

Advanced Economic Control of Electricity-based Space Heating Systems in Domestic Coalitions with Shared Intermittent Energy Resources

ATHANASIOS ARIS PANAGOPOULOS, UC Berkeley, USA

SASAN MALEKI, University of Southampton, UK

ALEX ROGERS, University of Oxford, UK

MATTEO VENANZI, Microsoft, UK

NICHOLAS R. JENNINGS, Imperial College, UK

Over the past few years, domestic heating automation systems (DHASs) that optimize the domestic space heating control process with minimum user-input, utilizing appropriate occupancy prediction technology, have emerged as commercial products (e.g. the smart thermostats from Nest and Honeywell). At the same time, many houses are being equipped with, potentially grid-connected, intermittent energy resources (IERs), such as rooftop photovoltaic systems and/or small wind turbine generators. Now, in many regions of the world, such houses can sell energy to the grid but at a lower price than the price of buying it. In this context, and given the anticipated increase in electrification of heating, the next generation DHASs need to incorporate advanced economic control (AEC). Such AEC can exploit the energy buffer that heating loads provide, in order to shift the consumption of electricity-based heating systems to follow the intermittent energy generation of the house. By so doing, the energy imported from the grid can be minimized and considerable monetary gains for the household can be achieved, without affecting the occupants' schedule. These benefits can be amplified still further in domestic coalitions, where a number of houses come together and share their IER generation to minimize their cumulative grid energy import.

Given the above, in this work we extend a state-of-the-art DHAS, to propose AdaHeat+, a practical DHAS, that, for the first time, incorporates AEC. Our work is applicable to both individual houses and domestic coalitions and comes complete with an allocation mechanism to share the coalition gains. Importantly, we propose an effective heuristic heating schedule planning approach for collective AEC which: (i) has a complexity that scales in a linear and parallelizable manner with the coalition size, and (ii) enables AdaHeat+ to handle the distinct preferences, in balancing heating cost and thermal discomfort, of the households. Our approach relies on stochastic IER power output predictions. In this context, we propose a simple and effective formulation for the site-specific calibration of such predictions based on adaptive Gaussian process modeling. Finally, we demonstrate the effectiveness of AdaHeat+ through real data evaluation, to show that collective AEC can improve heating cost-efficiency by up to 60%, compared to independent AEC (and even more when compared to no-AEC).

Categories and Subject Descriptors: I.2.11 [Artificial Intelligence-Distributed Artificial Intelligence]: Intelligent agents

General Terms: Design, Algorithms, Performance, Theory

Additional Key Words and Phrases: Advanced Economic Control, Control, Space Heating, Smart Grid, Energy Systems

ACM Reference Format:

Athanasios Aris Panagopoulos, Sasan Maleki, Alex Rogers, Matteo Venanzi, and Nicholas R. Jennings, 2015. Advanced Economic Control of Space Heating Loads in Domestic Coalitions with Shared Intermittent Energy Resources *ACM Trans. Intell. Syst. Technol.* V, N, Article A (January YYYY), 24 pages.

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

This work is supported by the UK Research Council for the ORCHID project, grant EP/I011587/1. Authors' addresses: A. A. Panagopoulos, UC Berkeley, 448 Sutardja Dai Hall, Berkeley Energy & Climate Institute, Berkeley, CA 94720, USA email: apanagopoulos@outlook.com; S. Maleki, Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, UK, email: sm9g09@ecs.soton.ac.uk; A. Rogers, Department of Computer Science, University of Oxford, Oxford, OX1 3QD, UK, email: alex.rogers@cs.ox.ac.uk; M. Venanzi, Microsoft, London, 2 Wharfedale Square, EC1N 2ST UK, email: mavena@microsoft.com; N. R. Jennings, Imperial College, Level 4, Faculty Building, South Kensington Campus, LONDON SW7 2AZ, email: n.jennings@imperial.ac.uk.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 2157-6904/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

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 is related to space heating.¹ As such, 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 CO₂ 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; Široký 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 [Široký et al. 2011].² To this end, programmable thermostats have been proposed as a means to enable more efficient heating control compared to traditional static thermostats, enabling 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, several 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; Shipworth et al. 2010]). 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.³

At the same time, 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].⁴ Moreover, in many regions, such as in several European countries and US states,⁵ 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.⁶ 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 [Sioshansi 2016].⁷ Now, 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

¹According to “Annual Energy Review 2011”, *US Energy Information Administration*, available on-line at www.eia.gov/totalenergy/data/annual/archive/038411.pdf (retrieved on 01/2016), and “United Kingdom Housing Energy Fact File”, *UK Department 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 01/2016).

²In 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 set-point with minimum oscillations [Wang and Ma 2008].

³www.nest.com; www.honeywell.com; and www.hivehome.com

⁴This 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].

⁵According to www.dsireusa.org and www.res-legal.eu/en (accessed on 01/2016)

⁶For 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.

⁷The 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].

they: (i) store their produced energy to follow their demand, or (ii) shift their electricity demand to follow their energy production. In either case, they would minimize their energy import and achieve financial gains. However, appropriate energy storage technology,⁸ that could facilitate the former is not available or cost efficient to install at the moment in domestic settings.⁹ Moreover, the latter approach 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]. In this context, and given the anticipated increase in electrification of heating [Hawkey 2015], the next generation DHASs needs to incorporate advanced economic control (AEC). Such AEC can exploit the aforementioned shifting potential in houses with electricity-based heating systems and grid-connected IERs in order to minimize the energy import and achieve respective financial gains. Importantly, these benefits of AEC can be amplified in domestic coalitions, where a number of houses come together to share their energy generation and coordinately shift their heating consumption in order to further minimize the energy drawn from the grid. Now, having this happen, requires the market to allow such groups to act as an individual, and can be facilitated by forming a micro-grid [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 smart metering [Niese et al. 2012]. That said, smart meters are already being installed in many countries [Krishnamurti et al. 2012] and dynamic on-line energy exchange is being explored through initiatives such as Piclo, Clickpower, and Powershop.¹⁰ Hence, 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].

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 2013; 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., [Panagopoulos et al. 2015; Shann and Seuken 2014; Rogers et al. 2011; Halvgaard et al. 2012; Oldewurtel et al. 2010]), 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]), 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 [Panagopoulos et al. 2012].

Against this background, in this work we propose a practical DHAS, AdaHeat+, that incorporates AEC. It considers, for the first time, issues associated with domestic grid-connected IERs and export tariffs. Our work is applicable to both coalitions and single houses (as an extreme case of a

⁸Existing energy storage technologies include, among others, battery, flywheels and super-capacitors [Liserre et al. 2010].

⁹In 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].

¹⁰www.openutility.com/piclo; www.clickpower.in; www.powershop.com.au

single-house coalition). In particular, we propose the formation of coalitions of houses that coordinate their heating system operation ahead of time, utilizing IER stochastic predictions, in order to minimize their energy import (or, more formally, to buy and sell energy to the grid when the tariffs are favorable),¹¹ as an aggregate. To do so, we extend a state-of-the-art single-house DHAS, AdaHeat [Panagopoulos et al. 2015], that considers simple economic control (i.e. accounting for variable energy import tariffs) and tackles many of the respective challenges (as further discussed in Sec. 2.1), to account for AEC (i.e., considering also domestic grid-connected IER generation capacity and export tariffs). We add to it by proposing an effective heuristic heating schedule planning approach for AEC in domestic coalitions. Our solution: (i) has a complexity that scales in a linear and parallelizable manner with the size of the coalition, and (ii) handles the distinct preferences, in balancing heating cost and thermal discomfort, of the households. Our planning approach relies on stochastic IER power output predictions. In this context, we utilize a state-of-the-art site-specific calibration technique, to improve such predictions, based on Gaussian process (GP) modeling, and propose a simple and effective respective formulation. Moreover, AdaHeat+ incorporates a practical cost allocation mechanism to share the gains of the coalition that respects individual rationality and allocation efficiency. Finally, we demonstrate the effectiveness of AdaHeat+ through real data evaluation in a contemporary market reality with flat import and export tariffs. Specifically, we show that collective AEC can improve heating cost-efficiency by up to 60%, compared to independent AEC, and even more when compared to AdaHeat's independent simple economic/heating control.

In more detail, our main contribution to the state-of-the-art is as follows:

- We are the first to show how AEC 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 AEC, 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 DHAS that accounts for AEC and demonstrate its effectiveness through evaluating with real data, in a contemporary market reality. We show that collective AEC can significantly improve heating cost-efficiency compared to independent AEC, and even more when compared to independent simple economic/heating control.

The remainder of this paper is organized as follows: First, in Sec. 2, we provide background material and literature review, relevant to our work. In Sec. 3 we detail our DHAS approach. Then, in Sec. 4, we provide a thorough evaluation of our DHAS approach (providing also an instantiation of our approach for our case study system) and, finally, Sec. 5 concludes. Table I summarizes the main abbreviations and symbols used in this paper.

Table I. List of main abbreviations and symbols.

AEC	Advanced Economic Control	C	Coalition set	T^{IN}	Indoor air temperature
DHAS	Domestic Heating Automation System	C^{Eff}	Heating system efficiency	T^{OUT}	Outside ambient temperature
GP	Gaussian Process	H	Planning horizon intervals set	T^{SP}	Set-point temperature
MPC	Model Predictive Control	J	Optimization unifying cost	u	Heating action
NWP	Numerical Weather Predictions	p^{Buy}	Import tariff	δ	Interval length
SEC	Simple Economic Control	p^{Sell}	Export tariff	λ	Weighting parameter
SHC	Simple Heating Control	R'	Future power output	τ	Interval id

¹¹ Although our work is motivated by the use of flat tariffs it will work equally well in variable tariff settings (as further discussed in Sec. 3). Nevertheless, evaluating the efficiency of our approach in such settings considers future work (Sec. 5).

2. BACKGROUND MATERIAL

In this section we provide the background material and literature review that is relevant to our work. In more detail, in Sec. 2.1, we provide a general discussion on space heating automation with a specific focus on the particular challenges that arise in domestic settings and its importance in facilitating AEC. Then, in Sec. 2.2, we provide a general discussion on IER stochastic prediction and Gaussian processes which are the bases of our case study IER power output prediction approach.

2.1. Domestic Space Heating Control Automation and Advanced Economic Control

Energy systems 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; Kummert et al. 2001]). More recently, however, with the onset of ever-increasing house instrumentation, more research has focused on domestic systems, DHASs, as well (e.g., [Mozer et al. 1997; Scott et al. 2011]). Indeed, such systems are now commercial products. Nevertheless, DHASs provide additional challenges over their, more explored, non-domestic counterparts.

In more detail, a particularly dominant challenge that arises in DHASs is tackling the occupancy schedule uncertainty. In contrast to non-domestic buildings with a generally fixed occupancy schedule, occupancy in domestic settings is highly dynamic. To this end, several approaches to predict occupancy have been proposed over time. These can, in general, be classified as either schedule-based or context-aware [Kleiminger et al. 2014]. The former predict occupancy based only on the history of the occupancy schedule, while the latter (also) utilize observations of the current context of the occupants. Although the latter have the potential to demonstrate higher predictive accuracy, as more information is being considered, they are typically more complex and require additional instrumentation that is always a concern in domestic settings. Moreover, schedule-based approaches have been shown to achieve high predictive accuracy and to be sufficient for effective domestic space heating control automation [Kleiminger et al. 2014; Scott et al. 2011; Panagopoulos et al. 2015].

Even the most accurate occupancy prediction algorithm is likely to have a degree of uncertainty over the predicted schedule. This is often modeled in the form of probabilistic occupancy estimates. In this context, balancing thermal discomfort and heating cost is inevitable in order not to significantly increase the heating cost. 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 cost with very small (or not-existent) improvements in the realized discomfort. Given this, effectively balancing heating cost and thermal discomfort with respect to the occupant preferences becomes an essential task for DHASs. In particular, only recent work in AdaHeat [Panagopoulos et al. 2015] has effectively tackled this challenge through appropriate planning and model predictive control (MPC), enabling an adaptive scheme to meet the occupant preferences. We should note here that the wide family of MPC (or receding horizon control) [Camacho and Alba 2013] has been shown to be very efficient and has been extensively used in space heating control automation (e.g., [Panagopoulos et al. 2015; Halvgaard et al. 2012; Oldewurtel et al. 2010; Scott et al. 2011]). For these reasons, it is also employed in this work. The particular success of MPC can be mainly attributed to two reasons: (i) an off-line computation of a control law is almost impossible in this setting [Panagopoulos et al. 2015]; and (ii) the slow nature of building thermal processes do not raise any of the stability issues that are often a concern with MPC [Široký et al. 2011].

Now, a detailed model of the system under control is crucial in MPC [Camacho and Alba 2013]. In this context, thermal modeling is an essential part of MPC-based DHASs. In general, a thermal model predicts the thermal response of a building based on: (i) the current thermal state of the building; (ii) the heating control actions to be executed; and, potentially, (iii) information variables regarding exogenous processes that affect the thermal process (e.g., incident solar radiation, outside or adjacent buildings' temperature). However, it is impossible to model all the factors that affect the thermal response of a building. This becomes even more apparent in domestic settings where the occupant's activity can drastically alter the thermal dynamics (e.g., opening a window, cooking,

or operating an auxiliary fan heater). That said, adaptive thermal modeling has been shown to be resilient and effective in such highly dynamic settings and several adaptive thermal modeling approaches have been proposed over time (e.g., [Ruano et al. 2006; Reece et al. 2014; Alam et al. 2014]). These can be classified into either *black-box* or *gray-box* approaches [Prívara et al. 2013]. The former use statistical or machine learning techniques (e.g., simple polynomial curve fits, neural networks, or support vector machines) to model the thermal dynamics with minimum need for prior knowledge of the space characteristics (e.g., [Huang et al. 2013; Ruano et al. 2006]). However, these approaches lack a physical interpretation and, hence, typically require a large amount of training data in order to demonstrate adequate performance [Li and Wen 2014]. On the other hand, adaptive *gray-box* approaches rely on simplified physical equations and, as such, do not suffer from the above drawbacks [Li and Wen 2014]. Such approaches have already shown their potential for thermal modeling in highly dynamic settings (e.g., [Reece et al. 2014; Alam et al. 2014]). Moreover, AdaHeat demonstrates the value of such an approach within the context of a DHAS.

As outlined above, AdaHeat tackles many of the challenges that arise in domestic settings and, hence, meets many of the requirements for facilitating advanced AEC: (i) it comes complete with an adequate schedule-based occupancy prediction algorithm; (ii) it effectively and efficiently balances heating cost and thermal discomfort, respecting the occupant preferences; (iii) it considers a robust gray-box adaptive thermal modeling; and (iv) it demonstrates high performance efficiency when compared to state-of-the-art DHASs (see [Panagopoulos et al. 2015]). However, AdaHeat has considerable limitations with respect to economic control: (i) it considers simple economic control, merely accounting for variable energy import tariffs; (ii) it does not exploit the potential that coalitions provide in (simple or advanced) economic control, and considers a single-house system; as such, (iii) it does not provide an effective approach in balancing thermal discomfort and heating cost when houses share their cost within a coalition, and (iv) it does not provide a cost allocation mechanism. Nevertheless, AdaHeat is the most suitable point of departure for this work.

2.2. Predicting IER Power Output and Gaussian Processes

An essential part of our AEC approach is a stochastic, short-term (up to 12 hours ahead), prediction of the shared IER power output. In general, the power output of IERs depends on the prevailing weather conditions. For instance, the power output of a photovoltaic system depends mostly on the irradiance incident to the photovoltaic module and the operating temperature of the module. In particular, the power output of a photovoltaic system increases with greater levels of incident solar irradiance and lower ambient temperatures (that lead to generally lower operating temperatures) [Luque and Hegedus 2011]. Furthermore, the power output of a wind turbine generator depends mostly on the prevailing wind speed and wind temperature, generally favoring higher wind speeds and lower wind temperatures (cooler air is denser, increasing power output) [Burton et al. 2011]. In addition, advanced IER models exist to transform the 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.¹² 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.

Now, a well-known state-of-the-art approach for handling non-linear regression problems that has widely and effectively been used for IER power output predictions (e.g., [Chen et al. 2014; Zamo et al. 2014]) is Gaussian process (GP) regression. Apart from its demonstrated efficiency compared to state-of-the-art hybrid IER power output prediction approaches ([Chen et al. 2014]), GP regression also provides stochastic predictions in a principled manner. Hence, it is also employed in this work. Formally, a GP can be defined as a stochastic process defining a distribution over functions $f : X \rightarrow Y \in \mathbb{R}$ such that any finite subset $X' \subset X$ is multivariate Gaussian distributed. Specifically, a GP is completely defined by a mean function, $m(\cdot)$ and a positive semi-definite covariance function (or kernel), $K(\cdot, \cdot)$. 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., $\mathbf{y} = f(\mathbf{x}) + \epsilon$, 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 $f \sim \mathcal{N}(m(\mathbf{x}|\theta), K(\mathbf{x}, \mathbf{x}|\theta) + \sigma^2 I)$, 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, the corresponding predictive posterior distribution can be evaluated in closed form as a multivariate Gaussian distribution, while the hyper-parameters are typically learned by maximizing the respective log marginal likelihood. This maximizer can be searched using a standard gradient-based optimization method—for more details see Rasmussen and Williams [2005].

3. ADVANCED ECONOMIC CONTROL OF ELECTRICITY-BASED SPACE HEATING SYSTEMS

In this work we propose a practical DHAS, AdaHeat+, for AEC of electricity-based space heating systems in domestic coalitions with shared IERs. 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. In the proposed scheme, the coalition C is able to buy energy from the grid (at any instance) at a higher import tariff, P^{Buy} , than the export tariff, P^{Sell} . In essence, this price difference motivates the coalition to use the shared IER generation to fulfill as much of its demand as possible. Although our work is motivated by the use of flat import and export tariffs (that are known in advance), our approach can work in conjunction with variable tariffs as well, even if they are not generally known in advance [Jenner et al. 2013], as long as an appropriate tariff prediction approach is considered. In this context, we propose a collective AEC scheme, where IER stochastic predictions are used in order for the coalition members, indexed by $i \in \{1, 2, \dots, |C|\}$, to coordinate their heating system operation ahead in time, in a collective manner. In particular, given derived stochastic IER power output predictions along with detailed house models, we employ an adaptive MPC approach for optimizing the heating control process of the coalition, that effectively considers up-to-date, weather variables, IER power output (and potentially P^{Buy} and P^{Sell}) predictions, in real time; and adaptive gray-box thermal modeling (to account for the dynamic domestic thermal characteristics, as discussed in Sec. 2.1). In the context of our MPC, we utilize a heuristic heating

¹²That 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.

schedule planning approach, that we propose, which enables the practical applicability of AdaHeat+ (in contrast to optimal planning, as further discussed in Sec. 3.3). Finally, AdaHeat+ incorporates a cost allocation mechanism to share the coalition gains that respects individual rationality. That said, AdaHeat+ is applicable both to coalitions and individual houses (as single-house coalitions).

In the following sections, we first describe our stochastic IER prediction approach in general (a specific instantiation is provided for our case study of shared wind turbine generators in Sec. 4.2.1). Then we detail the house modeling which is used in our MPC approach. Finally, we detail our MPC approach, along with the heuristic planning approach and the cost allocation mechanism.

3.1. Predicting the Power Output of the Shared IERs

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 $|I|$ (potentially dependent) stochastic processes R_i where $i \in \{1, 2, \dots, |I|\}$. Each one of these sub-processes considers the power output of each one of the shared IERs and is valued within the range $[0, r_i^{\max}]$, where r_i^{\max} is the maximum power output that can be achieved, at any time, by the IER indexed by i . Now, letting t stand for time, R_C will be 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, we can define $F_C^{-1} : [0, 1] \leftarrow [0, r_C^{\max}]$ as the quantile function corresponding to CDF F_C , where r_C^{\max} is the maximum power output that can be achieved in total, and at any time, by the shared IERs.

In this context, a variety of stochastic modeling approaches, including Kalman filters [Louka et al. 2008] and Gaussian processes [Jiang et al. 2010], can be used to model the power output of each one of the shared IERs (or directly the cumulative IER power output) in order to provide probabilistic estimates of the cumulative future power output, henceforth denoted R'_C [Zamo et al. 2014] (as discussed in Sec. 2.2). In Sec. 4.2.1 we provide an instantiation of IER stochastic prediction based on adaptive GP modeling for our case study of shared wind turbine generators.

3.2. Modeling Individual Houses

A detailed model of the system under control is crucial in the context of MPC. Following the component-based approach of AdaHeat, we modeled our system by deriving, for each house in the coalition: (i) a thermal comfort model of the occupants (incorporating an appropriate occupancy prediction algorithm), (ii) a thermal model of the house, and (iii) a heating cost model. We now proceed to detail each of these components/models.

3.2.1. Thermal Comfort Model. Optimal thermal comfort has been defined as “the condition of the mind in which satisfaction is expressed with the thermal environment”.¹³ In accordance with this definition, thermal comfort is affected by several interacting and less tangible factors such as differences in mood or activity, and environmental conditions [Djongyang et al. 2010]. For simplicity, we follow the (piecewise linear) intuitive AdaHeat modeling, directly extending it to account for the total comfort of the coalition. In particular, we assume comfort depends only on the inside air temperature, T^{IN} ; and any discomfort experienced, when the house is occupied, to be the absolute deviation of T^{IN} from the user-provided set-point temperature, T^{SP} . More formally, assuming all variables remain constant during interval τ of length δ , the expected thermal discomfort of C is:

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

where $O_{i,t}$ and $T_{i,t}^{IN}$ is the occupancy probability and inside temperature at time t (the starting time of interval τ) for each house $i \in C$, respectively (assumed constant during interval τ as noted above).¹⁴

¹³According to the ANSI/ASHRAE Standard 55, “Thermal Environment Conditions for Human Occupancy”, 2013.

¹⁴Note that T_i^{SP} is also defined according to house i , as different houses might have different set-point preferences.

That said, more advanced comfort modeling (e.g., [Auffenberg et al. 2015; Langevin et al. 2013]) can be directly incorporated in AdaHeat+ by simply adjusting the above equation accordingly.¹⁵ In addition, different modeling approaches can be used for each household to maximally reflect its preferences. Nevertheless, this particular choice does not generally affect the results of this work.

Now, to derive occupancy predictions, we employ the PreHeat approach which provides probabilistic estimates based on an average of K past similar days (selected based on a Hamming distance similarity estimation) [Scott et al. 2011]. We chose this scheduled-based approach (Sec. 2.1) since it has low instrumentation requirements, and has demonstrated high efficiency (median predictive accuracies of $\sim 80\%$) in comparison to the state-of-the-art [Kleiminger et al. 2014]. Moreover, it has been shown to be adequate in the context of an advanced DHASs [Panagopoulos et al. 2015].

3.2.2. Thermal Model. As discussed in Sec. 2.1, the thermal characteristics of houses are highly dynamic (due to the occupants' activity) and adaptive gray-box thermal modeling has been shown to effectively account for this dynamism (as discussed in Sec. 2.1). As such, in this work we use an adaptive gray-box thermal modeling approach. In particular, we employ a standard building thermal modeling formulation, 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 , is:

$$T_{i,t+1}^{IN} = T_{i,t}^{IN} + \left(u_{i,t} r_i^{in} - r_i^{out} (T_{i,t}^{IN} - T_{i,t}^{OUT}) + c_i \right) \delta \quad (2)$$

where, r_i^{in} stands for the heat provided by the heater; r_i^{out} for the leakage rate; $u_{i,t}$ for the heating action; 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 can be estimated with a range of respective techniques (such as Kalman filters, and appropriate least squares fitting [Soderstrom and Stoica 1989])—in Sec. 4.2.2 we provide a specific instantiation of our thermal modeling for our case study houses based on moving training window least-squares estimation [Soderstrom and Stoica 1989]. We note that it is possible to use a more complex formulation (and/or potentially a different formulation for each house) within our general framework [Panagopoulos et al. 2015].¹⁵ Nevertheless, this choice does not generally affect the results of this work as long as adequate thermal modeling is considered.

3.2.3. Heating Cost Model. In this work we extend the AdaHeat heating cost modeling to account for IER power availability along with import and export tariffs. Specifically, the amount of energy consumed by space heaters is the energy provided to the space over the efficiency of the system. As such, the aggregate input power consumption of coalition C , at any time instance t , is:

$$\text{Cons}_C(\cdot_t) = \sum_{i=1}^{|C|} \text{Cons}_i(\cdot_t), \text{ where } \text{Cons}_i(\cdot_t) = \frac{\text{Pwr}_i(u_{i,t}, T_{i,t}^{IN})}{C_i^{Eff}} \quad (3)$$

where $\text{Pwr}_i(u_{i,t}, T_{i,t}^{IN})$ stands for the output power of the space heating system according to the heating action and the building thermal state and C_i^{Eff} stands for the heating system efficiency (for each house $i \in C$).¹⁶ Given the export and import energy tariff difference, i.e., $P^{Buy} > P^{Sell}$, and assuming all variables remain constant to their value at time t (the starting time of interval τ) during interval τ of length δ , the expected consumption cost of coalition C (given the predicted cumulative IER power output $R'_{C,t}$) is:

$$\mathbb{E} [\text{Cost}_C(\cdot_\tau)] = \mathbb{E} \left[\max \left(0, \text{Cons}_C(\cdot_t) - R'_{C,t} \right) P_t^{Buy} + \min \left(0, \text{Cons}_C(\cdot_t) - R'_{C,t} \right) P_t^{Sell} \right] \delta \quad (4)$$

¹⁵In 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. 2010]), the dynamic programming planning approach of AdaHeat+ (further discussed in Sec. 3.3) does not generally raise constraints on the form of the components' modeling.

¹⁶ C_i^{Eff} can consider a simple scalar or be 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 [Chua et al. 2010].

A closed form of Eq. 4 is provided in Appendix A.1 in accordance with our GP-based IER modeling.

3.3. Controlling the Heating Systems of the Coalition

We now describe our control approach for collective AEC that utilizes the above components. As outlined in Sec. 2.1, we employ an adaptive MPC approach, 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 $O_{i,t}$, $R'_{C,t}$, $T_{i,t}^{OUT}$ (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.

In more detail, 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 , it corresponds to a set of intervals, noted H , where $|H| = \Delta/\delta$. During each interval, τ , all environmental conditions are assumed constant. Now, balancing heating cost and thermal discomfort is a two-objective optimization problem, and a weighted sum approach can be used to reduce the complexity and provide Pareto optimal solutions (although, potentially, not all) [Marler and Arora 2004]. As such, the optimization objective is to find the action matrix, \mathbf{U} , 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 :

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

where \mathbf{U}_i stands for the set of feasible heating control actions for each house $i \in C$, 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.

Now, solving the above objective leads to the optimal plan for AEC and corresponds to a shortest path finding problem in the emerging directed acyclic graph, G . In particular, each node of G considers a distinct tuple that contains all the necessary information (i.e., the current thermal state of the coalition and the interval id)¹⁷ to predict the next thermal state of the coalition according to a coalition heating action combination. Furthermore, each of the graph edges considers a tuple that contains the initial and successor node and a weigh corresponding to the expected unifying cost of the coalition during the particular interval. As such, the optimal plan for AEC corresponds to a shortest path finding problem in the 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 $O(|V_G| + |E_G|)$ time, with respect to the number of edges, $|E_G|$, and vertices, $|V_G|$, as proposed in [Panagopoulos et al. 2015]. 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 branch 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., $\prod_{i=1}^{|C|} |\mathbf{U}_i|$. 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. This fact leads to a complex and non-linear outcome of the λ population that disables their adjustment through an adaptive boolean feedback procedure (as achieved in [Panagopoulos et al. 2015]). This, in turn, undermines the applicability of optimal planning in realistic settings.

Given the above, we propose a heuristic planning approach for collective AEC that enhances the practical applicability of AdaHeat+. In particular, we divide equally all available IER energy among

¹⁷We note here that any exogenous information variables (such as the outside temperature, the occupancy and the future IER power output estimates) can be inferred by the time interval alone, and, hence, need not to be included in the respective tuple.

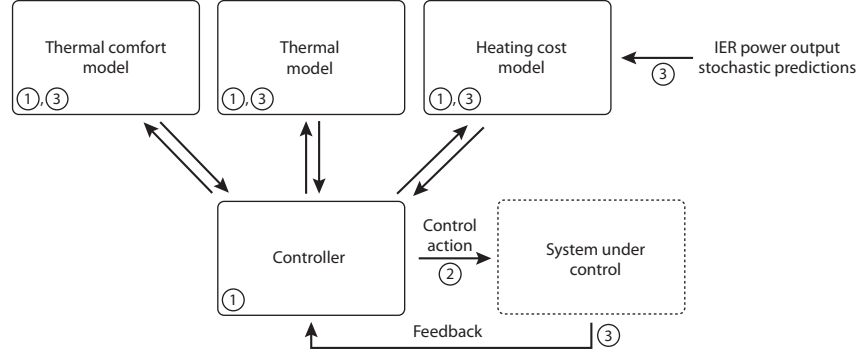


Fig. 1. AdaHeat+ heuristic approach flow diagram for each house. Our heuristic approach enables us to optimize the heating control process of the coalition in a distributed manner. As such, at each house: ① the controller plans a house-specific heating schedule over a fixed horizon into the future utilizing the specific house models and respective predictions—importantly, the controller utilizes stochastic predictions of the house-allocated share of the cumulative IER power output; then ② the first action of the heating schedule is executed; and ③ respective feedback is acquired and the system models and predictions are updated. Finally, the procedure is repeated, shifting the planning horizon into the future.

the members and plan their heating schedule, given the respective IER predictions, independently. More formally, the planning objective for each house $i \in C$ is:

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

where $\mathbb{E} [\text{Cost}_i^*(\cdot_\tau)]$, essentially, corresponds to Eq. 4 for single house coalitions, adjusted to consider an even share of the total IER energy availability. In particular:

$$\mathbb{E} [\text{Cost}_i^*(\cdot_\tau)] = \mathbb{E} \left[\max \left(0, \text{Cons}_i(\cdot_\tau) - R'_{C,i}/C \right) P^{\text{Buy}} + \min \left(0, \text{Cons}_{i,i}(\cdot_\tau) - R'_{C,i}/C \right) P^{\text{Sell}} \right] \delta \quad (5)$$

This optimization can be solved, independently for each house, as a shortest path finding problem (similar to optimal planning, above). As such, the complexity of our heuristic approach scales in a linear and parallelizable manner with the coalition size, that is $O(|C|)$, enabling it to be executed in separate units in each member (as further illustrated in Figure 1). Now, although an even share of the total IER energy availability is allocated to each house, some might consume less than originally allocated to them, while others might require more. Nevertheless, in the context of our collective AEC, planning 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 near-optimal) in such a genuine domestic scenario (further discussed in 4.4).¹⁸

Apart from its efficiency (demonstrated in Sec. 4.4) 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

¹⁸Considering coalitions where industrial or commercial entities are also participating (and, hence, the energy requirement differences are expected to be significant among the members) is imminent future work (as further discussed in Sec. 5).

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 (as achieved in [Panagopoulos et al. 2015], further illustrated in Sec. 4.4).

3.4. Allocating Gains to Coalition Members

Our work incorporates a cost allocation mechanism to share the realized coalition gains that respects individual rationality [Gibbons 1992], motivating, as such, the households to join a coalition. Appropriately, the outside option that we compare to is independent AEC, where each household optimizes its heating control process (considering both export tariffs and its own IER generation capacity) independently.¹⁹ Now, in order to identify the gains of collective AEC compared to independent AEC (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 AEC leads to the same discomfort for the house as collective AEC. Subsequently, the difference between the corresponding estimated heating cost of independent AEC (for all members, in total) and the cost of collective AEC 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 AEC, λ is in a monotonous relationship with discomfort (also discussed in Sec. 4.4). Hence, the specific λ 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.²⁰

Given this initial identification of the gains of collective AEC compared to independent AEC, we proceed with the respective allocation. Now, any allocation of 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 $i \in C$, over the day (as realized following collective AEC). More formally, we use the ratio:

$$\frac{\left(\int_{t=0}^{t=D} R_{i,t} dt + \theta_0 \right) / \left(\int_{t=0}^{t=D} \text{Cons}_i(\cdot_t) dt + \theta_1 \right)}{\sum_{i=1}^{|C|} \left[\left(\int_{t=0}^{t=D} R_{i,t} dt + \theta_0 \right) / \left(\int_{t=0}^{t=D} \text{Cons}_i(\cdot_t) dt + \theta_1 \right) \right]} \quad (6)$$

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),²¹ 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 satisfies individual rationality and allocation efficiency.

¹⁹Independent AEC considers AdaHeat+ for single-house coalitions and, hence, it is proposed for the first time in this work.

²⁰In 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 AEC reduces to independent AEC, with the only difference being that the members share their energy before buying from the grid. As such, the cost of independent AEC can be projected based on the realized consumption and IER production of each member.

²¹This 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 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.




Fig. 2. Instance of our GP-based hybrid wind turbine generator power output prediction approach.

4. EVALUATION

Here, we provide a thorough evaluation of AdaHeat+. We first describe the evaluation case study and how we collected the necessary data. Then, we describe the respective instantiation of our approach. Finally, we detail the experimental setup and the evaluation results.

4.1. Case Study and Data Collection

For our case study, we modeled 30 houses in Mablethorpe, Lincolnshire, UK, 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 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,²² for the location of our interest (i.e., Mablethorpe, UK).

4.2. Instantiating our Approach

In the following paragraphs, we detail the instantiation of AdaHeat+ for our case study: First, we detail our IER stochastic prediction instantiation for our case of wind turbine generators; and, then, we detail the instantiation of the individual member modeling and control.

4.2.1. Stochastic Wind Turbine Generator Power Output Prediction. 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 (Sec. 4.1) and the RENES²³ simulation platform [Panagopoulos et al. 2012] 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

²²www.uk.weather.com and www.wunderground.com

²³www.intelligence.tuc.gr/renes

to lie in the same region (and have the same WTG^{nom}), we use the same ground truth and derived NWP-based predictions, assuming the same sigmoid parameters (for simplicity) for all IERs.

Now, as discussed in Sec. 2.1, 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 through a simple and effective formulation, that we propose, within the context of adaptive GP regression. 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 $(x_1, y_1), \dots, (x_n, y_n)$ and NWP-based prediction points $(x_1, g_1), \dots, (x_n, g_n), \dots, (x_m, g_m)$ (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 $(z_1, x_1), \dots, (z_n, x_n)$. 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., $K_{\text{SE}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp((\mathbf{x} - \mathbf{x}')/2l^2)$, 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 $[0, r_C^{\text{max}}]$. 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 Fig. 2, 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




Fig. 3. Instantiation evaluation results.

estimates are under or over confident, respectively. As can be seen in Fig. 3(a), 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 Sec. 4.4 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; Yan et al. 2014].

4.2.2. Individual Members Modeling and Control. Regarding thermal comfort modeling, we set the set-point temperature, T_i^{SP} , of each house to the actual one as captured in the MyJoulo dataset for all 30 houses (Sec. 4.1). Moreover, for the occupancy schedule we utilized the one-month’s worth of data of the PreHeat original deployment (Sec. 4.1). 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.

Regarding thermal modeling, we trained our model using a moving-training-window least-squares fitting on the 30 houses of the MyJoulo dataset (Sec. 4.1). Specifically, our thermal modeling approach (Eq. 2), can be easily transformed in (least squares form) as $(T_{i,t+1}^{IN} - T_{i,t}^{IN}) / \delta = u_{i,t} r_i^{in} + r_i^{out} (T_i^{OUT} - T_{i,\tau}^{IN}) + c_i$, where, $(T_{i,t+1}^{IN} - T_{i,t}^{IN}) / \delta$ is the dependent variable/quantity, $u_{i,t}$ and $(T_i^{OUT} - T_{i,\tau}^{IN})$ 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 Fig. 3(b). It can be seen that this simple approach provides adequate predictive accuracy for our MPC approach which is in alignment to the literature [Panagopoulos et al. 2015]. 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.

Regarding heating cost modeling, we assumed for simplicity that the heating systems do not consider heating pump technology. As such, we set $C_i^{Eff} = 1$ and $\text{Pwr}_i(u_i, T_{i,t}^{IN}) = u_i \text{HS}^{\text{nom}}$ (i.e., $\text{Pwr}_i(0, T_{i,t}^{IN}) = 0$, and $\text{Pwr}_i(1, T_{i,t}^{IN}) = \text{HS}^{\text{nom}} \forall i \in C$.²⁴ Furthermore, following the 2015 UK market rates, we set $P^{Buy} = 11.25$ p/kWh and $P^{Sell} = 4.5$ p/kWh (p is British penny) which correspond to $P^{Buy} / P^{Sell} = 2.5$. Now, given $R'_C \sim N(\mu, \sigma)$, $R_C \in [0, r_C^{\max}]$, and the straightforward conditions that $r_C^{\max} > 0$, and $\text{Cons}(\cdot) \geq 0$, Eq. 4 can be transformed in closed form (Appendix A.1). Finally, regarding the controller, the slow nature of building thermal dynamics [Široký et al. 2011] dictates a relatively long planning horizon, Δ , and allows a long planning interval, δ . As such, we set $\Delta = 3.5$ hours and $\delta = 10$ minutes after experimenting with various populations [Camacho and Alba 2013].²⁵

4.3. Experimental Setup

In our evaluation procedure, we compare our heuristic collective AEC approach, in a contemporary market reality with flat import and export tariffs, against two benchmark approaches: (i) independent AEC, 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, using AdaHeat. As such, we evaluate the benefits

²⁴We note that setting $C_i^{Eff} = 1$ ensures that we do not favor our approach via modeling questionably inefficient systems.

²⁵Notably, this choice does not set any particular restrictions on the billing interval of a house (or a coalition), since smart meters generally enable the data acquisition interval within a home network to be set independently (and also provide considerable flexibility in this respect) [Toledo 2013].

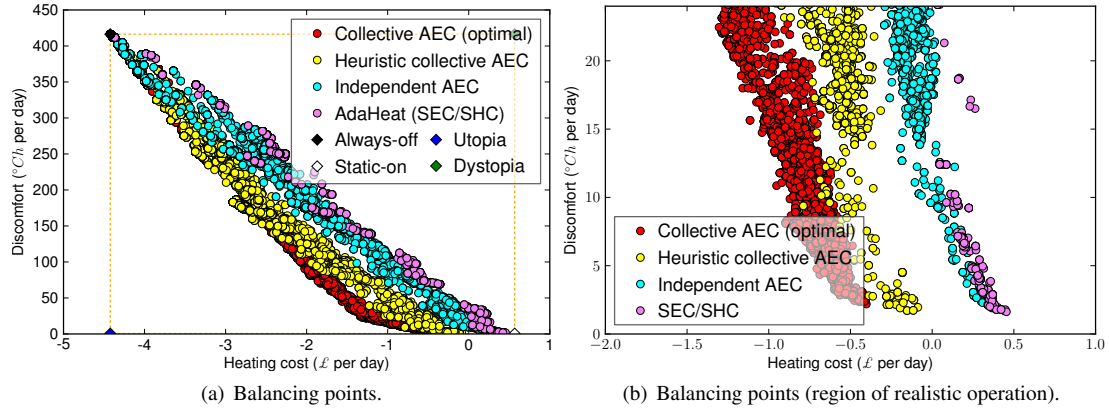


Fig. 4. Initial evaluation results (two-house coalition)

of collective AEC, compared to independent AEC, as well as the benefits of independent AEC compared to (independent) simple economic/heating control. We note here that although independent AEC is used as a benchmark against heuristic collective AEC, it considers our approach for single house coalitions. Moreover, in the case of flat import tariffs, simple economic control, that merely considers import tariffs (see Sec. 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 Feb. 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 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 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 AEC considering optimal planning, for the case of a two-house coalition (where the respective optimization is feasible, Sec. 3.4). 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 NWP-based predictions. Note here that for single house coalitions, heuristic collective AEC, optimal collective AEC, and independent AEC become the same. Hence, we evaluate the benefits of our hybrid approach without being affected by any heuristic planning losses. Finally, we also note here that since AdaHeat+ considers an extension of AdaHeat to account for AEC (and to exploit the respective coalition potential), the general improvements of AdaHeat with respect to simple economic and simple heating control (compared to the state-of-the-art) illustrated in Panagopoulos et al. [2015] are retained.

4.4. Evaluation Results

Fig. 4(a) illustrates our evaluation results for a two-house coalition (where optimal planning can feasibly be used, as discussed in Sec. 3.3), considering the collective discomfort and cost of the coalition per day (negative “cost” indicates profit). Given the fact that extreme discomfort preferences are not expected in real settings, Fig. 4(b) focuses on the region where the discomfort preferences are arguably more likely. We note here, that since this is a two-objective optimization, points closer to the origin indicate higher Pareto efficiency. Within this context, Fig. 4 illustrates that our heuristic collective AEC approach has considerably higher Pareto efficiency when compared to independent AEC and even higher when compared to AdaHeat, i.e., simple economic/heating control

Table II. Cumulative evaluation results.

	(a) Heating cost-efficiency of each approach				(b) Hybrid contribution	
	Collective AEC (optimal)	Heuristic collective AEC	Independent AEC	AdaHeat (SEC/SHC)	Hybrid predictions	NWP-based predictions
Heating cost-efficiency	1.48	1.46	1.17	1.03	1.26	1.10
Standard mean error	.002	.011	.010	.006	.081	.060

(denoted in all figures as SEC/SHC).²⁶ Furthermore, independent AEC 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 AEC 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 estimated the heating cost-efficiency of each approach. To do so, we evaluate two simple static timer heating policies; namely *static-on* and *always-off*, illustrated in Fig. 4(a) (for concreteness). In essence, *always-off* is the heating policy where there is no heating at all, while *static-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. Intuitively, *always-off* considers a possible worst-case scenario in terms of discomfort, Disc^{\max} , while it considers the best-case scenario in terms of heating cost, Cost^{\min} . In contrast, *static-on* considers a possible worst-case scenario in terms of cost, Cost^{\max} , and the best-case scenario in terms of discomfort, Disc^{\min} . Now, we utilized 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 *static-on* and *always-off* policies create a 4-point imaginary window that frames all solution points of our approaches (as seen in Fig. 4(a)). Given this window, we normalized 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 calculated 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 $A := \{P_1, P_2, \dots, P_{|A|}\}$, where the coordinates of any given point P_x is a tuple $(\text{Cost}_{P_x}, \text{Disc}_{P_x})$ considering the respective cost and discomfort values, the heating cost-efficiency of the approach is calculated as $\frac{1}{|A|} \sum_{i=1}^{|A|} \frac{1 - (\text{Disc}_{P_i} - \text{Disc}^{\min}) / (\text{Disc}^{\max} - \text{Disc}^{\min})}{(\text{Cost}_{P_i} - \text{Cost}^{\min}) / (\text{Cost}^{\max} - \text{Cost}^{\min})}$. Table II(a) reports the heating cost-efficiency of each approach. It can be seen that the results are consistent with the observations in Fig. 4(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.

Now, the above results indicate a significant improvement in the heating cost-efficiency of collective AEC compared to independent AEC (and even more when compared to simple economic/heating control) which can lead to considerable monetary gains for the coalition members. Fig. 6 illustrates the monetary gains of using collective AEC, for each one of the houses in the coalition (after the cost allocation), compared to: (i) independent AEC 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 AEC compared to independent (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 AEC fluctuate considerably along with the heating cost (since, for instance, there cannot be any gains if heating is

²⁶Fig. 4 also shows how our approach is able to capture a wide and smooth range of discomfort-cost balancing points (as AdaHeat), enabling, as such, a wide range of user preferences to be met [Panagopoulos et al. 2015].




Fig. 5. Two-house coalition energy exchange.




Fig. 6. Two-house coalition monetary gains.

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). Fig. 5 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 we can see, the energy consumption flattens significantly (towards 0) even for a small coalition of only 2 houses, compared to independent AEC. Now, as discussed in Sec. 4.2.1 here we also identify the benefits of incorporating our hybrid wind turbine generator predictive approach to AdaHeat+. We do so by evaluating AdaHeat+ for a single house coalition both with NWP-based and our hybrid approach (Table II(b)). Not surprisingly, we see that our hybrid approach considerably increases the cost-efficiency of AdaHeat+, as more accurate predictions are considered.

As discussed in Sec. 3.3, AdaHeat+ also enables the coalition members to balance worst-case heating cost and thermal discomfort through a simple boolean feedback procedure. To demonstrate this, Fig. 7 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 cost and discomfort are in a monotonous relationship with λ (irrespective 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 Sec. 3.3, the realized heating cost depends also on the λ population of the other coalition members. To demonstrate this, Fig. 8 illustrates the realized cost and discomfort balancing





Fig. 7. Balancing worst-case cost and discomfort.

Fig. 8. Fixed λ evaluation results.




Fig. 9. Heating cost-efficiency vs coalition size.




Fig. 10. Collective compared to independent AEC.

of each house given a fixed $\lambda = 90$ and for various λ populations for the other house.²⁷ 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 an informed decision. Evaluating the user behavior with respect to this balancing requires a real world trial (that considers future work—see Sec. 5).

Regarding scalability, we evaluate how our approach scales with the size of the coalition. Fig. 9 illustrates how the heating cost-efficiency scales with the size of the coalition up to 30 houses (the bars consider standard mean error). As we can see the cost-efficiency of simple economic/heating control and independent AEC remain relatively constant with the size of the coalition, as there is no energy exchange within the coalition. In contrast, the collective AEC 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 AEC over independent AEC in Fig. 10 illustrates that the efficiency improvement flattens at around 27 houses, following a logarithmic trend. As can be seen in Fig. 11, this is the size 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 Fig. 5). Fig. 10 also illustrates the mean monetary gains of the houses, compared to independent AEC (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 (Sec. 3.4). 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 AEC 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 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 (Sec. 4.1). Finally, in order to further identify the financial gains of AdaHeat+, we evaluate a specific instance of our approach given fixed discomfort-cost preferences. Fig. 12 illustrates how the mean cost of the houses scales with the size of the coalition given a fixed discomfort preference of $\sim 12^\circ\text{Ch}$ for all houses. According to the houses' occupancy schedule this discomfort could corresponds to a $1\text{-}2^\circ\text{C}$ 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 AEC leads to considerable savings (as a fraction of the daily heating cost) in this realistic scenario (the

²⁷For the effective range of λ for each house consult Fig. 7.





Fig. 11. 27-house coalition energy exchange.

Fig. 12. Cost with fixed mean discomfort at $\sim 12^\circ\text{C}$.

bars correspond to standard mean error capturing the diverse share of gains allocated to each house). As a final note, our evaluation results suggest that, due to their shifting potential (as discussed in Sec. 1), heating loads can effectively support collective AEC even if similar occupancy schedules and/or household preferences are considered, 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.

5. CONCLUSIONS

This work is the first to show how DHASs can be enhanced with AEC to increase the economic efficiency of houses with electricity-based space heating systems and grid-connected IERs. It also illustrates the potential of domestic coalitions in this context. Specifically, we propose AdaHeat+, an advanced DHAS that extends a state-of-the-art DHAS (retaining its practical advantages) to account for AEC and to exploit the respective coalition potential. Importantly, AdaHeat+: (i) is applicable to both coalitions and individual houses, (ii) is able to handle different preferences in balancing cost and discomfort of individual households, (iii) incorporates a computationally practical allocation mechanism to share the collective gains that respects individual rationality and allocation efficiency, and (iv) has a complexity that scales linearly and in a parallelizable manner with the size of the coalition (due to an effective heuristic planning approach for collective AEC that we propose). Our approach has efficiency and low computational complexity that allows it to be applicable in real settings with limited computational resources, minimum instrumentation, and operating time constraints. We demonstrate its effectiveness through real data evaluation and show that collective AEC can improve heating cost-efficiency by up to $\sim 60\%$ (leading to considerable monetary gains), compared to independent AEC that we also propose, and even more when compared to simple economic/heating control. In the context of our AEC schema we also propose a simple and effective formulation for hybrid wind turbine generator stochastic prediction based on adaptive Gaussian process modeling to improve the short-term accuracy of NWP-based predictions.

Imminent future work will involve a real world trial. In such a scenario we also intend to evaluate broader and less homogeneous coalitions where also commercial and/or industrial members are allowed to participate. By doing so we aim to identify the potential gains of AdaHeat+ for various heating technologies (and particularly for heat-pumps) and evaluate different IER energy allocation schemes within our heuristic approach that can potentially outperform our even allocation approach in such settings (as discussed in Sec. 3.3). There, we also intend to collect valuable feedback from the users and evaluate their behavior with respect to balancing worst-case heating cost and thermal discomfort, and also explore the concept of optimal decision making regarding profitable coalition formation and fairness. Considering variable export and import tariffs (and respective stochastic price predictions in the context of AdaHeat+) is an additional future work direction.

REFERENCES

- Muddasser Alam, Athanasios A. Panagopoulos, Alex Rogers, Nicholas R. Jennings, and James Scott. 2014. Applying Extended Kalman Filters to Adaptive Thermal Modelling in Homes: Poster Abstract. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*. ACM, 214–215.
- Muddasser Alam, Sarvapali D Ramchurn, and Alex Rogers. 2013. 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*. International Foundation for Autonomous Agents and Multiagent Systems, 731–738.
- Frederik Auffenberg, Sebastian Stein, and Alex Rogers. 2015. A Personalised Thermal Comfort Model Using a Bayesian Network. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*. AAAI Press, 2547–2553.
- Peder Bacher and Henrik Madsen. 2011. Identifying Suitable Models For The Heat Dynamics of Buildings. *Energy and Buildings* 43, 7 (2011), 1511–1522.
- Godfrey Boyle (Ed.). 2012. *Renewable Energy: Power for a Sustainable Future (3rd Ed.)*. Oxford University Press and Open University.
- Barbara G. Brown, Richard W. Katz, and Allan H. Murphy. 1984. Time Series Models to Simulate and Forecast Wind Speed and Wind Power. *Journal of Climate and Applied Meteorology* 23, 8 (1984), 1184–1195.
- Phil Bulman. 2015. Tesla's Powerwall Battery Production Requires Super-charged Supply Chain. *Renewable Energy Focus* 16, 56 (2015), 126–127.
- Tony Burton, David Sharpe, Nick Jenkins, and Ervin Bossanyi. 2011. *Wind Energy Handbook*. John Wiley & Sons, Ltd.
- Duncan S. Callaway. 2009. 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 (2009), 1389–1400.
- Eduardo F. Camacho and Carlos B. Alba. 2013. *Model Predictive Control*. Springer.
- Georgios Chalkiadakis, Valentin Robu, Alex Kota, Ramachandra and Rogers, and Nicholas R Jennings. 2011. 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)*, Vol. 2. 787–794.
- Niya Chen, Zheng Qian, Ian T. Nabney, and Xiaofeng Meng. 2014. Wind Power Forecasts Using Gaussian Processes and Numerical Weather Prediction. *IEEE Transactions on Power Systems* 29, 2 (2014), 656–665.
- Kian J. Chua, Siaw K. Chou, and Wenming Yang. 2010. Advances In Heat Pump Systems: A Review. *Applied Energy* 87, 12 (2010), 3611–3624.
- Thomas H. Cormen, Clifford Stein, Ronald L. Rivest, and Charles E. Leiserson. 2001. *Introduction to Algorithms* (2nd ed.). McGraw-Hill Higher Education.
- Frits de Nijs, Matthijs T. J. Spaan, and Mathijs M. de Weerd. 2015. Best-Response Planning of Thermostatically Controlled Loads under Power Constraints. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*. 615–621.
- Noël Djongyang, René Tchinda, and Donatien Njomo. 2010. Thermal Comfort: A Review Paper. *Renewable and Sustainable Energy Reviews* 14, 9 (2010), 2626–2640.
- Junqiao H. Dudley and Mary A. Piette. 2008. 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 (2008), 1–4.
- Donald R. Farris and James L. Melsa. 1978. Energy Savings For A Solar Heated And Cooled Building Through Adaptive Optimal Control. In *IEEE Conference on Decision and Control Including The 17th Symposium on Adaptive Processes*. 206–213.
- Peter X. Gao and Srinivasan Keshav. 2013. Optimal Personal Comfort Management Using SPOT+. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM, 1–8.
- Robert Gibbons. 1992. *A Primer In Game Theory*. Harvester Wheatsheaf.
- Rasmus Halvgaard, Niels K. Poulsen, Henrik Madsen, and John B. Jørgensen. 2012. Economic Model Predictive Control For Building Climate Control In A Smart Grid. In *IEEE PES Innovative Smart Grid Technologies*. 1–6.
- He Hao, Borhan M. Sanandaji, Kameshwar Poolla, and Tyrone L. Vincent. 2013. 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*. 551–558.
- David Hawkey. 2015. *Sustainable Urban Energy Policy*. Taylor & Francis.
- Hao Huang, Lei Chen, Morteza Mohammadzadeh, Eric Hu, and Minlei Chen. 2013. Multi-zone Temperature Prediction In A Commercial Building Using Artificial Neural Network Model. In *The 10th IEEE International Conference on Control and Automation*. 1896–1901.
- Mark Z. Jacobson, Mark A. Delucchi, Guillaume Bazouin, Zack A. F. Bauer, Christa C. Heavey, Emma Fisher, Sean B. Morris, Diniana J. Y. Piekutowski, Taylor A. Vencill, and Tim W. Yeskoo. 2015. 100% Clean And Renewable Wind, Water, And Sunlight (WWS) All-Sector Energy Roadmaps For The 50 United States. *Energy & Environmental Science* 8 (2015), 2093–2117.

- Steffen Jenner, Felix Groba, and Joe Indvik. 2013. Assessing The Strength And Effectiveness of Renewable Electricity Feed-in Tariffs In European Union Countries. *Energy Policy* 52 (2013), 385–401. Special Section: Transition Pathways to a Low Carbon Economy.
- Xiaoqian Jiang, Bing Dong, Le Xie, and Latanya Sweeney. 2010. Adaptive Gaussian Process for Short-Term Wind Speed Forecasting. In *Proceedings of the 19th European Conference on Artificial Intelligence*. IOS Press, 661–666.
- Lalarukh Kamal and Yasmin Zahra Jafri. 1997. Time Series Models To Simulate And Forecast Hourly Averaged Wind Speed In Quetta, Pakistan. *Solar Energy* 61, 1 (1997), 23–32.
- Wilhelm Kleiminger, Friedemann Mattern, and Silvia Santini. 2014. Predicting Household Occupancy For Smart Heating Control: A Comparative Performance Analysis Of State-Of-The-Art Approaches. *Energy and Buildings* 85 (2014), 493–505.
- Tamar Krishnamurti, Daniel Schwartz, Alexander Davis, Baruch Fischhoff, Wändi B. de Bruin, Lester Lave, and Jack Wang. 2012. Preparing For Smart Grid Technologies: A Behavioral Decision Research Approach To Understanding Consumer Expectations About Smart Meters. *Energy Policy* 41 (2012), 790–797.
- Michaël Kummert, Philippe André, and Jacques Nicolas. 2001. Optimal Heating Control In A Passive Solar Commercial Building. *Solar Energy* 69, Supplement 6 (2001), 103–116.
- Jared Langevin, Jin Wen, and Patrick L. Gurian. 2013. Modeling Thermal Comfort Holistically: Bayesian Estimation Of Thermal Sensation, Acceptability, And Preference Distributions For Office Building Occupants. *Building and Environment* 69 (2013), 206–226.
- Gong Li and Jing Shi. 2010. On Comparing Three Artificial Neural Networks For Wind Speed Forecasting. *Applied Energy* 87, 7 (2010), 2313–2320.
- Xiwan Li and Jin Wen. 2014. Review Of Building Energy Modeling For Control And Operation. *Renewable and Sustainable Energy Reviews* 37 (2014), 517–537.
- Marco Liserre, Thilo Sauter, and John Y. Hung. 2010. Future Energy Systems: Integrating Renewable Energy Sources into the Smart Power Grid Through Industrial Electronics. *IEEE Industrial Electronics Magazine* 4, 1 (2010), 18–37.
- Petroula Louka, George Galanis, Nils Siebert, George Kariniotakis, Petros Katsafados, Ioannis Pytharoulis, and George Kallos. 2008. Improvements In Wind Speed Forecasts For Wind Power Prediction Purposes Using Kalman Filtering. *Journal of Wind Engineering and Industrial Aerodynamics* 96, 12 (2008), 2348–2362.
- Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field, and Kamin Whitehouse. 2010. 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)*. ACM, 211–224.
- Antonio Luque and Steven Hegedus. 2011. *Handbook of Photovoltaic Science and Engineering*. Wiley.
- R. Timothy Marler and Jasbir S. Arora. 2004. Survey Of Multi-Objective Optimization Methods For Engineering. *Structural and multidisciplinary optimization* 26, 6 (2004), 369–395.
- Mohammad A. Mohandes, Talal O. Halawani, Shafiqur Rehman, and Ahmed A. Hussain. 2004. Support Vector Machines For Wind Speed Prediction. *Renewable Energy* 29, 6 (2004), 939–947.
- Michael C. Mozer, Lucky Vidmar, and Robert H. Doder. 1997. The Neurothermostat: Predictive Optimal Control of Residential Heating Systems. *Advances in Neural Information Processing Systems* (1997), 953–959.
- Astrid Niese, Sebastian Lehnhoff, Martin Tröschel, Mathias Uslar, Carsten Wissing, H Jürgen Appelrath, and Michael Sonnenschein. 2012. Market-Based Self-Organized Provision Of Active Power And Ancillary Services: An Agent-Based Approach For Smart Distribution Grids. In *Complexity in Engineering*. IEEE, 1–5.
- Frauke Oldewurtel, Andreas Ulbig, Alessandra Parisio, Göran Andersson, and Manfred Morari. 2010. Reducing Peak Electricity Demand In Building Climate Control Using Real-Time Pricing And Model Predictive Control. In *The 49th IEEE Conference on Decision and Control*. 1927–1932.
- Martin J. Osborne and Ariel Rubinstein. 1994. *A Course in Game Theory*. MIT Press.
- Rosalía Pacheco, Javier Ordóñez, and Germán Martínez. 2012. Energy Efficient Design of Building: A Review. *Renewable and Sustainable Energy Reviews* 16, 6 (2012), 3559–3573.
- Omid Palizban, Kimmo Kauhaniemi, and Josep M. Guerrero. 2014. Microgrids In Active Network Management Part I: Hierarchical Control, Energy Storage, Virtual Power Plants, And Market Participation. *Renewable and Sustainable Energy Reviews* 36 (2014), 428–439.
- Athanasios A. Panagopoulos, Muddasser Alam, Alex Rogers, and Nicholas R. Jennings. 2015. AdaHeat: A General Adaptive Intelligent Agent for Domestic Heating Control. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1295–1303.
- Athanasios A. Panagopoulos, Georgios Chalkiadakis, and Eftichios Koutroulis. 2012. Predicting the Power Output of Distributed Renewable Energy Resources within a Broad Geographical Region. In *The 20th European Conference on Artificial Intelligence*. 981–986.
- Marius Paulescu, Eugenia Paulescu, Paul Gravila, and Viorel Badescu. 2012. *Weather Modeling and Forecasting of PV Systems Operation*. Springer London.

- Therese Peffer, Marco Pritoni, Alan Meier, Cecilia Aragon, and Daniel Perry. 2011. How People Use Thermostats In Homes: A Review. *Building and Environment* 46, 12 (2011), 2529–2541.
- Samuel Přívara, Jiří Cigler, Zdeněk Váňa, Frauke Oldewurtel, Carina Sagerschnig, and Eva Žáčková. 2013. Building Modeling As A Crucial Part For Building Predictive Control. *Energy and Buildings* 56, 0 (2013), 8–22.
- Sarvapali Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nicholas R. Jennings. 2012. Putting the “Smarts” into the Smart Grid: A Grand Challenge for Artificial Intelligence. *Commun. ACM* (2012).
- Carl E. Rasmussen and Christopher K. I. Williams. 2005. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning Series)*. The MIT Press.
- Steven Reece, Siddhartha Ghosh, Alex Rogers, Stephen Roberts, and Nicholas R. Jennings. 2014. Efficient State-Space Inference Of Periodic Latent Force Models. *The Journal of Machine Learning Research* 15, 1 (2014), 2337–2397.
- Alex Rogers, Siddhartha Ghosh, Reuben Wilcock, and Nicholas R. Jennings. 2013. 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*. ACM, 1–8.
- Alex Rogers, Sasan Maleki, Siddhartha Ghosh, and Nicholas R. Jennings. 2011. Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming. In *The 2nd International Workshop on Agent Technology for Energy Systems*. 71–78.
- Antonio E. Ruano, Eduardo M. Crispim, Eusébio Z. E. Conceição, and Ma Manuela J. R. Lúcio. 2006. Prediction Of Building’s Temperature Using Neural Networks Models. *Energy and Buildings* 38, 6 (2006), 682–694.
- James Scott, Alice J. Bernheim Brush, John Krumm, Brian Meyers, Michael Hazas, Stephen Hodges, and Nicolas Villar. 2011. Preheat: Controlling Home Heating Using Occupancy Prediction. In *Proceedings of the 13th International Conference on Ubiquitous Computing*. ACM, 281–290.
- Mike Shann and Sven Seuken. 2014. Adaptive Home Heating Under Weather And Price Uncertainty Using GPS And MDPS. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 821–828.
- Michelle Shipworth, Steven K. Firth, Michael I. Gentry, Andrew J. Wright, David T. Shipworth, and Kevin J. Lomas. 2010. Central Heating Thermostat Settings And Timing: Building Demographics. *Building Research & Information* 38, 1 (2010), 50–69.
- Fereidoon P. Sioshansi. 2016. *Future of Utilities - Utilities of the Future: How Technological Innovations in Distributed Energy Resources Will Reshape the Electric Power Sector*. Elsevier Science.
- Jan Šíroký, Frauke Oldewurtel, Jiří Cigler, and Samuel Přívara. 2011. Experimental Analysis of Model Predictive Control For An Energy Efficient Building Heating System. *Applied Energy* 88, 9 (2011), 3079–3087.
- Torsten S. Soderstrom and Peter G. Stoica. 1989. *System Identification*. Prentice Hall.
- Saurabh S. Soman, Hamidreza Zareipour, Om Malik, and Paras Mandal. 2010. A Review Of Wind Power And Wind Speed Forecasting Methods With Different Time Horizons. In *North American Power Symposium, 2010*. 1–8.
- Fabio Toledo. 2013. *Smart Metering Handbook*. Pennwell Corporation.
- Jacopo Torriti, Mohamed G. Hassan, and Matthew Leach. 2010. Demand Response Experience In Europe: Policies, Programmes And Implementation. *Energy* 35, 4 (2010), 1575–1583.
- Daniel Urieli and Peter Stone. 2013. A Learning Agent For Heat-Pump Thermostat Control. In *Proceedings of The 2013 International Conference on Autonomous Agents And Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1093–1100.
- Shengwei Wang and Zhenjun Ma. 2008. Supervisory And Optimal Control Of Building HVAC Systems: A Review. *HVAC&R Research* 14, 1 (2008), 3–32.
- Juan Yan, Kang Li, Er-Wei Bai, Jing Deng, and Aoife M. Foley. 2016. Hybrid Probabilistic Wind Power Forecasting Using Temporally Local Gaussian Process. *IEEE Transactions on Sustainable Energy* 7, 1 (2016), 87–95.
- Juan Yan, Zhile Yang, Kang Li, and Yusheng Xue. 2014. A Variant Gaussian Process for Short-Term Wind Power Forecasting Based on TLBO. *Intelligent Computing in Smart Grid and Electrical Vehicles* (2014), 165.
- Ahmad Zahedi. 2010. Australian Renewable Energy Progress. *Renewable and Sustainable Energy Reviews* 14, 8 (2010), 2208–2213.
- Michael Zamo, Olivier Mestre, Philippe Arbogast, and Olivier Pannekoucke. 2014. 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 (2014), 804–816.

Received ; revised ; accepted

APPENDIX**A.1. Closed Form Expected Cost**

Here, we provide a closed form of Eq. 4, given $R' \sim N(\mu, \sigma)$:

$$\mathbb{E}[\text{Cost}(\cdot)] = \begin{cases} P^{Sell} \left(\frac{\sigma}{K} \sqrt{\frac{2}{\pi}} \left(e^{\frac{(\mu - r_C^{\max})^2}{-2\sigma^2}} - e^{\frac{-\mu^2}{2\sigma^2}} \right) - \mu \right) & \text{Cons}(\cdot)=0 \\ \frac{\text{Cons}(\cdot) - \mu}{K} \left(\text{erf} \left(\frac{\text{Cons}(\cdot) - \mu}{\sqrt{2}\sigma} \right) (P^{Buy} - P^{Sell}) + \text{erf} \left(\frac{\mu}{\sqrt{2}\sigma} \right) P^{Buy} - \text{erf} \left(\frac{\mu - r_C^{\max}}{\sqrt{2}\sigma} \right) P^{Sell} \right) & \text{Cons}(\cdot) < r_C^{\max} \\ + \frac{\sigma}{K} \sqrt{\frac{2}{\pi}} \left(P^{Buy} \left(e^{\frac{(\text{Cons}(\cdot) - \mu)^2}{-2\sigma^2}} - e^{\frac{-\mu^2}{2\sigma^2}} \right) + P^{Sell} \left(e^{\frac{(\mu - r_C^{\max})^2}{-2\sigma^2}} - e^{\frac{(\text{Cons}(\cdot) - \mu)^2}{-2\sigma^2}} \right) \right) & \wedge \text{Cons}(\cdot) > 0 \\ P^{Buy} \left(\frac{\sigma}{K} \sqrt{\frac{2}{\pi}} \left(e^{\frac{(\mu - r_C^{\max})^2}{-2\sigma^2}} - e^{\frac{-\mu^2}{2\sigma^2}} \right) + \text{Cons}(\cdot) - \mu \right) & \text{Cons}(\cdot) \geq r_C^{\max} \\ & \wedge \text{Cons}(\cdot) > 0 \end{cases} \quad (7)$$

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