purpose:

- Chronological vertical single axis tracking (VSAT).
- Chronological azimuth-altitude dual axis tracking (AADAT).
- A yearly optimal fixed-orientation system.

In more detail, the *chronological* AADAT calculates the sun positions as prescribed in the work of Reda and Andreas (2004), and then orients the PVS so as to point towards the sun, irrespective of weather conditions. For the *chronological* VSAT, we used the same procedure to calculate the PVS azimuthal positions; while the slope angle was fixed to its yearly optimal value for VSAT tracking, given the location's latitude, as prescribed in (Li et al. 2011). Finally, we used the equations provided in Chang (2009) to calculate the yearly optimal slope position for *fixed-orientation* south-facing panels at our location of interest.

Finally, we note here that, in the absence of probabilistic weather forecasting reports (and respective online providers), we used deterministic archival weather data for both the weather predictions and ground-truth. As such, in our experiments, the accuracy of the weather forecasts does not affect the efficiency of our methods. Moreover, in the absence of weather prediction uncertainty, the evaluation results of the proposed low-cost *Myopic* method are equivalent to a tracking system where an expensive sensor arrangement along with a closed-loop controller is used to orient the solar panel towards the maximum incident solar irradiance. As such, by comparing *STPI* against *Myopic* we effectively also compare our sensor-less, low-cost STPI against sensor-based ST.

6.3 Evaluation Results

The evaluation results of our experiments are collectively reported in Table 6.1, while the net gains alone are also plotted in Figure 6.2 to enhance comparison. All energy values are in kWh, and correspond to PVS net energy gain. Tracking consumption is also reported inside parenthesis (when applicable).

In all experiments, our methods clearly outperform the baseline ones. It is also worth noting that, in general, as the system's degrees of freedom are increased, so do the positive system efficiency effects from using our methods (i.e., compared to fixed-orientation systems, the net energy gain increases when using *Myopic* or *STPI* with one rotation axis; and it increases even more when using these methods with two axes of rotation). By contrast, the benefits from using chronological ST often decrease when moving from fixed-orientation to one and, further on, to two rotation axes, as the additional system abilities are not fully exploited.

Regarding the methods' individual performance, not surprisingly, *next-day optimal fixed-orientation* significantly outperforms the yearly optimal one, as the former is specialized for the particular day. In addition, *Myopic* gives a significant advantage over chronological tracking, in both VSAT and AADAT, as it also considers the weather conditions. At the same time, *STPI* does consistently better than *Myopic*, even though not by a wide margin. This low improvement margin is not surprising: in an appropriately designed, sizable PVS, like the one considered for our case study, the tracking consumption is *much* lower than the energy produced (less than 1% (Mousazadeh et al. (2009))). Thus, the net energy gains achieved by methods that take consumption into account, are not expected to differ dramatically from those achieved by methods that maximize power generation notwithstanding consumption needs. This fact is confirmed from our evaluation results: an improvement from using *STPI* instead of *Myopic* is present in all days and tracking systems, but is more substantial for high prevailing wind speeds, and especially for dual-axis tracking. However, over time (i.e., within a long operating time window and/or for clusters composed of many PVSs put together), even small improvements like the ones observed are significant. Even for an average-sized PV park of 2MW nominal power, one would be able to, annually, gain over €1500 more by using *STPI*, compared to *Myopic* (and over €10000 compared to chronological AADAT).[4]

Now, our evaluation results suggest that our case study PVS experiences a higher improvement from using *STPI* compared to *Myopic* in cloudy and windy days. This is not surprising since in such settings our PVS exhibits a higher consumption over production ratio (compared to sunny days with generally low prevailing wind speeds) and, hence, a higher potential improvement from using *STPI* which attempts to approximate the optimal policy. That said, smaller PVSs, or not very efficiently designed ones, would generally exhibit a higher consumption over production ratio, compared to our case study system for all weather conditions. Hence, such systems would generally exhibit a higher improvement from using *STPI* compared to *Myopic*. To illustrate the above, here we contact an additional experiment where an identical PVS is considered with, however, only half of the original generation capacity. The respective evaluation results are collectively reported in Table 6.2, while the net gains alone are also plotted in Figure 6.3 for concreteness. As can be seen, the relative net energy gain improvement of using *STPI* compared to *Myopic* is higher compared to the original case study PVS for all days considered (i.e., both real and fictional) and for both tracking architectures (i.e., both single axis and dual axis tracking). Of course, the net gain is expected to further improve if the actual optimal policy is computed. However, *Myopic* is already near-optimal, as argued above, and, when compared to *Myopic*, *STPI* is shown to already be achieving higher net energy gains and substantially lower consumption (of up to ∼90% reduction even in our original case study PVS).

As a final note, *STPI* is expected to yield increased benefits when one considers a more detailed consumption model. In particular, the tracking cost is not limited to the motor consumption; there is also the maintenance cost, which should, ideally, also be taken into account. Moreover, real-world buy and sell energy prices will most probably have different values. These requirements can be readily incorporated in our model, by simply modifying Equation 5.5. Specifically, for a grid-connected PVS, Equation 5.5 can be replaced by:

$$R_a(s, s') = Prod(s, s')P^{Sell} - Cons(s, s')P^{Buy} - t^{op}c_m \tag{6.1}$$

Here, c_m denotes maintenance cost given operating time t^{op}, and P^{Sell} and P^{Buy} denote the sell and buy energy prices. The maintenance cost can be estimated considering the original price of the tracking system along with its life expectancy in maximum operating time, as well as any additional maintenance cost per operating time such as the lubrication cost. Nevertheless, such detailed modeling is out of the scope of this work and a respective evaluation would require a real world trial (discussed in Chapter 7).

[4]These estimations are obtained by appropriately extrapolating the mean gains reported in Table 6.1 in accordance to the typical mean electricity export tariffs in Crete in 2014 (`https://www.dei.gr/en`).

Dataset		Fixed-Orientation		Single Axis ST (VSAT)			Dual Axis ST (AADAT)		
	Day	Year-Opt	Next-Day-Opt	Chronological	Myopic	STPI	Chronological	Myopic	STPI
Real	1	31.520 (-)	**32.448** (-)	32.533 (.027)	32.791 (.019)	**32.794** (.015)	32.021 (.061)	33.033 (.078)	**33.070** (.037)
	2	49.736 (-)	**50.275** (-)	52.036 (.029)	52.042 (.028)	**52.046** (.023)	51.624 (.063)	52.326 (.087)	**52.360** (.049)
	3	67.301 (-)	**68.921** (-)	71.037 (.039)	72.977 (.057)	**72.985** (.048)	73.434 (.091)	74.003 (.106)	**74.027** (.080)
	4	11.736 (-)	**11.748** (-)	11.623 (.019)	11.738 (.031)	**11.754** (.010)	11.465 (.037)	11.788 (.059)	**11.822** (.021)
Fictional	1	31.520 (-)	**32.448** (-)	32.530 (.030)	32.784 (.026)	**32.790** (.019)	31.899 (.183)	32.730 (.381)	**32.972** (.121)
	2	49.736 (-)	**50.275** (-)	52.034 (.031)	52.040 (.030)	**52.045** (.023)	51.515 (.172)	52.034 (.379)	**52.247** (.156)
	3	67.301 (-)	**68.921** (-)	71.018 (.059)	72.961 (.074)	**72.977** (.055)	73.264 (.261)	73.706 (.404)	**73.862** (.237)
	4	11.736 (-)	**11.748** (-)	11.615 (.026)	11.729 (.041)	**11.751** (.010)	11.411 (.090)	11.567 (.280)	**11.747** (.032)

TABLE 6.1: Evaluation results (all values are in kWh, and correspond to PVS net energy gain; tracking consumption in parenthesis).

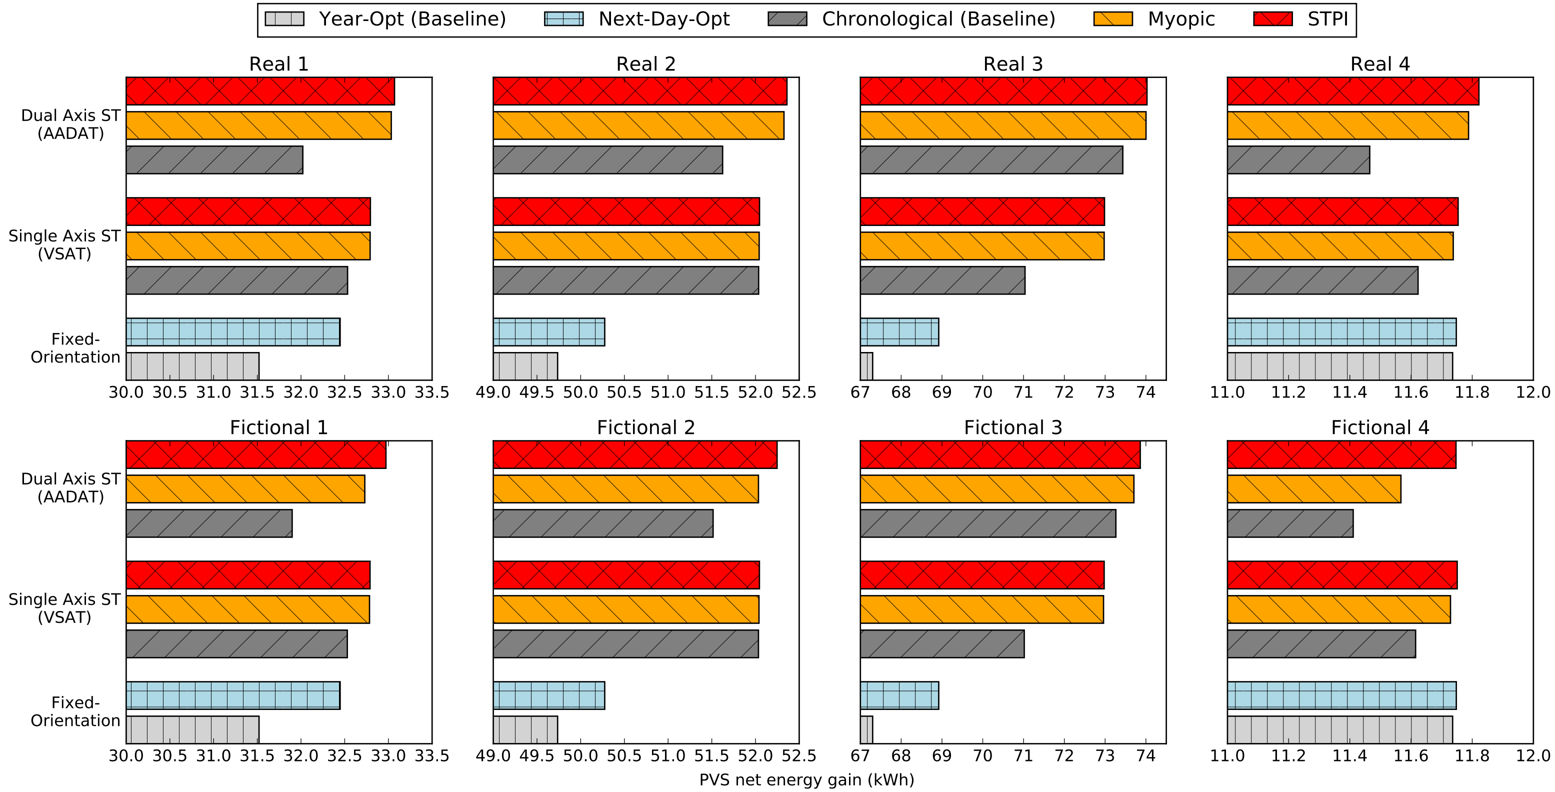

FIGURE 6.2: Evaluation results (bar chart).

Dataset		Fixed-Orientation		Single Axis ST (VSAT)			Dual Axis ST (AADAT)		
	Day	Year-Opt	Next-Day-Opt	Chronological	Myopic	STPI	Chronological	Myopic	STPI
Real	1	15.760 (-)	**16.224** (-)	16.253 (.027)	16.386 (.019)	**16.390** (.013)	15.980 (.061)	16.478 (.078)	**16.517** (.034)
	2	24.868 (-)	**25.137** (-)	26.003 (.029)	26.007 (.028)	**26.012** (.021)	25.781 (.063)	26.120 (.087)	**26.156** (.048)
	3	33.650 (-)	**34.460** (-)	35.499 (.039)	36.460 (.057)	**36.469** (.045)	36.672 (.091)	36.948 (.106)	**36.974** (.078)
	4	5.868 (-)	**5.874** (-)	5.802 (.019)	5.854 (.031)	**5.874** (.003)	5.714 (.037)	5.865 (.059)	**5.901** (.017)
Fictional	1	15.760 (-)	**16.224** (-)	16.250 (.030)	16.379 (.026)	**16.386** (.016)	15.858 (.183)	16.175 (.381)	**16.427** (.101)
	2	24.868 (-)	**25.137** (-)	26.002 (.031)	26.005 (.030)	**26.012** (.022)	25.672 (.172)	25.827 (.379)	**26.059** (.105)
	3	33.650 (-)	**34.460** (-)	35.480 (.059)	36.443 (.074)	**36.460** (.056)	36.501 (.261)	36.651 (.404)	**36.821** (.215)
	4	5.868 (-)	**5.874** (-)	5.794 (.026)	5.844 (.041)	**5.872** (.007)	5.660 (.090)	5.644 (.280)	**5.865** (.009)

TABLE 6.2: Evaluation results, half generation capacity (all values are in kWh, and correspond to PVS net energy gain; tracking consumption in parenthesis).

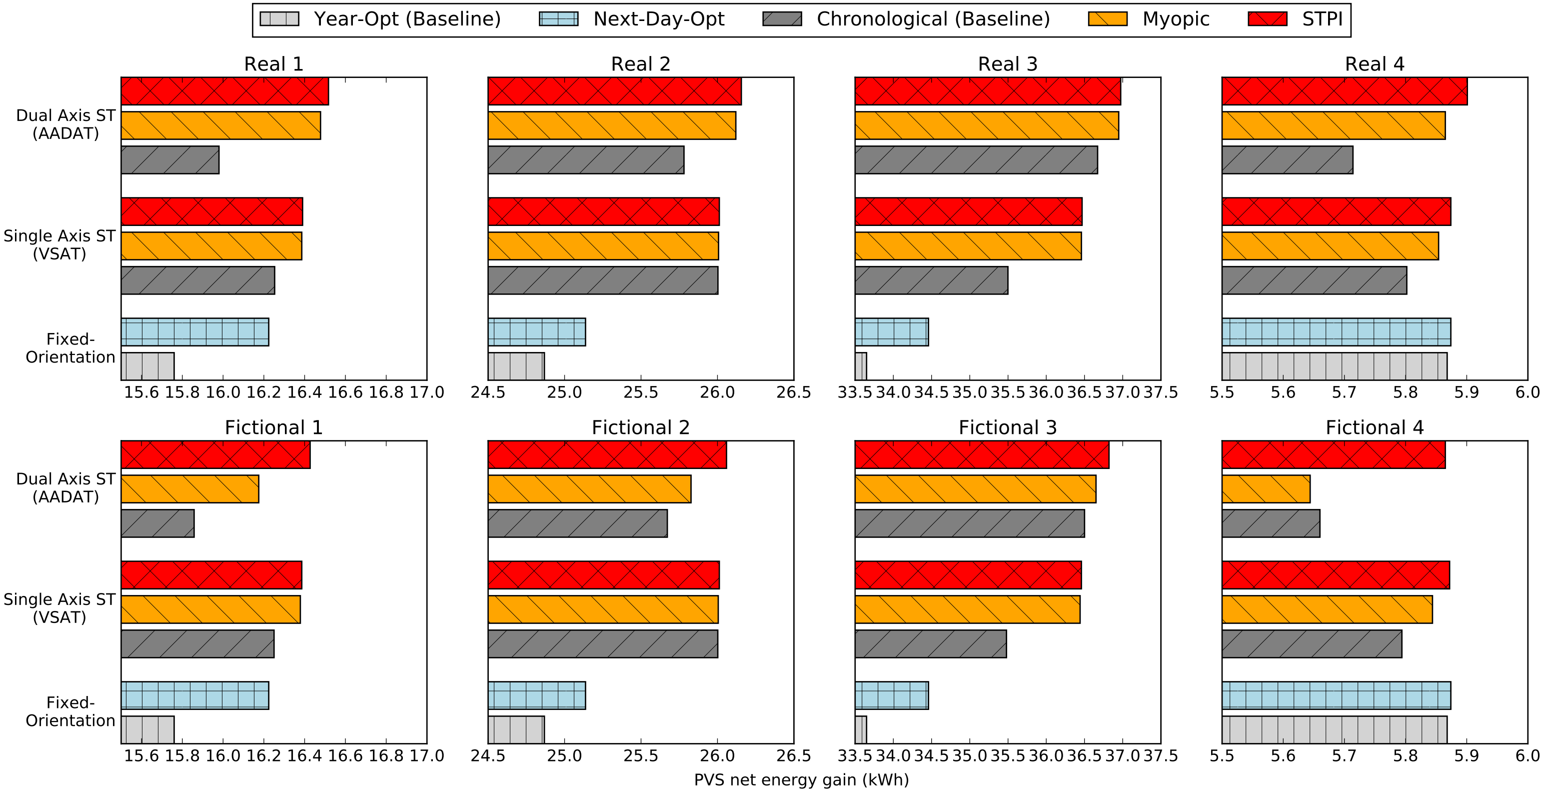

Figure 6.3: Evaluation results, half generation capacity (bar chart).

6.4 Summary

In this chapter we provided a real-data-based evaluation of our ST approach, PreST, and demonstrated its efficiency against commonly employed conventional ST techniques (i.e., chronological, sensor-based, and/or fixed-orientation) in various limit-case weather patterns. In particular, we first described our evaluation set-up justifying the choice of our case study location and the various weather patterns. Further on, we describe in detail the case study of our evaluation and how we collected the necessary data. Finally we reported our evaluation results and showed that *chronological* ST does not fully exploit the additional system abilities, while our methods clearly outperform the baseline ones. In particular, we showed that: (i) the *next-day optimal fixed-orientation* method significantly outperforms the yearly optimal one; (ii) the *Myopic* method gives a significant advantage over *chronological* ST; and (iii) STPI does consistently better than the Myopic/Sensor-based method—even though, not by a wide margin. Finally, we justified why the small improvements of STPI compared to Myopic are not surprising, and showed that despite their seemingly insignificance they can lead to considerable monetary gains for an average-sized PV park, while smaller PVSs (or not very efficiently designed ones), compared to the case-study system, would exhibit higher expected improvement from using *STPI* compared to Myopic. As such, we showed how our novel predictive solar tracking approach; PreST, can meet the optimality requirements of efficient and effective solar tracking (as stated in Section 1.2.2) without relying on expensive equipments or data (and, hence, meeting the respective applicability requirement, Section 1.2.2). Moreover, PreST performed adequately in conjunction with the popular azimuth-altitude dual axis trackers, AADAT, and vertical single axis trackers, VSAT, and, due to its parameterizable and generic nature, can be used in conjunction with many other ST systems typically employed in practice (as discussed in Section), meeting the generality requirement 1.2.2. Hence, all requirements stated in 1.2.2 have been met.

Chapter 7

Conclusions and Future Work

In this thesis we preoccupied ourselves with efficient control of domestic space heating systems and intermittent energy resources (IERs). We proposed specialized systems that increase the operational efficiency of domestic space heating systems and intermittent energy resources with minimum cost. Regarding the former, we proposed a new general domestic heating automation system (DHAS), AdaHeat, while, regarding the latter, we proposed a novel, low-cost and generic, dynamic programming predictive solar tracking approach, PreST.

In more detail, AdaHeat effectively accounts for: (i) simple heating control, (ii) simple economic control, as well as, (iii) advanced economic control that exploits the potential for operating as a coalition. As such, AdaHeat is able to also effectively and efficiently consider electricity-based heating systems even in the presence of import and/or export tariff variability, as well as in the presence of domestic IERs. We note that this is the first DHAS to incorporate advanced economic control and, the first to exploit the respective coalition potential. AdaHeat employs model predictive control (MPC), utilizing adaptive gray-box thermal modeling and a new general heating planning algorithm that fully exploits the probabilistic occupancy estimates employing dynamic programming. As such, AdaHeat is able to effectively account for the highly dynamic thermal characteristics of houses and is also general enough to consider a diverse range of heating systems typically employed in domestic settings. In this context, and due to its component-based structure, AdaHeat considers as a general framework where specific models can be inserted to give particular characteristics. Importantly, AdaHeat adapts to the user preferences in balancing cost and discomfort as it relies on only one parametrization factor that can be learned on-line. In the context of collective advanced economic control, AdaHeat also incorporates: (i) a practical allocation mechanism to share the collective gains of the coalition and (ii) a heuristic planning approach that has a complexity that scales in a linear and parallelizable manner with the coalition size, enhancing, as such, its applicability.

Table 7.1: AdaHeat: Requirements evaluation

Requirement	Met	Comments
1. Minimal user-input	✓	As required, our system relies to the minimum extent on user-input and, in particular, on a simple boolean feedback.
2. Reliable thermal modeling	✓	AdaHeat adaptive thermal modeling effectively handled the dynamic thermal characteristics of our case study systems. Moreover, adaptive thermal modeling has been reported to be resilient and effective in domestic settings in several studies (see Section 2.2)
3. Dealing with occupancy uncertainty	✓	AdaHeat efficiently handles the probabilistic occupancy estimates coming from the appropriate predictive systems.
4. Pareto efficiency	✓	Our system demonstrated an adequate performance in terms of Pareto efficiency and distribution of solutions compared to state-of-the-art benchmark systems.
5. Matching the user preferences	✓	Our system is able to match the user preferences in balancing thermal discomfort and heating cost (for more details see the subsequent requirements below).
5(a) Flexibility	✓	AdaHeat has been shown to capture a sufficiently wide, and evenly distributed, range of balancing points between heating cost and thermal discomfort that allows a variety of user preferences to be captured.
5(b) Usability	✓	AdaHeat matches the user preferences in balancing heating cost and thermal discomfort via an adaptive procedure that requires a simple boolean feedback by the user.
5(c) Adaptability	✓	The on-line procedure of AdaHeat in matching the user preferences, in trading heating cost and thermal discomfort, enables our system to directly adapt to potentially time-varying user preferences.
6. Generality	✓	Due to the general approach followed, our system is able to work in conjunction with a diverse range of heating systems and respective technologies that are typically employed in domestic settings.
7. Applicability	✓	The efficient application of AdaHeat in the various case-study systems considered in our evaluation illustrates its applicability.
8. Integrate simple economic control	✓	AdaHeat is able to consider variable energy import tariffs in the case of electricity-based heating.
9. Integrate advanced economic control	✓	Our system effectively integrates advanced economic control in the case of electricity-based heating and domestic IERs (for more details see the subsequent requirements below).
9(a) Coalition potential	✓	Our system is able to exploit the coalition potential that emerges in advanced economic control due to an effective and efficient respective heuristic planning approach.
9(a) Cost allocation	✓	Our system comes complete with a practical cost allocation to share the collective gains of collective advanced economic control among the coalition members that respects individual rationality and allocation efficiency.

In this thesis, we evaluated AdaHeat with respect to simple heating and simple economic control with data coming from a real house that employs underfloor heating. We showed that this particular case-study constitutes a challenging testbed for our approach both in terms of thermal modeling and control. Our evaluation results illustrated the benefits of incorporating adaptive gray-box thermal modeling in DHASs and the effectiveness of our approach in balancing heating cost and thermal discomfort. In the context of our evaluation procedure, we also provided a comprehensive comparison of existing state-of-the-art DHAS approaches against AdaHeat. By doing so, we showed that the latter leads to a more stable performance, in terms of Pareto efficiency, in various operation settings. In addition, we provided significant insights into the evaluated DHASs' usability in various operating settings, revealing the high usability of AdaHeat. Regarding advanced economic control, we demonstrated the effectiveness of our respective approach through evaluating with real data, in a contemporary market reality. In this context, we showed that collective advanced economic control can significantly improve heating cost-efficiency compared to independent advanced economic control, and even more when compared to independent simple economic/heating control. The above suggest that all the respective requirements as stated in Section 1.1.2 have been met. Table 7.1 provides a detailed evaluation of this work against the stated requirements.

As discussed above, in this work we also preoccupied ourselves with IER efficient control. In particular, we formulated solar tracking, ST, as a dynamic programming problem and introduce PreST, a dynamic programming approach for predictive ST. The exact solution to the dynamic programming formulation would provide the optimal ST trajectories. However, for reasons of computability, we approximate the optimal solution by a policy iteration method that we propose, that is suitable for large Markov decision processes (MDPs), along with specialized variants, utilizing freely available weather forecasts. Importantly, our methods make use of a generic and parameterizable tracker power consumption model that we developed. All our methods come with optimality or near-optimality guarantees, and our next-day policy comes complete with an expected PVS power output estimation, which is crucial for the smooth integration of PVSs into the electrical grid (as discussed in Section 6.4).

We demonstrated the efficiency of PreST against a number of commonly employed conventional ST techniques (i.e., chronological, sensor-based, and/or fixed-orientation). In particular, we showed that: (i) our next-day optimal fixed-orientation method significantly outperforms the yearly optimal one (enabling efficient manual ST); (ii) our *Myopic* method gives a significant advantage over the commonly employed *chronological* ST; and (iii) our proposed policy iteration method (i.e, STPI—Chapter 5) does consistently better, even though not by a wide margin, compared to, the also proposed by us, Myopic method. Finally, we showed that despite the small margin such small improvements are significant in real settings. Regarding our respective research requirements, all requirements as stated in Section 1.2.2 have been met. Table 7.2 provides a detailed evaluation of this work against the stated requirements.

Table 7.2: ST: Requirements evaluation

Requirement	Met	Comments
1. Generality	✓	Due to the general approach followed and the parameterizable ST consumption model proposed, our approach is able to work in conjunction with a diverse range of ST architectures that are typically employed in practice (e.g., HSAT; TSAT; TTDAT; or AADAT systems)
2. Applicability	✓	Our ST approach utilizes available weather forecasts that can actually come from online providers for free, and does not rely on expensive equipment, sensors or data. Moreover, our ST approach has appropriate computational requirements and efficiency (as discussed in Chapter 5) that allows it to be widely applicable in real settings with limited budget availability, limited resources, and operating time constrains.
3. Performance optimality	✓	All our ST methods come complete with optimal or near-optimal performance guaranties that lead to highly efficient ST performance.
3(a) Considering the prevailing weather conditions	✓	Our ST approach considers the weather conditions, utilizing available weather forecasts, to optimize the tracking trajectory.
3(b) Considering the consumption cost	✓	Our ST approach accounts for the tracking cost via an appropriate and generic ST consumption model that we devise.
3(c) Considering the maintenance cost	✓	Our ST consumption model is general enough to consider a modeling of the maintenance cost of the tracking system itself, in a straightforward manner (as discussed in Section 6.3)

Regarding the impact of this work, in this thesis we provided two distinct practical systems, namely AdaHeat and PreST, to increase the operational efficiency of domestic space heating systems and IERs, respectively. The results of this thesis could be used directly for the development of respective products with significant societal and commercial value. In particular, as discussed in Section 1.3.1, AdaHeat overcomes many of the limitations and shortcomings of previous DHAS approaches. Importantly, AdaHeat is also the first DHAS to integrate advanced economic control in order to maximally account for the emerging electricity market reality (as discussed in Section 1.3.1). As such, AdaHeat can serve as the guideline for the development of the next generation respective products. In this context, the results presented in this thesis provide significant lessons to this end. Furthermore, PreST is a novel ST approach that overcomes several limitations of previous approaches with minimum cost (as discussed in Section 6.4). In this context, PreST can serve as the basis for the development of web-based tools for efficient predictive ST—giving rise to a new form of cheep and efficient ST. In particular, the cost reduction of a smart house in the UK utilizing both PreST and AdaHeat that is employed with: (i) a 4kW space heating system, (ii) a 6kW wind turbine generator and (iii) a typical $72m^2$ azimuth-altitude dual axis tracking (AADAT) photovoltaic system can be estimated to over $25\pounds$ per month when participating in a collective advanced economic control schema and utilizes STPI (instead of employing simple economic/heating control and utilizing chronological AADAT tracking, respectively).[1] Last but not least, when taken together the contributions presented in this thesis (outlined in Section 1.3) consider a significant advance in the state-of-the-art with great value for the scientific community. Although of great importance, this value is not limited to the scope of efficient control within the broad energy sustainability agenda. In particular, apart from extending our understanding on the applicability of numerous artificial intelligence (AI) and AI-related techniques in the domain of our interest and providing specialized solutions and approaches, in this work we also proposed efficient control and modeling approaches with far more broader value for the scientific community. Importantly, we propose a new policy iteration approximation algorithm that considers the first alternative optimization dynamic programming algorithm for MDPs and is suitable for tackling MDPs with large state-action spaces. Nevertheless, despite these advances, many open problems remain with respect to efficient control towards an energy sustainable future. Given this, we identify three promising directions for future research to extend the scope of our work:

- First and foremost, a real world trial of both AdaHeat and PreST could be held as the most concrete way to evaluate our systems in real settings and to effectively consider any potentially missing detail. In particular, even though in our respective evaluation procedures we consider real data and account for numerous aspects of a real settings scenario, only a real-world trial could ensure that no essential detail is omitted. In addition, a real world trial would enable to explore the behavior of the users and collect valuable feedback from them. Furthermore, with respect to AdaHeat, a real world trial would enable to also explore the concept of optimal decision making regarding profitable coalition formation, and the respective market behavior, in real settings (with respect to collective advanced economic control). Moreover, depending on the region of interest, such trials can also consider variable export and import tariffs that are not perfectly known in advance and respective stochastic price predictions, which is a straightforward extension of AdaHeat (as discussed in Section 3.1.3).

- Furthermore, in recent years, a growing number of electric vehicles are emerging in the streets (Sperling (2013)). In this context, the owners of such vehicles could use the energy capacity of their vehicles in the context of domestic space heating advanced economic control. In particular, the respective energy capacity can serve as an energy buffer for profitable energy shifting within the AdaHeat framework. Although experimental work that consider the energy capacity of electric vehicles in the context of simple economic control is starting to emerge (Nguyen and Le

[1] Notably, to estimate the gains of PreST in the UK only the results considering the fully cloudy day in Crete have been utilized (i.e, the real day 4 in Table 6.1), utilizing the 2015 UK market electricity rates (see Section 4.2.2).

(2014)), the case of advanced economic control has not been explored yet. The most prominent challenge towards such an extension is to ensure the usability of the electric vehicle (i.e., to be adequately charged when the user needs it) while exploiting its energy capacity for profitable energy shifting. Such an extension could further enhance the efficiency of AdaHeat with respect to economic control in the case of electricity-based heating considering an emerging societal scenario.

- Given the ever growing electrification of transportation (Sperling (2013)) more and more attempts towards developing energy self-sustainable vehicles are emerging. To this end, vehicle-integrated IERs aim to provide the necessary energy to the vehicle whilst in motion. Such IERs are being integrated in numerous vehicle types ranging from cars, trains and boats, to even airplanes and helicopters (Corkish et al. (2013)). In this context, predictive IER control can enhance the efficiency and effectiveness of such IERs. In more detail, in such settings the possible route that a vehicle is following needs to also be taken into account in terms of maximizing the IER power output and minimizing the traveling energy consumption. For instance, if multiple paths can be followed to reach an intended destination the most favorable in terms of IER generation and traveling consumption can be chosen to minimize the net traveling cost. In addition, in the case of vehicle-integrated solar tracking systems, the dynamic position of the vehicle needs to also be taken into account in optimizing the tracking trajectory followed within the PreST framework. The most prominent challenge towards such an extension is to ensure the arrival of the vehicle to its destination within a feasible time and with minimum cost. This challenge is further exacerbated by potentially unpredicted rapid weather changes that need to, dynamically, be taken into account.

By meeting the above discussed challenges and extensions, the results related to efficient control of domestic space heating systems and IERs in this thesis can be further increased, towards meeting the aim for an energy sustainable future.

Appendix A

Closed Form Calculation of Expected Heating Cost

Estimating the expected heating cost in the case of advanced economic control is essential for developing respective systems. The particular closed form calculation is subject to the stochastic modeling of the (shared) intermittent energy resource power output followed. In this context, in Equation A.1 we provide a closed form of Equation 3.11, given our Gaussian-process-based respective modeling, i.e., $R' \sim N(\mu, \sigma)$.

$$
\begin{aligned}
&\mathbb{E}\left[\text{Cost}(\cdot)\right] = \\
&= \begin{cases}
P^{Sell}\left(\frac{\sigma}{K}\sqrt{\frac{2}{\pi}}\left(e^{\frac{(\mu - r_C^{\max})^2}{-2\sigma^2}} - e^{\frac{-\mu^2}{2\sigma^2}}\right) - \mu\right) & {\scriptstyle \text{Cons}(\cdot)=0} \\
\frac{\text{Cons}(\cdot)-\mu}{K}\left[\text{erf}\left(\frac{\text{Cons}(\cdot)-\mu}{\sqrt{2}\sigma}\right)(P^{Buy} - P^{Sell}) + \text{erf}\left(\frac{\mu}{\sqrt{2}\sigma}\right)P^{Buy}\right. & \\
\left. - \text{erf}\left(\frac{\mu - r_C^{\max}}{\sqrt{2}\sigma}\right)P^{Sell}\right] + \frac{\sigma}{K}\sqrt{\frac{2}{\pi}}\left[P^{Buy}\left(e^{\frac{(\text{Cons}(\cdot)-\mu)^2}{-2\sigma^2}} - e^{\frac{-\mu^2}{2\sigma^2}}\right)\right. & {\scriptstyle \text{Cons}(\cdot)<r_C^{\max} \wedge \text{Cons}(\cdot)>0} \\
\left. + P^{Sell}\left(e^{\frac{(\mu - r_C^{\max})^2}{-2\sigma^2}} - e^{\frac{(\text{Cons}(\cdot)-\mu)^2}{-2\sigma^2}}\right)\right] & \\
P^{Buy}\left(\frac{\sigma}{K}\sqrt{\frac{2}{\pi}}\left(e^{\frac{(\mu - r_C^{\max})^2}{-2\sigma^2}} - e^{\frac{-\mu^2}{2\sigma^2}}\right) + \text{Cons}(\cdot) - \mu\right) & {\scriptstyle \text{Cons}(\cdot)\geq r_C^{\max} \wedge \text{Cons}(\cdot)>0}
\end{cases}
\end{aligned}
\tag{A.1}
$$

where $K = \text{erf}\left(\frac{\mu}{\sqrt{2}\sigma}\right) - \text{erf}\left(\frac{\mu - r_C^{\max}}{\sqrt{2}\sigma}\right)$ and $\text{erf}(x) = \frac{2}{\sqrt{\pi}}\int_0^x e^{-t^2}dt$ (i.e., the error function).

Appendix B

Moment of Inertia of a Potentially Tilted Cuboid

Calculating the moment of inertia of a tilted cuboid is essential for many practical applications, such as modeling a solar tracking system (as discussed in Section 5.3.2). In this work we provide a *general* equation for calculating the moment of inertia of a potentially tilted, respecting its axes of rotation, cuboid. This equation is used within our solar tracking consumption model, proposed in Section 5.3.2), i.e., Equation 5.10.

B.1 General Cuboid Inertia Equation

The moment of inertia of a potentially tilted cuboid, I_c, can be calculated as:

$$I_c = \frac{m_c}{12}\left(l^2\cos^2(\beta) + d^2\sin^2(\beta) + w^2\right) \tag{B.1}$$

where m_c is the mass of the cuboid, β stands for its slope angle, and l, w and d stand for its length, width and depth respectively. The dimensions and the tilte angle of the cuboid are all defined with respect to its axis of rotation, as seen in Figure B.1.

B.2 Deriving the General Equation

In order to calculate the moment of inertia of a tilted cuboid, I_c, we calculated the moment of inertia of the imaginary cuboid that exactly contains the cuboid in question and we then subtracted the moment of inertia of the extra right angled prisms.

In particular the moment of inertia of the cuboid can be calculated through $I_c = I_{ALL} - (I_{R_1} + I_{R_2} + I_{R_3} + I_{R_4})$, where I_{ALL} stands for the moment of inertia of the imaginary cuboid, and I_{R_1}, I_{R_2}, I_{R_3} and I_{R_4} for the moment of inertia of the imaginary prisms, as

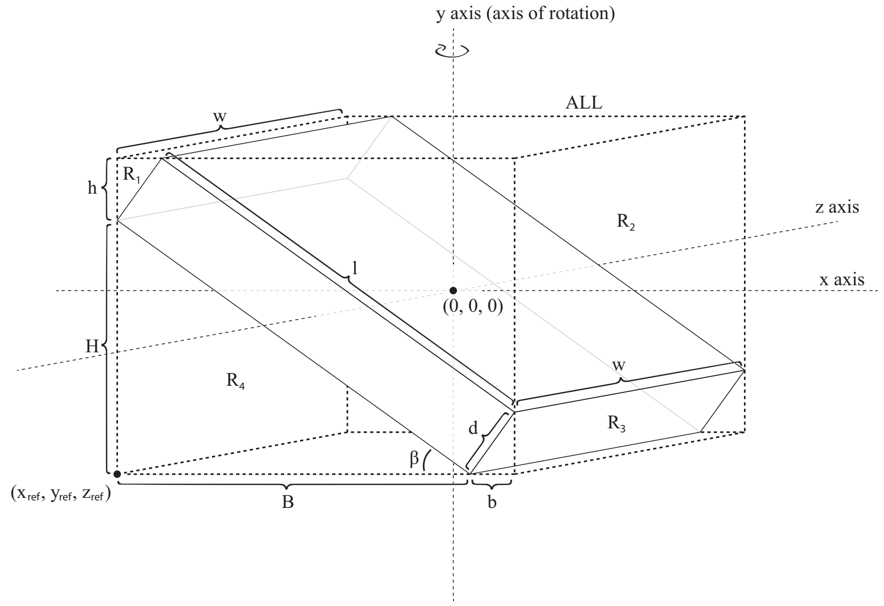

FIGURE B.1: Tilted Cuboid

illustrated in Figure B.1. Note, however, that $I_{R_1} = I_{R_3}$ and $I_{R_2} = I_{R_4}$, and hence our calculation can be simplified as:

$$I_c = I_{ALL} - (2I_{R_3} + 2I_{R_4}) \tag{B.2}$$

Also note here, that this equation is suitable for non-tilted cuboids as well, as in that case $I_{R_1} = I_{R_2} = I_{R_3} = I_{R_4} = 0$ and hence $I_c = I_{All}$ (and hence so does Eq B.1). In the following paragraphs we calculate the I_{ALL}, I_{R_3} and I_{R_4} above, in order to derive with I_c.

B.2.1 I_{ALL}

I_{ALL} can be calculated from the non-tilted cuboid Equation B.3 (Myers (1962)).

$$I_{ALL} = \frac{m_{ALL}}{12}\left((B+b)^2 + w^2\right) \tag{B.3}$$

where m_{ALL} is the mass of the imaginary cuboid and the quantity $(B+b)$ stands for the length of the imaginary cuboid, as seen in Figure B.1. The mass of the imaginary cuboid can be calculated based on its volume V_{ALL}, and its density ρ, via $m_{ALL} = \rho V_{ALL}$. The density of the imaginary cuboid will be the same as the density of the cuboid in question and can be calculated as $\rho = \frac{m_c}{lwd}$. The volume can be calculated as $V_{ALL} = (B+b)(H+h)w$, where the quantity $(H+h)$ stands for the depth of the imaginary cuboid, as seen in Figure B.1. H, h, B and b can be computed based on the dimensions and the tilt angle of the cuboid in question through Equations B.4, B.5, B.6 and B.7, respectively.

$$H = sin(\beta)l \tag{B.4}$$

$$h = cos(\beta)d \tag{B.5}$$

$$B = cos(\beta)l \tag{B.6}$$

$$b = sin(\beta)d \tag{B.7}$$

B.2.2 I_{R_3}

The moment of inertia of the right angled prism R_3 with respect to an axis of rotation that passes through its centroid and is parallel to the axis of rotation of the cuboid in question can be calculated through Equation B.8 (Myers (1962)) below:

$$I_{R_3}^{centroid} = \frac{m_{R_3}}{36}(2b^2 + 3w^2) \tag{B.8}$$

where m_{R_3} is the mass of the prism which can be calculated as $m_{R_3} = \frac{bhw}{2}\rho$.

Now the centroid of the right angled prism is defined as:

$$(-(x_{ref} + \tfrac{B}{3}), -(y_{ref} + \tfrac{H}{3}), -(z_{ref} + \tfrac{w}{2}))$$

As such, from Equation B.8 and by applying the parallel axes theorem (Kane and Levinson (1985)), we will have:

$$I_{R_3} = \frac{m_{R_3}}{36}(2b^2 + 3w^2) + m_{R_3}D_{R_3}^2 \tag{B.9}$$

where D_{R_3} is the distance of the centroid from the axis of rotation:

$$D_{R_3} = \frac{B + b}{2} - \frac{b}{3} \tag{B.10}$$

B.2.3 I_{R_4}

With the same method in computing I_{R_3} above, we will have:

$$I_{R_4} = \frac{m_{R_4}}{36}(2B^2 + 3w^2) + m_{R_4}D_{R_4}^2 \tag{B.11}$$

where $m_{R_4} = \frac{BHw}{2}\rho$ and for D_{R_4}:

$$D_{R_4} = \frac{B + b}{2} - \frac{B}{3} \tag{B.12}$$

By combining Equations B.2, B.9, B.11 and B.3, and by applying simple arithmetic operations we derive the general Equation B.1 for calculating the moment of inertia for potentially tilted cuboids.

References

N. Abas, A. Kalair, and N. Khan. Review of fossil fuels and future energy technologies. *Futures*, 69:31–49, 2015.

U. T. Aksoy and M. Inalli. Impacts of some building passive design parameters on heating demand for a cold region. *Building and Environment*, 41(12):1742–1754, 2006.

M. Alam, E. H. Gerding, A. Rogers, and S. D. Ramchurn. A scalable interdependent multi-issue negotiation protocol for energy exchange. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pages 1098–1104, 2015.

M. Alam, A. A. Panagopoulos, A. Rogers, N. R. Jennings, and J. Scott. Applying extended Kalman filters to adaptive thermal modelling in homes: Poster abstract. In *Proceedings of the 1st ACM International Conference on Embedded Systems for Energy-Efficient Buildings (Buildsys '14)*, pages 214–215, 2014.

M. Alam, S. D. Ramchurn, and A. Rogers. Cooperative energy exchange for the efficient use of energy and resources in remote communities. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 731–738. International Foundation for Autonomous Agents and Multiagent Systems, 2013.

A. T. Alan, E. Costanza, S. D. Ramchurn, J. Fischer, T. Rodden, and N. R. Jennings. Tariff agent: Interacting with a future smart energy system at home. *ACM Transactions on Computer-Human Interaction*, 23(4):25:1–25:28, 2016a.

A. T. Alan, M. Shann, E. Costanza, S. D. Ramchurn, and S. Seuken. It is too hot: An in-situ study of three designs for heating. In *Proceedings of the 31st ACM Conference on Human Factors in Computing Systems (CHI 2016)*, pages 5262–5273, 2016b.

K. K. Andersen, H. Madsen, and L. H. Hansen. Modelling the heat dynamics of a building using stochastic differential equations. *Energy and Buildings*, 31(1):13–24, 2000.

N. Armaroli and V. Balzani. The future of energy supply: Challenges and opportunities. *Angewandte Chemie International Edition*, 46(1–2):52–66, 2007.

F. Auffenberg, S. Stein, and A. Rogers. A personalised thermal comfort model using a bayesian network. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pages 2547–2553, 2015.

M. Augier and D. J. Teece, editors. *The Palgrave Encyclopedia of Strategic Management*. Macmillan Publishers, 2014.

P. Bacher and H. Madsen. Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings*, 43(7):1511–1522, 2011.

R. Bellman. *Dynamic Programming*. Princeton University Press, 1st edition, 1957.

A. Bemporad and M. Morari. Robust model predictive control: A survey. In *Robustness in identification and control*, pages 207–226. Springer, 1999.

D. P. Bertsekas. Dynamic programming and suboptimal control: A survey from ADP to MPC. *European Journal of Control*, 11(4–5):310–334, 2005.

A. Betz. *Introduction to the Theory of Flow Machines*. Elsevier, 2014.

J. C. Bezdek and R. J. Hathaway. Some notes on alternating optimization. In *2002 Advances in Soft Computing AFSS*, pages 288–300. Springer, 2002.

C. Böhringer, A. Löschel, and T. F. Moslener, Ulfand Rutherford. Eu climate policy up to 2020: An economic impact assessment. *Energy Economics*, 31:S295–S305, 2009.

G. Boyle, editor. *Renewable Energy: Power for a Sustainable Future*. Oxford University Press and Open University, 3rd edition, 2012.

J. Braun and G. Diderrich. Near-optimal control of cooling towers for chilled-water systems. *ASHRAE (American Society of Heating, Refrigerating and Air-Conditioning Engineers) Transactions*, 96(CONF-9006117–), 1990.

B. G. Brown, R. W. Katz, and A. H. Murphy. Time series models to simulate and forecast wind speed and wind power. *Journal of Climate and Applied Metereology*, 23 (8):1184–1195, 1984.

P. Bulman. Tesla's powerwall battery production requires 'super-charged' supply chain. *Renewable Energy Focus*, 16(5–6):126–127, 2015.

C. M. Burt, X. Piao, F. Gaudi, B. Busch, and N. Taufik. Electric motor efficiency under variable frequencies and loads. *Journal of irrigation and drainage engineering*, 134(2): 129–136, 2008.

T. Burton, D. Sharpe, N. Jenkins, and E. Bossanyi. *Wind Energy Handbook*. John Wiley & Sons, Ltd, 2011.

D. S. Callaway. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Management*, 50(5):1389–1400, 2009.

E. F. Camacho and C. B. Alba. *Model Predictive Control*. Springer, 2013.

E. Camacho and C. Bordons. *Model Predictive Control in the Process Industry*. Advances in industrial control. Springer-Verlag, 1995.

P. Chakraborty, M. Marwah, M. F. Arlitt, and N. Ramakrishnan. Fine-grained photovoltaic output prediction using a bayesian ensemble. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI 2012)*, pages 274–280, 2012.

G. Chalkiadakis, V. Robu, A. Kota, Ramachandra andRogers, and N. R. Jennings. Cooperatives of distributed energy resources for efficient virtual power plants. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, volume 2, pages 787–794, 2011.

T. P. Chang. The sun's apparent position and the optimal tilt angle of a solar collector in the northern hemisphere. *Solar Energy*, 83(8):1274–1284, 2009.

N. Chen, Z. Qian, I. T. Nabney, and X. Meng. Wind power forecasts using Gaussian processes and numerical weather prediction. *IEEE Transactions on Power Systems*, 29(2):656–665, 2014.

K. J. Chua, S. K. Chou, and W. Yang. Advances in heat pump systems: A review. *Applied Energy*, 87(12):3611–3624, 2010.

D. A. Coley and J. Penman. Second order system identification in the thermal response of real buildings. paper ii: Recursive formulation for on-line building energy management and control. *Building and Environment*, 27(3):269–277, 1992.

R. Corkish, M. Green, M. Watt, and S. Wenham. *Applied Photovoltaics*. Taylor & Francis, 2013.

T. H. Cormen, C. E. Leiserson, and C. Rivest, Ronald L andStein. *Introduction to Algorithms*, volume 2. MIT press Cambridge, 2001.

P. Couchman, B. Kouvaritakis, et al. MPC for stochastic systems. In *Assessment and Future Directions of Nonlinear Model Predictive Control*, pages 255–268. Springer, 2007.

K. Dalamagkidis, D. Kolokotsa, K. Kalaitzakis, and G. Stavrakakis. Reinforcement learning for energy conservation and comfort in buildings. *Building and Environment*, 42(7):2686–2698, 2007.

S. Dasgupta, F. Suwandi, S. Sahoo, and S. Panda. Dual axis sun tracking system with PV cell as the sensor, utilizing hybrid electrical characteristics of the cell to determine insolation. In *Proceedings of the 2010 IEEE International Conference on Sustainable Energy Technologies (ICSET 2010)*, pages 1–5, 2010.

M. Davis and M. Maschler. The kernel of a cooperative game. *Naval Research Logistics Quarterly*, 12(3):223–259, 1965.

F. de Nijs, M. T. J. Spaan, and M. M. de Weerdt. Best-response planning of thermostatically controlled loads under power constraints. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI 2015)*, pages 615–621, 2015.

K. S. Deffeyes. *Hubbert's Peak: The Impending World Oil Shortage*. Princeton University Press, 2008.

K. Deng, P. Barooah, P. G. Mehta, and S. P. Meyn. Building thermal model reduction via aggregation of states. In *Proceedings of the 2010 American Control Conference (ACC 2010)*, pages 5118–5123, 2010.

M. Diesendorf. *Greenhouse Solutions with Sustainable Energy*. A UNSW Press book. University of New South Wales Press, 2007.

T. Dietz, G. T. Gardner, J. Gilligan, P. C. Stern, and M. P. Vandenbergh. Household actions can provide a behavioral wedge to rapidly reduce us carbon emissions. In *Proceedings of the 2009 National Academy of Sciences (PNAS 2009)*, volume 106, pages 18452–18456, 2009.

N. Djongyang, R. Tchinda, and D. Njomo. Thermal comfort: A review paper. *Renewable and Sustainable Energy Reviews*, 14(9):2626–2640, 2010.

B. Dong, K. P. Lam, and C. Neuman. Integrated building control based on occupant behavior pattern detectionand local weather forecasting. In *Proceedings of the 12th International Building Simulation Conference (BS 2011)*, pages 14–17, 2011.

H. Doukas, K. D. Patlitzianas, K. Iatropoulos, and J. Psarras. Intelligent building energy management system using rule sets. *Building and Environment*, 42(10):3562–3569, 2007.

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):1246–1261, 2009.

J. H. Dudley and M. A. Piette. Solutions for summer electric power shortages: Demand response and its application in air conditioning and refrigerating systems. *Refrigeration, Air Conditioning, & Electric Power Machinery*, 29(1):1–4, 2008.

R. Dugad and U. B. Desai. A tutorial on hidden markov models. Technical report, Signal Processing and Artifical Neural Networks Laboratory, Dept. of Electrical Engineering, Indian Institute of Technology, 1996.

J.-P. Ebert, B. Stremmel, and A. Wiederhold, Eckhardt andWolisz. An energy-efficient power control approach for WLANs. *Journal of Communications and Networks*, 2(3): 197–206, 2000.

D. Einstein, E. Worrell, and M. Khrushch. Steam systems in industry: Energy use and energy efficiency improvement potentials. Technical report, Lawrence Berkeley National Laboratory, 2001.

C. Ellis, M. Hazas, and J. Scott. Matchstick: A room-to-room thermal model for predicting indoor temperature from wireless sensor data. In *Proceedings of the 12th IEEE/ACM International Conference on Information Processing in Sensor Networks (IPSN 2013)*, pages 31–42, 2013.

C. Ellis, J. Scott, M. Hazas, and J. Krumm. Earlyoff: Using house cooling rates to save energy. In *Proceedings of the 4th ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys '12)*, pages 39–41, 2012.

V. L. Erickson, S. Achleitner, and A. E. Cerpa. POEM: Power-efficient occupancy-based energy management system. In *Proceedings of the 12th IEEE/ACM International Conference on Information Processing in Sensor Networks (IPSN 2013)*, pages 203–216, 2013.

V. L. Erickson and A. 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, 2010.

T. Esram, P. L. Chapman, et al. Comparison of photovoltaic array maximum power point tracking techniques. *IEEE Transactions on Energy Conversion EC*, 22(2):439, 2007.

P. O. Fanger et al. *Thermal comfort. Analysis and Applications in Environmental Engineering.* Copenhagen: Danish Technical Press., 1970.

J. Farla, K. Blok, and L. Schipper. Energy efficiency developments in the pulp and paper industry: Across-country comparison using physical production data. *Energy policy*, 25(7):745–758, 1997.

D. Farris and J. Melsa. Energy savings for a solar heated and cooled building through adaptive optimal control. In *Proceedings of the 29th IEEE Conference on Decision and Control (CDC 1978) including the 17th Symposium on Adaptive Processes*, pages 206–213, 1978.

R. Z. Freire, G. H. Oliveira, and N. Mendes. Predictive controllers for thermal comfort optimization and energy savings. *Energy and Buildings*, 40(7):1353–1365, 2008.

H. Fukushima, T.-H. Kim, and T. Sugie. Adaptive model predictive control for a class of constrained linear systems based on the comparison model. *Automatica*, 43(2): 301–308, 2007.

S. F. Fux, A. Ashouri, M. J. Benz, and L. Guzzella. EKF based self-adaptive thermal model for a passive house. *Energy and Buildings*, 68, Part C(0):811–817, 2014.

P. X. Gao and S. Keshav. Optimal personal comfort management using SPOT+. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings (BuildSys '13)*, pages 1–8, 2013a.

P. X. Gao and S. Keshav. SPOT: A smart personalized office thermal control system. In *Proceedings of the 4th ACM International Conference on Future Energy Systems (e-Energy '13)*, pages 237–246, 2013b.

M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman & Co., 1979.

R. Gibbons. *A Primer In Game Theory.* Harvester Wheatsheaf, 1992.

M. Giuntoli and D. Poli. Optimized thermal and electrical scheduling of a large scale virtual power plant in the presence of energy storages. *IEEE Transactions on Smart Grid*, 4(2):942–955, 2013.

D. Y. Goswami and F. Kreith. *Energy Efficiency and Renewable Energy Handbook.* Taylor & Francis, 2nd edition, 2015.

F. Gottron. Energy efficiency and the rebound effect: Does increasing efficiency decrease demand. In *CRS Report for Congress*, pages 1–4, 2001.

M. A. Green, K. Emery, Y. Hishikawa, W. Warta, and E. D. Dunlop. Solar cell efficiency tables (version 45). *Progress in Photovoltaics: Research and Applications*, 23(1):1–9, 2015.

L. A. Greening, D. L. Greene, and C. Difiglio. Energy efficiency and consumption — the rebound effect — a survey. *Energy Policy*, 28(6–7):389–401, 2000.

M. S. Grewal and A. P. Andrews. *Kalman Filtering: Theory and Practice using MATLAB.* John Wiley & Sons, 2011.

M. Gupta, S. S. Intille, and K. Larson. Adding GPS-control to traditional thermostats: An exploration ofpotential energy savings and design challenges. In *Proceedings of the 7th International Conference on Pervasive Computing (Pervasive '9)*, pages 95–114, 2009.

C. Haiad, J. Peterson, P. Reeves, and J. Hirsch. Programmable thermostats installed into residential buildings: Predicting energy savings using occupant behavior & simulation. Technical report, Southern California Edison, 2004.

R. Haines and M. Myers. *HVAC Systems Design Handbook*. McGraw-Hill Education, fifth edition, 2009.

R. Halvgaard, N. K. Poulsen, H. Madsen, and J. B. Jørgensen. Economic model predictive control for building climate control in a smart grid. In *Proceedings of the 3rd IEEE PES Conference on Innovative Smart Grid Technologies (ISGT 2012)*, pages 1–6, 2012.

H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. A generalized battery model of a collection of thermostatically controlled loads for providing ancillary service. In *Proceedings of the 51th Annual Allerton Conference on Communication, Control and Computing (Allerton 2013)*, pages 551–558, 2013.

J. Hartman. *Turbocharging Performance Handbook*. MotorBooks International, 2007.

D. Hawkey. *Sustainable Urban Energy Policy*. Taylor & Francis, 2015.

S. S. Haykin. *Kalman Filtering and Neural Networks*. John Wiley & Sons, Inc., 2001. ISBN 0471369985.

G. P. Henze, C. Felsmann, and G. Knabe. Evaluation of optimal control for active and passive building thermal storage. *International Journal of Thermal Sciences*, 43(2): 173–183, 2004.

L. Hirth. The market value of variable renewables: The effect of solar wind power variability on their relative price. *Energy Economics*, 38:218–236, 2013.

M. I. Hoffert, K. Caldeira, G. Benford, D. R. Criswell, C. Green, H. Herzog, A. K. Jain, H. S. Kheshgi, K. S. Lackner, J. S. Lewis, et al. Advanced technology paths to global climate stability: Energy for a greenhouse planet. *Science*, 298(5595):981–987, 2002.

R. Howard. *Dynamic Programming and Markov Processes*. Published jointly by the Technology Press of the Massachusetts Instituteof Technology and, 1960.

H. Huang, L. Chen, M. Mohammadzaheri, E. Hu, and MinleiChen. Multi-zone temperature prediction in a commercial building usingartificial neural network model. In *Proceedings of the 10th IEEE International Conference on Control and Automation (ICCA 2013)*, pages 1896–1901, 2013.

M. Z. Jacobson, M. A. Delucchi, Z. A. F. Bazouin, Guillaume andBauer, C. C. Heavey, E. Fisher, S. B. Morris, D. J. Y. Piekutowski, and T. W. Vencill, Taylor A. andYeskoo. 100% clean and renewable wind, water, and sunlight, (WWS) all-sector energy roadmaps for the 50 United States. *Energy & Environmental Science*, 8:2093–2117, 2015.

A. B. Jaffe and R. N. Stavins. The energy-efficiency gap: What does it mean? *Energy Policy*, 22(10):804–810, 1994.

S. Jenner, F. Groba, and J. Indvik. Assessing the strength and effectiveness of renewable electricity feed-in tariffs in european union countries. *Energy Policy*, 52:385–401, 2013.

X. Jiang, B. Dong, L. Xie, and L. Sweeney. Adaptive Gaussian process for short-term wind speed forecasting. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI 2010)*, pages 661–666, 2010.

L. Kamal and Y. Z. Jafri. Time series models to simulate and forecast hourly averaged windspeed in quetta, pakistan. *Solar Energy*, 61(1):23–32, 1997.

T. Kane and D. Levinson. *Dynamics: Theory and Applications*. McGraw-Hill Series in Mechanical Engineering. McGraw-Hill Ryerson, Limited, 1985.

S. Karjalainen. Thermal comfort and use of thermostats in Finnish homes and offices. *Building and Environment*, 44(6):1237–1245, 2009.

K. Keesman. *System Identification: An Introduction*. Advanced Textbooks in Control and Signal Processing. Springer, 2011.

W. Kempton. Two theories of home heat control. *Cognitive Science*, 10(1):75–90, 1986.

M. Kintner-Meyer and A. Emery. Optimal control of an HVAC system using cold storage and building thermal capacitance. *Energy and Buildings*, 23(1):19–31, 1995.

W. Kleimingera, F. Matterna, and S. Santinib. Predicting household occupancy for smart heating control: A comparativeperformance analysis of state-of-the-art approaches. Technical report, ETH Zurich, 2013.

G. Kok, S. H. Lo, G.-J. Y. Peters, and R. A. Ruiter. Changing energy-related behavior: An intervention mapping approach. *Energy Policy*, 39(9):5280–5286, 2011.

R. Kota, G. Chalkiadakis, V. Robu, A. Rogers, and N. R. Jennings. Cooperatives for demand side management. In *Proceedings of the 7th European Conference on Artificial Intelligence (ECAI 2012)*, pages 969–974, 2012.

M. V. Kothare, V. Balakrishnan, and M. Morari. Robust constrained model predictive control using linear matrix inequalities. *Automatica*, 32(10):1361–1379, 1996.

T. Krishnamurti, D. Schwartz, A. Davis, B. Fischhoff, W. B. de Bruin, L. Lave, and J. Wang. Preparing for smart grid technologies: A behavioral decision research approach to understanding consumer expectations about smart meters. *Energy Policy*, 41:790–797, 2012.

N. R. Kristensen, H. Madsen, and S. B. Jørgensen. Parameter estimation in stochastic grey-box models. *Automatica*, 40(2):225–237, 2004.

J. Krumm and A. Brush. Learning time-based presence probabilities. In *Proceedings of the 9th International Conference on Pervasive Computing (Pervasive '11)*, pages 79–96, 2011.

J. Krumm and E. Horvitz. Predestination: Inferring destinations from partial trajectories. In *Proceedings of the 8th ACM International Conference on Ubiquitous Computing (UbiComp '06)*, pages 243–260, 2006.

M. Kummert, P. André, and J. Nicolas. Optimal heating control in a passive solar commercial building. *Solar Energy*, 69, Supplement 6:103–116, 2001.

K. Kurokawa. *Energy from the Desert: Practical Proposals for Very Large ScalePhotovoltaic Systems*. Earthscan LLC, 2012.

J. Langevin, J. Wen, and P. L. Gurian. Modeling thermal comfort holistically: Bayesian estimation of thermal sensation, acceptability, and preference distributions for office building occupants. *Building and Environment*, 69:206–226, 2013.

J. H. Lee. Model predictive control: Review of the three decades of development. *International Journal of Control, Automation and Systems*, 9(3):415–424, 2011.

J. H. Lee and N. L. Ricker. Extended Kalman filter based nonlinear model predictive control. *Industrial & Engineering Chemistry Research*, 33(6):1530–1541, 1994.

S. Lee, J. Speight, and S. Loyalka. *Handbook of Alternative Fuel Technologies, Second Edition*. Green Chemistry and Chemical Engineering. Taylor & Francis, 2014.

Y. I. Lee and B. Kouvaritakis. Robust receding horizon predictive control for systems with uncertain dynamics and input saturation. *Automatica*, 36(10):1497–1504, 2000.

N. S. Lewis. Toward cost-effective solar energy use. *science*, 315(5813):798–801, 2007.

G. Li and J. Shi. On comparing three artificial neural networks for wind speed forecasting. *Applied Energy*, 87(7):2313–2320, 2010.

X. Li and J. Wen. Review of building energy modeling for control and operation. *Renewable and Sustainable Energy Reviews*, 37:517–537, 2014.

Z. Li, X. Liu, and R. Tang. Optical performance of vertical single-axis tracked solar panels. *Renewable Energy*, 36(1):64–68, 2011.

M. Liserre, T. Sauter, and J. Y. Hung. Future energy systems: Integrating renewable energy sources into the smart power grid through industrial electronics. *IEEE Industrial Electronics Magazine*, 4(1):18–37, 2010.

B. Liu and R. Jordan. Daily insolation on surfaces tilted towards equator. *Solar Energy*, 10:53, 1961.

S. Liu and G. P. Henze. Experimental analysis of simulated reinforcement learning control for active and passive building thermal storage inventory: Part 1. Theoretical foundation. *Energy and Buildings*, 38(2):142–147, 2006.

P. Louka, G. Galanis, N. Siebert, G. Kariniotakis, P. Katsafados, I. Pytharoulis, and G. Kallos. Improvements in wind speed forecasts for wind power prediction purposes using Kalman filtering. *Journal of Wind Engineering and Industrial Aerodynamics*, 96 (12):2348–2362, 2008.

J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The smart thermostat: Using occupancy sensors to save energy in homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys '10)*, pages 211–224, 2010.

A. Luque and S. Hegedus. *Handbook of Photovoltaic Science and Engineering.* Wiley, 2011.

P. Mann, L. Gahagan, and M. B. Gordon. Tectonic setting of the world's giant oil and gas fields. *AAPG Memoir*, 78:15–105, 2003.

R. T. Marler and J. S. Arora. Survey of multi-objective optimization methods for engineering. *Structural and multidisciplinary optimization*, 26(6):369–395, 2004.

R. T. Marler and J. S. Arora. The weighted sum method for multi-objective optimization: new insights. *Structural and multidisciplinary optimization*, 41(6):853–862, 2010.

D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

D. Mayne and J. Rawlings. Correspondence: Correction to "Constrained model predictive control: Stability and optimality". *Automatica*, 37(3):483–, 2001.

A. Meier, C. Aragon, T. Peffer, D. Perry, and M. Pritoni. Usability of residential thermostats: Preliminary investigations. *Building and Environment*, 46(10):1891–1898, 2011.

T. Michalak, J. Tyrowicz, P. McBurney, and M. Wooldridge. Exogenous coalition formation in the e-marketplace based on geographical proximity. *Electronic Commerce Research and Applications*, 8(4):203–223, 2009.

M. A. Mohandes, T. O. Halawani, S. Rehman, and A. A. Hussain. Support vector machines for wind speed prediction. *Renewable Energy*, 29(6):939–947, 2004.

M. Morari and J. Lee. Model predictive control: Past, present and future. *Computers & Chemical Engineering*, 23(4):667–682, 1999.

N. Morel, M. Bauer, M. El-Khoury, and J. Krauss. Neurobat, a predictive and adaptive heating control system usingartificial neural networks. *International Journal of Solar Energy*, 21(2–3):161–201, 2001.

P.-D. Moroşan, R. Bourdais, D. Dumur, and J. Buisson. Building temperature regulation using a distributed model predictive control. *Energy and Buildings*, 42(9):1445–1452, 2010.

H. Mousazadeh, A. Keyhani, A. Javadi, H. Mobli, K. Abrinia, and A. Sharifi. A review of principle and sun-tracking methods for maximizing solar systems output. *Renewable and Sustainable Energy Reviews*, 13(8):1800–1818, 2009.

M. C. Mozer, L. Vidmar, and R. H. Dodier. The neurothermostat: Predictive optimal control of residential heating systems. In *Proceedings of the 10th Annual Conference on Advances in Neural Information Processing Systems (NIPS 1997)*, volume 7, pages 953–959, 1997.

J. A. Myers. *Handbook of Equations for Mass and Area Properties of Various Geometrical Shapes.* NAVWEPS report 7827. U.S. Naval Ordnance Test Station, 1962.

D. T. Nguyen and L. B. Le. Joint optimization of electric vehicle and home energy scheduling considering user comfort preference. *IEEE Transactions on Smart Grid*, 5 (1):188–199, 2014.

A. Niese, S. Lehnhoff, M. Tröschel, M. Uslar, C. Wissing, H. J. Appelrath, and M. Sonnenschein. Market-based self-organized provision of active power and ancillary services: An agent-based approach for smart distribution grids. In *2012 IEEE Complexity in Engineering (COMPENG 2012)*, pages 1–5, 2012.

S. Oberthür and H. E. Ott. *The Kyoto Protocol: international climate policy for the 21st century.* Springer Science & Business Media, 1999.

F. Oldewurtel, D. Gyalistras, M. Gwerder, C. Jones, A. Parisio, V. Stauch, B. Lehmann, and M. Morari. Increasing energy efficiency in building climate control using weather forecasts and model predictive control. In *Proceedings of the 10th REHVA World Congress on Sustainable Energy Use in Buildings (Clima 2010)*, number EPFL-CONF-169735, 2010a.

F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, 45:15–27, 2012.

F. Oldewurtel, A. Parisio, C. N. Jones, M. Morari, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and K. Wirth. Energy efficient building climate control using stochastic model predictive control and weather predictions. In *Proceedings of the 2010 American Control Conference (ACC 2010)*, pages 5100–5105, 2010b.

F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson, and M. Morari. Reducing peak electricity demand in building climate control using real-time pricing and model predictive control. In *Proceedings of the 49th IEEE Conference on Decision and Control (CDC 2010)*, pages 1927–1932, 2010c.

Z. O'Neill, S. Narayanan, and R. Brahme. Model-based thermal load estimation in buildings. In *Proceedings of the 4th IBPSA-USA National SimBuild Conference (SimBuild 2010)*, pages 474–481, 2010.

M. J. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT Press, 1994.

R. Pacheco, J. Ordóñez, and G. Martínez. Energy efficient design of building: A review. *Renewable and Sustainable Energy Reviews*, 16(6):3559–3573, 2012.

O. Palizban, K. Kauhaniemi, and J. M. Guerrero. Microgrids in active network management—part i: Hierarchical control, energy storage, virtual power plants, and market participation. *Renewable and Sustainable Energy Reviews*, 36:428–439, 2014.

A. A. Panagopoulos, G. Chalkiadakis, and E. Koutroulis. Predicting the power output of distributed renewable energy resources within a broad geographical region. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012)*, pages 981–986, 2012.

A. A. Panagopoulos. A novel method for predicting the power output of distributed renewable energy resources. Diploma thesis, Technical University of Crete, 2013.

A. A. Panagopoulos, M. Alam, A. Rogers, and N. R. Jennings. AdaHeat: A general adaptive intelligent agent for domestic heating control. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, pages 1295–1303, 2015a.

A. A. Panagopoulos and G. Chalkiadakis. Moment of inertia of potentially tilted cuboids. Technical report, University of Southampton, 2015.

A. A. Panagopoulos, G. Chalkiadakis, and N. R. Jennings. Towards optimal solar tracking: A dynamic programming approach. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI 2015)*, pages 695–701, 2015b.

P. Patrinos, P. Sopasakis, and H. Sarimveis. A global piecewise smooth newton method for fast large-scale model predictive control. *Automatica*, 47(9):2016–2022, 2011.

M. G. Patterson. What is energy efficiency?: Concepts, indicators and methodological issues. *Energy Policy*, 24(5):377–390, 1996.

M. Paulescu, E. Paulescu, P. Gravila, and V. Badescu. *Weather Modeling and Forecasting of PV Systems Operation.* Green Energy and Technology. Springer London, 2012.

T. Peffer, M. Pritoni, A. Meier, C. Aragon, and D. Perry. How people use thermostats in homes: A review. *Building and Environment*, 46(12):2529–2541, 2011.

J. A. Peterka and R. G. Derickson. Wind load design methods for ground-based heliostats and parabolic dish collectors. Technical report, Sandia National Labs, USA, 1992.

J. A. Peterka, N. Hosoya, B. Bienkiewicz, and J. E. Cernak. Wind load reduction for heliostats. Technical report, Colorado State University, USA, 1986.

A. Picard, R. S. Davis, M. Gläser, and K. Fujii. Revised formula for the density of moist air (CIPM-2007). *Metrologia*, 45(2):149, 2008.

S. Prívara, J. Cigler, Z. Váňa, F. Oldewurtel, C. Sagerschnig, and E. Žáčeková. Building modeling as a crucial part for building predictive control. *Energy and Buildings*, 56 (0):8–22, 2013.

S. Prívara, J. Široký, L. Ferkl, and J. Cigler. Model predictive control of a building heating system: The first experience. *Energy and Buildings*, 43(2–3):564–572, 2011.

S. Prívara, Z. Váňa, E. Žáčeková, and J. Cigler. Building modeling: Selection of the most appropriate model for predictive control. *Energy and Buildings*, 55:341–350, 2012.

D. Pudjianto, C. Ramsay, and G. Strbac. Virtual power plant and system integration of distributed energy resources. *Renewable power generation, IET*, 1(1):10–16, 2007.

M. L. Puterman and M. C. Shin. Modified policy iteration algorithms for discounted markov decision problems. *Management Science*, 24(11):1127–1137, 1978.

M. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley, 2014.

S. Qin and T. A. Badgwell. A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7):733–764, 2003.

P. Radecki and B. Hencey. Online building thermal parameter estimation via unscented Kalman filtering. In *Proceedings of the 2012 American Control Conference (ACC 2012)*, pages 3056–3062, 2012.

P. Radecki and B. Hencey. Online thermal estimation, control, and self-excitation of buildings. In *Proceedings of the 52nd IEEE Conference on Decision and Control (CDC 2013)*, pages 4802–4807, 2013.

S. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Putting the "smarts" into the smart grid: A grand challenge for artificial intelligence. *Communications of the ACM*, 55(4):86–97, 2012.

S. D. Ramchurn, M. Osborne, O. Parson, T. Rahwan, S. Maleki, S. Reece, T. D. Huynh, M. Alam, J. E. Fischer, T. Rodden, et al. Agentswitch: towards smart energy tariff selection. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 981–988. International Foundation for Autonomous Agents and Multiagent Systems, 2013.

S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings. Agent-based control for decentralised demand side management in the smart grid. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, volume 1, pages 5–12, 2011.

C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning Series)*. The MIT Press, 2005.

I. Reda and A. Andreas. Solar position algorithm for solar radiation applications. *Solar Energy*, 76(5):577–589, 2004.

K. Roebuck. *Energy Storage: High-impact Strategies - What You Need to Know: Definitions, Adoptions, Impact, Benefits, Maturity, Vendors*. Emereo Publishing, 2012a.

K. Roebuck. *Photovoltaics (PV): High-impact Strategies - What You Need to Know: Definitions, Adoptions, Impact, Benefits, Maturity, Vendors*. Emereo Publishing, 2012b.

A. Rogers, S. Ghosh, R. Wilcock, and N. R. Jennings. A scalable low-cost solution to provide personalised home heating advice to households. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings (BuildSys '13)*, pages 1–8, 2013.

A. Rogers, S. Maleki, S. Ghosh, and N. R. Jennings. Adaptive home heating control through Gaussian process prediction and mathematical programming. In *Proceedings of the 2nd International Workshop on Agent Technology for Energy Systems (ATES 2011)*, pages 71–78, 2011.

T. H. Roos. A wind loading correlation for an isolated square heliostat part 2: Moments and side forces. In *Proceedings of the 1st Southern African Solar Energy Conference (SASEC 2012)*, pages 1–8, 2012.

A. E. Ruano, E. M. Crispim, E. Z. E. Conceiçao, and M. M. J. R. Lúcio. Prediction of building's temperature using neural networks models. *Energy and Buildings*, 38(6): 682–694, 2006.

W. Saad, Z. Han, H. V. Poor, and T. Basar. Game-theoretic methods for the smart grid: An overview of microgrid systems, demand-side management, and smart grid communications. *IEEE Signal Processing Magazine*, 29(5):86–105, 2012.

S. Scellato, M. Musolesi, C. Mascolo, V. Latora, and A. T. Campbell. NextPlace: A spatio-temporal prediction framework for pervasive systems. In *Proceedings of the 9th International Conference on Pervasive Computing (Pervasive '11)*, pages 152–169, 2011.

D. Schmeidler. The nucleolus of a characteristic function game. *SIAM Journal on applied mathematics*, 17(6):1163–1170, 1969.

P. O. Scokaert and J. B. Rawlings. Feasibility issues in linear model predictive control. *AIChE Journal*, 45(8):1649–1659, 1999.

J. Scott, A. J. Bernheim Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, and N. Villar. Preheat: Controlling home heating using occupancy prediction. In *Proceedings of the 13th ACM International Conference on Ubiquitous Computing (UbiComp '11)*, pages 281–290, 2011.

M. Sengupta, P. Gotseff, and T. Stoffel. Evaluation of photodiode and thermopile pyranometers for photovoltaic applications. *27th EUPVSEC*, 1:3705–3708, 2012.

M. Shann and S. Seuken. Adaptive home heating under weather and price uncertainty using GPS and MDPS. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2014)*, pages 821–828, 2014.

L. S. Shapley and M. Shubik. Quasi-cores in a monetary economy with nonconvex preferences. *Econometrica: Journal of the Econometric Society*, pages 805–827, 1966.

M. Shipworth, S. K. Firth, M. I. Gentry, A. J. Wright, D. T. Shipworth, and K. J. Lomas. Central heating thermostat settings and timing: Building demographics. *Building Research & Information*, 38(1):50–69, 2010.

F. P. Sioshansi. *Future of Utilities - Utilities of the Future: How Technological Innovations in Distributed Energy Resources will Reshape the Electric Power Sector*. Elsevier Science, 2016.

J. Široký, F. Oldewurtel, J. Cigler, and S. Prívara. Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy*, 88 (9):3079–3087, 2011.

K. A. Small and K. Van Dender. The effect of improved fuel economy on vehicle miles traveled: estimating the rebound effect using US state data, 1966–2001. Technical report, University of California Energy Institute, USA, 2005.

T. S. Soderstrom and P. G. Stoica. *System Identification*. Prentice Hall International Series In Systems And Control Engineering. Prentice Hall, 1989.

S. S. Soman, H. Zareipour, O. Malik, and P. Mandal. A review of wind power and wind speed forecasting methods with different time horizons. In *Proceedings of the 2010 North American Power Symposium (NAPS 2010)*, pages 1–8, 2010.

D. Sperling. *Future drive: Electric vehicles and sustainable transportation*. Island Press, 2013.

D. Sperling and D. Gordon. *Two billion cars: Driving toward sustainability.* Oxford University Press, 2009.

G. Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12): 4419–4426, 2008.

R. Sutton and A. Barto. *Reinforcement Learning: An Introduction.* A Bradford book. Bradford Book, 1998.

J. Torriti, M. G. Hassan, and M. Leach. Demand response experience in europe: Policies, programmes and implementation. *Energy*, 35(4):1575–1583, 2010.

D. Urieli and P. Stone. A learning agent for heat-pump thermostat control. In *Proceedings of The 12th International Conference on Autonomous Agents And Multi-Agent Systems (AAMAS 2013)*, pages 1093–1100, 2013.

J. Van Nunen. A set of successive approximation methods for discounted markovian decision problems. *Zeitschrift fuer operations research*, 20(5):203–208, 1976.

J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton university press, 2007.

P. Waide, B. Lebot, and M. Hinnells. Appliance energy standards in europe. *Energy and Buildings*, 26(1):45–67, 1997.

S. Wang and Z. Ma. Supervisory and optimal control of building HVAC systems: A review. *HVAC&R Research*, 14(1):3–32, 2008.

Y. Wang and S. Boyd. Fast model predictive control using online optimization. *IEEE Transactions on Control Systems Technology*, 18(2):267–278, 2010.

E. Worrell, N. Martin, and L. Price. Potentials for energy efficiency improvement in the us cement industry. *Energy*, 25(12):1189–1214, 2000.

J. Yan, K. Li, E.-W. Bai, J. Deng, and A. M. Foley. Hybrid probabilistic wind power forecasting using temporally local Gaussian process. *IEEE Transactions on Sustainable Energy*, 7(1):87–95, 2016.

J. Yan, Z. Yang, K. Li, and Y. Xue. A variant Gaussian process for short-term wind power forecasting based on TLBO. In *Proceedings of the 2014 International Conference on Life System Modeling and Simulation (LSMS 2014) and the 2014 International Conference on Intelligent Computing for Sustainable Energy and Environment (ICSEE 2014)*, pages 165–174, 2014.

R. Yao. *Design and management of sustainable built environments.* Springer, 2013.

Y. Ye, Y. Zheng, Y. Chen, J. Feng, and X. Xie. Mining individual life pattern based on location history. In *Proceedings of the 10th IEEE International Conference on Mobile Data Management: Systems, Services and Middleware (MDM 2009)*, pages 1–10, 2009.

L. Zadeh. Optimality and non-scalar-valued performance criteria. *IEEE Transactions on Automatic Control*, 8(1):59–60, 1963.

A. Zahedi. Australian renewable energy progress. *Renewable and Sustainable Energy Reviews*, 14(8):2208–2213, 2010.

M. Zamo, O. Mestre, P. Arbogast, and O. Pannekoucke. A benchmark of statistical regression methods for short-term forecasting of photovoltaic electricity production. part ii: Probabilistic forecast of daily production. *Solar Energy*, 105(0):804–816, 2014.

O. Zehner. *Green illusions: The dirty secrets of clean energy and the future of environmentalism.* University of Nebraska Press, 2012.

S. Zionts. Multiple criteria mathematical programming: an updated overview and several approaches. In *Multiple Criteria Decision Making and Risk Analysis Using Microcomputers*, pages 7–60. Springer, 1989.