Lawrence Berkeley National Laboratory

Title

Dealing with expected thermal discomfort

Permalink

https://escholarship.org/uc/item/05c575fc

Authors

Panagopoulos, Athanasios Aris

Katsigiannis, Michail

Pritoni, Marco

et al.

Publication Date

2023-12-13

Peer reviewed

Dealing with Expected Thermal Discomfort

Athanasios Aris Panagopoulos, University of California, Berkeley
(apanagopoulos@outlook.com)
Michail Katsigiannis, Technical University of Crete (mkatsigiannis@isc.tuc.gr)
Marco Pritoni, Lawrence Berkeley National Laboratory (marco.pritoni@gmail.com)
Gabe Fierro, University of California, Berkeley (gtfierro@cs.berkeley.edu)
Daniel Lengyel, University of California, Berkeley (daniel.lengyel@berkeley.edu)
Therese E Peffer, University of California, Berkeley (tpeffer@berkeley.edu)
Georgios Chalkiadakis, Technical University of Crete (gchalkiadakis@isc.tuc.gr)
David Culler, University of California, Berkeley (culler@berkeley.edu)

ABSTRACT

Optimizing the heating ventilation and air condition (HVAC) control process generally considers a trade-off between thermal discomfort and energy cost. Technically, this can be represented as a multi-objective optimization problem, where the occupants' thermal preferences and the cost of energy consumed are two conflicting objectives. When the occupancy schedule is known in advance this problem typically reduces to minimizing the energy cost while meeting the thermal comfort constraint set by a user, typically defined as a setpoint-temperature-based comfort band. However, this approach is inadequate in dynamic occupancy settings where the optimization process is driven by occupancy estimates that are inherently uncertain. In particular, ensuring that the room temperature lays within a typically narrow comfort band, even when there is small probability of occupancy, significantly increases HVAC energy usage. In this case, accepting some probability of discomfort, is essential to reduce energy cost. We investigate three optimization HVAC control algorithms that deal with dynamic occupancy and we provide a comparative analysis to show the advantages and disadvantages of each approach with respect to their efficiency, effectiveness, applicability and usability in different settings.

1. Introduction

Heating ventilation and air condition (HVAC) represents one of the biggest shares of residential and commercial buildings' energy consumption (EIA 2015). As such, increasing the HVAC efficiency can lead to considerable saving on the monthly energy bills and carbon emissions. Improving HVAC operation using optimization algorithms requires less cost and disruption compared to more intrusive energy efficiency improvements, such as replacing building insulation and installing new HVAC equipment. For this reason, optimizing the HVAC control process has been heralded as a key means for energy efficiency improvements in buildings (Dounis and Caraiscos 2009) towards an energy sustainable future.

Optimizing the HVAC control process generally considers a trade-off between thermal discomfort and energy cost. Technically, this can be represented as a multi-objective optimization problem, where the occupants' thermal preferences and the cost of energy are the two main and conflicting objectives. Other objectives include safety (e.g., min temperature allowed) or regulatory constraints. In fixed occupancy scenarios, such as in commercial buildings with a fixed working schedule, optimizing the HVAC control process typically involves meeting the thermal comfort constraint set by the user with the minimum energy cost. These requirements are typically

defined as a fixed setpoint temperature comfort band for particular time periods (e.g., the next 24 hours). However, this approach is inadequate in dynamic occupancy settings, such as in homes and many office buildings, where the optimization process is driven by occupancy estimates that are inherently uncertain. In particular, ensuring that the room temperature lays within a narrow comfort band, even when there is a small probability of occupancy, significantly increases HVAC energy usage. In this case, accepting some probability of discomfort, is essential to reduce energy cost. This is especially true in residential settings which are less likely to have multiple zone sensors. As such, defining the amount of discomfort acceptable according to the user preferences is an additional challenge in such settings.

A number of approaches for optimizing the HVAC control process have been proposed over time that deal with probabilistic occupancy estimates (Gao and Keshav 2013; Lu et al. 2010; Panagopoulos et al. 2017; Panagopoulos, A.A. 2016; Panagopoulos et al. 2015; Scott et al. 2011; Urieli and Stone 2013). One can distinguish two main lines of research: the thresholding approach, and the weighting sum approach. The former relies on thresholding the probabilistic occupancy estimates (i.e., defining a probability threshold above which the occupancy is considered certain) and optimizing the HVAC control process on the derived deterministic occupancy schedule. The latter utilizes the probabilistic occupancy schedules as such and relies on a weighting parameter to balance expected thermal discomfort and heating cost (Panagopoulos et al. 2015). Each one of these approaches comes with cons and pros with respect to their efficiency, applicability and usability. Nevertheless, these advantages and disadvantages have not been thoroughly understood and evaluated. Importantly, the critical question of how to meet the user preferences in balancing heating cost and thermal discomfort is usually disregarded in designing intelligent algorithms despite the extensive literature covering discussing this challenge (de Dear and Brager, 2002, Rutkowski 1997). A notable exception, Panagopoulos et al. 2015 is our point of departure and provides preliminary results in this direction.

In this paper we provide a qualitative and quantitative comparison of these approaches evaluating their efficiency, applicability and usability. Importantly, we also propose a new approach that relies on variable bounding and show that it is able to capture optimal solutions with minimum user input. Our study shows that the weighted sum formulation is a clear winner when considering the range of different user preferences that is able to capture. However, all approaches capture optimal solutions with minimum user input. As such, choosing one approach over the other depends on the application requirements with respect to user input. In this work we also show that in stochastic occupancy settings the setpoint temperature set by the user should consider the origin of the discomfort metric, i.e., the region where the occupant feels absolute thermal comfort, and not a parameter to balance discomfort and cost as this balancing needs to happen while respecting the expected occupancy probabilities.

The rest of the paper is structured as follows: In Section 2 we provide background material. Then, in Section 3 we detail the new variable bounding approaches as well as the rest of the approaches evaluated in this work. In Section 4 we discuss the evaluation results. Finally, Section 5 concludes.

2. Background Material

The state-of-the-art of HVAC control process optimization approaches rely on model predictive control (MPC). MPC considers a wide family of control approaches that, share the following 3 basic steps (Camacho and Alba 2013):

I. Every period (such as, every 15 minutes), an HVAC control schedule (i.e., a sequence of control actions, such as thermostat setpoints) is planned over a finite horizon into the future (e.g., the next 24 hours) using models of the system dynamics (e.g., the thermal response of the building to weather, occupancy and HVAC operation)

II. The first action of the planned schedule is executed.

III. The procedure is repeated shifting the planning horizon into the future.

A common variant of this basic algorithm includes an additional step where the model parameters are updated (i.e., learned) that is called adaptive MPC. Although adaptive MPC can suffer from stability issues (i.e., tendency for variables to swing around setpoints or go out of control) (Camacho and Alba 2013), this is not the case in HVAC control due to the slow nature of thermal dynamics (Siroky et al. 2011).

Depending on the building under control, the models that consider the system dynamics can include a thermal model, a discomfort model, a local weather model, an occupancy prediction model, an energy cost model, a renewable generation model, an energy consumption model of the rest of the building loads. Almost all such systems incorporate the first three and, in the case of buildings with a dynamic occupancy schedule, an occupancy prediction model.

In more detail, the thermal model essentially is expressed with a function that links the current thermal state (i.e., indoor temperature) of the building with the future thermal state (e.g., after 10 or 15 minutes) given an HVAC control action (e.g., a thermostat setpoint) and potentially a set of variables such as the outside temperature and the incident solar radiation. More formally a thermal model is:

$$\mathbf{x}_{t+1} = TM(\mathbf{x}_t, \mathbf{a}, \mathbf{i})$$

where $\mathbf{x}$ is the thermal state vector. $\mathbf{a}$ is the HVAC control action vector and $\mathbf{i}$ is the additional variables vector.

Furthermore, the thermal discomfort model essentially provides a quantitative metric of the discomfort experienced by the occupants. In doing so, it relies on the user-provided comfort band (e.g., heat setpoint: 68 °F, cool setpoint: 72 °F) and a metric of the deviation of the inside temperature from the bounds of the comfort-band. For instance, a very simple discomfort metric can be:

$$Disc = \begin{array}{ll} (T^{in} - T^{up})^2 & \text{if } T^{in} > T^{up} \\ (T^{in} - T^{down})^2 & \text{if } T^{in} < T^{down} \\ 0 & \text{otherwise} \end{array}$$

where Disc is the instantaneous discomfort, T^{in} stands for the indoor temperature, and T^{up} and T^{down} for the upper and lower limit of the comfort band, respectively. Discomfort is experienced only when the building is occupied. As such, in order to predict thermal discomfort and plan an optimal HVAC control schedule, one needs to also predict occupancy. Most of today's HVAC control process optimization methods for dynamic occupancy settings incorporate an occupancy prediction algorithm that utilizes available information, such as passive infrared (PIR) sensor signals, global positioning system (GPS) signals coming from mobile devices as well as WiFi traits to detect and predict occupancy (Kleimingera et al. 2013). Supported by such predictive approaches the expected instantaneous discomfort can then be calculated as:

$$\mathbb{E}[Disc] = \begin{array}{ll} O(T^{in} - T^{up})^2 & \text{if } T^{in} > T^{up} \\ O(T^{in} - T^{down})^2 & \text{if } T^{in} < T^{down} \\ 0 & \text{otherwise} \end{array}$$

where O stands for the probabilistic occupancy estimates.

Within this framework the basic difference of HVAC control process optimization approaches in dealing with expected discomfort lies within planning (i.e. step I in the MPC procedure above) and, in particular, on how the HVAC control schedule is calculated. In this context, one can distinguish two main lines of thought. The first is the thresholding approach, and the second is the weighting sum approach, which we detail in the following paragraphs.

3. Dealing with Expected Thermal Discomfort

As discussed in Section 2. several approaches have been proposed in dealing with expected discomfort, the main two approaches are the thresholding approach, and the weighted sum balancing one. In this section we first discuss the thresholding approach and then the weighted sum balancing one. Then we detail the new variable bounding approach.

3.1 Thresholding

As discussed in Section 1, thresholding the occupancy probabilities relies on an arbitrarily selected threshold to derive a deterministic occupancy schedule. A common threshold of choice is the 0.5 value where all occupancy probabilities higher than the threshold are forced to 1. Similarly, all occupancy probabilities lower than the threshold are forced to 0. An example of probabilistic occupancy estimate vector is provided below:

$$|0.3|0.2|0.5|0.6|0.6|0.7|0.5|0.5|0.4|$$

were each number between | | corresponds to a particular interval (e.g., 15-minute) within a day. By setting the threshold to 0.5 the stochastic occupancy schedule is turned into a deterministic vector where 1 corresponds to certainty of occupancy, while 0 corresponds to certainty of vacancy:

$$|0|0|1|1|1|1|1|1|0|$$

Henceforth, the optimization algorithm takes into account user-provided comfort bounds only when occupancy is anticipated (i.e., periods with deterministic occupancy equal to 1) and not otherwise. By doing so, such approaches do not cool or heat the space when the probability of occupancy is low. However, the algorithm will behave in the same way in time intervals when the probabilistic occupancy estimates are very close to 0 and in intervals when the probabilistic occupancy estimates are just below the threshold (e.g., 0.49 in the case of a 0.5 threshold). This is not optimal in many cases, as the optimization algorithm should preheat/precool a space in order for it to be closer to being comfortable when the probability of occupancy is higher. Despite these shortcomings, this approach is intuitive and simple enough that enables straightforward human-

computer interaction. In particular, the user can visualize the anticipated deterministic occupancy schedule which can help him decide when to override the intelligence. Although the user can change the threshold, the expected user interaction consists in populating the comfort bands (e.g., heating and cooling setpoints) for each day, which is a straightforward task.

3.2 Weighted Sum Balancing

As discussed in Section 1, another approach on dealing with expected discomfort is the weighted sum balancing approach. This approach utilizes the probabilistic occupancy schedules and relies on a weighting parameter to balance expected thermal discomfort and heating/cooling cost. Effectively, the comfort constraint becomes part of the optimization function. In particular, the majority of these approaches aim to minimize the weighted sum, J, of cost and discomfort of the following (or an equivalent) form:

$$J = (1 - \lambda)Cost + \lambda E[Disc]$$

where $lambda$ is the weighting parameter that ranges from 0 to 1 (excluding 0 and 1 since one objective is entirely omitted in these limit cases). Values closer to 1 indicate that one values discomfort more compared to cost while values closer to 0 indicate that one values cost more. Such approaches deal with probabilistic occupancy estimates in a mathematically concrete manner and provide Pareto optimal solutions (i.e., a set of solutions that have an optimal balance of cost and discomfort; the same cost cannot be achieved with less discomfort and/or the same discomfort cannot be achieved with less cost). Notably, the value of λ that perfectly captures the user preferences is hard to identify since cost and discomfort are measured in different units and their relationship is hard to be interpreted by a user. Nevertheless, since $lambda$ is a single parameter, it can be populated adaptively until the user preferences are met (see Panagopoulos et al. 2015 for more details on this schema).

Some hybrid approaches rely on both thresholding and weighted sum balancing (e.g., SPOT+ (Gao and Keshav 2013)). Nevertheless, such approaches have been shown to deteriorate in all categories of efficiency, usability and applicability compared with either weighted sum balancing or thresholding and hence are out of the focus of this work (Panagopoulos et al. 2015).

3.3 Variable Bounding

In this work we propose a new approach in dealing with expected thermal discomfort which relies on variable bounding. In particular, we propose an approach where the comfort bands are dynamically adjusted according to the occupancy probabilities. As such, when the occupancy probabilities are close to 1 the comfort band bounds are close to the user-provided values while when the probabilities tend to 0 the bounds become wider. This allows for greater deviation from the original comfort band as the occupancy probability becomes smaller. Nevertheless, the question of how wide should the bounds become in relation to the occupancy probability is not trivial. One can suggest that when the occupancy probability is 0 then the bounds should become infinitely wide. However, a practical limit to the boundaries is useful to avoid extremely uncomfortable indoor environment and facilitate a linear expansion function. For instance, allowing too wide a band during no occupancy can make it very difficult to bring temperature back

to comfortable conditions by the time occupancy is expected. More formally, the bands can be expanded according to the following formula:

$$T^{up'} = OT^{up} + (1-O)T^{max\ up}$$

$$T^{down'} = OT^{down} + (1-O)T^{min\ down}$$

where $T^{up'}$ and $T^{down'}$ stand for the upper and lower expanded band limit respectively while $T^{max\ up}$ and $T^{min\ down}$ stand for the upper and lower absolute limits, respectively. These upper and lower limits can also consider the safety thermal requirements of the building. Such safety requirements can ensure that the water pipes of the building do not freeze or even that the delicious chocolates in the kitchen cabinet do not melt.

4. Evaluation

In this section we provide a thorough evaluation of the above detailed approaches with respect to their efficiency, usability and applicability. In particular, we first discuss our evaluation case study and then we provide the results and a detailed discussion. Then we highlight the importance of conceptually separating the metric measuring the origin (i.e., cause) of the discomfort from the balancing parameter (e.g., the threshold, λ and the expanded band limits) which we consider to be a flaw in previous approaches.

### 4.1 Evaluation Case Study and Instantiation

For our evaluation case study, we consider a small building in Berkeley, California USA. The building has 4 thermal zones. We choose one zone for our evaluation that we call T2 zone. The zone is instrumented using wireless occupancy and temperature sensors which collect data and send them to a local server running our building operating system (Fierro and Culler 2015). Occupancy is predicted using an algorithm developed during an earlier project (Scott et al. 2011) that achieves ~80% accuracy (Figure 1). As can be seen in Figure 1 the miss rate for the zones considered in our case study ranges from 5% to 25% throughout the 8 hours ahead predictive horizon considered. This translates into a 95% to 75% success rate which is consistent with the efficiency achieved in previous implementations of the algorithm.

Thermal response is predicted using a simple linear formulation and, in particular:

$$T_{t+\Delta}^{in} = T_t^{in} + (c_1 T_t^{in} h + c_2 T_t^{in} c + c_3(T_t^{out} - T_t^{in}))\Delta$$

where, $T_{t+\Delta}^{in}$ is the estimated indoor temperature after Δ amount of time while T_t^{in} and T^{out} is the current indoor and outdoor temperature respectively. The parameters c_1, c_2 and c_3 are the regression coefficients to be estimated. The thermal model is estimated through least-squares fitting regression. The predictive accuracy of the model is reported in Figure 2.

The root-mean square error (RMSE) for 7 hour ahead predictions is about 2.5F for T4 which is consistent with state-of-the-art (Alam M. et al. 2017).

Figure 1: Occupancy prediction approach evaluation for all zones

Figure 2: Thermal model prediction evaluation for all zones

For our evaluation we use data for one day during the winter and evaluate our approach for this particular day through an iterative procedure, ensuring that the thermal state at the beginning and end of the day is the same. As such, we are able to provide long term evaluation results (assuming that the same day repeats itself) in feasible time. This enabled us to evaluate all approaches considered with a wide population of the corresponding balancing parameter (i.e. the threshold, the λ parameter, and the $T^{max\ up}$ $T^{min\ down}$ values, respectively). In more detail, we evaluated all approaches with the following ranges:

- **Thresholding**: A threshold within the 0-1 range with a step of 0.1.
- **Weighted Sum Balancing:** λ within the 0-1 range with a step of 0.02.
- **Variable Bounding:** $T^{max\ up}$ and $T^{min\ down}$ from the comfort band limits to 10 times the safety requirements with a variable step of 0.5 to 2F.

4.2 Main Results and Discussion

Our evaluation results are collectively reported in Figure 3. The cumulative cost in Figure 1 considers the absolute dollars spend for the day while the cumulative discomfort is measured in F^2 minutes. The latter is able to capture both the deviation form the desired temperature as well as the duration of this deviation. All approaches have similar Pareto efficiency (points closer to the origin indicate a higher Pareto efficiency). Nevertheless, the weighted sum balancing algorithm captures a wider and more evenly distributed range of balancing points, while the other two approaches seem to capture the balancing points over a narrow region (on the top left of Figure 3). Notably, the thresholding approach is also able to capture a point of minimum cost with maximum discomfort, differently from the variable bounding approach. However, this solution corresponds to the operation where HVAC is completely switched off and the cost is zero, that is trivial and should be disregarded. Given this, the thresholding solution has no practical benefits over the variable bounding one with respect to the range of balancing points (and, hence, user preferences) that is able to capture.

Since both the variable bounding and the thresholding approach produce solution points in a very narrow region, there is limited value in populating the balancing parameters with values other than the intuitive 0.5 for the threshold and the intuitive safety requirements for the bounds. Conversely, the weighted balancing approach captures a wide range of balancing points depending on the λ parameter. As such, populating the λ parameter in accordance to the user preferences is an important task. Figure 4 depicts the relationship between λ and comfort, while Figure 5 shows the relationship between λ and cost. Both cost and discomfort have a generally monotonous relationship with λ which confirms that one can progressively increase or reduce λ until the preferences are met in an adaptive manner (as also discussed in Panagopoulos et al. 2015).

Although the weighted sum balancing allows users to specify a wider range of balancing points, both the thresholding and the variable bounding approaches produce Pareto efficient solutions. As such, when a fixed operation is needed around the point of minimum discomfort, all approaches are able to produce efficient solutions. In addition, both the thresholding and the variable bounding approaches do not need users to populate any parameters.

©2018 ACEEE Summer Study on Energy Efficiency in Buildings

Figure 3: Thermal model prediction evaluation of three methods for all zones

Figure 4: Balancing Discomfort – Weighted Sum Balancing

Figure 5: Balancing Cost – Weighted Sum Balancing

4.3 The origin of the discomfort metric

The fact that the variable bounding approach provides Pareto optimal solutions gives rise to an interesting discussion with respect to the functional aspects of setpoint temperatures in fixed and dynamic occupancy settings. In static occupancy settings, setpoints are typically used to balance cost and thermal discomfort. For instance, a user reduces a setpoint temperature not only to feel more comfortable but also to minimize the cost of heating despite the fact that this action might introduce some thermal discomfort. Nevertheless, such an approach is not optimal in dynamic occupancy settings when the occupancy is not known in advance with absolute accuracy. In stochastic settings the setpoint temperature should be adjusted in accordance to the occupancy probability as supported by the demonstrated Pareto optimality of the variable bounding approach. This gives rise to a differentiation between the parameter to balance discomfort and cost and the setpoint temperatures in stochastic occupancy settings. In stochastic occupancy settings the setpoint temperature should only consider the origin of the discomfort metric, i.e., the region where the occupant feels absolute thermal comfort, and not a parameter to balance discomfort and cost. This is the case since the latter balancing needs to happen according to the occupancy probability, for instance through the parameters of one of the three approaches evaluated in this work (i.e., the threshold, the weighting parameter λ and the variable bounding limits).

5. Conclusion

In this work we evaluated three approaches for optimizing the HVAC control process in stochastic occupancy settings and proposed a new approach that relies on variable bounding. Our study confirms that the weighted sum formulation is superior with respect to the range of user preferences that is able to capture. Nevertheless, we show that all approaches evaluated are able to capture optimal solutions with minimum user input. As such, choosing one approach over the other depends on the project requirements and the desired interaction with the user. In particular, the weighted sum balancing approach provides increased flexibility in balancing discomfort and cost but requires an ongoing interaction with the occupants. In more detail the balancing parameter λ is not straightforward to be populated by the user as the users are not very prone in understanding mathematical relationships between quantifications of discomfort and cost (Panagopoulos et al. 2015). As such, an adaptive approach is required to populate the balancing parameter λ. Nevertheless, both the thresholding and the variable bounding approaches provide Pareto optimal solutions with minimum user. As such, in cases where minimum user interaction is desired any of these two methods can be used according to which one makes more intuitive sense to the user. Future work includes experimentation with nonlinear expansion of the bounds in the variable bounding approach. Depending on the non-linear function utilized, variable bounding can become more or less risky in introducing discomfort while minimizing cost.

References

Alam M., Panagopoulos A.A., Rogers A., Jennings N.R. and Scott J. 2014. Applying extended kalman filters to adaptive thermal modelling in homes. Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings.

Camacho E. F. and C. B. Alba. *Model predictive control.* Springer, 2013.

Dounis A. I. and C. Caraiscos. 2009. *Advanced control systems engineering for energy and comfort management in a building environment—a review.* Renewable and Sustainable Energy Reviews, 13(6):1246–1261

De Dear, R.J. and Brager, G.S., 2002. Thermal comfort in naturally ventilated buildings: revisions to ASHRAE Standard 55. Energy and buildings, 34(6), pp.549-561.

Energy Information Administration (EIA). 2015. *RESIDENTIAL ENERGY CONSUMPTION SURVEY (RECS).* https://www.eia.gov/consumption/residential/index.php

Fierro, G., D. Culler. 2015. *An Extensible Building Operating System.* Report UCB/EECS-2015-197 EECS Department, University of California, Berkeley http://www2.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-197.html

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

Lu J., T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. 2010. *The smart thermostat: Using occupancy sensors to save energy in homes.* In Proc. of the 8th ACM Conf. on Embedded Networked Sensor Systems, pages 211–24. ACM, 2010.

Panagopoulos, A.A., et al. 2015. *AdaHeat: A general adaptive intelligent agent for domestic heating control.* Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

Panagopoulos A.A., Maleki S., Rogers A., Venanzi M. and Jennings N.R. 2017. "Advanced economic control of electricity-based space heating systems in domestic coalitions with shared intermittent energy resources." ACM Transactions on Intelligent Systems and Technology (TIST) 8, no. 4 (2017): 59.

Panagopoulos, A.A. 2016. Efficient control of domestic space heating systems and intermittent energy resources. PhD diss., University of Southampton

Rutkowski, H., 1997. Manual RS: Comfort, Air Quality, and Efficiency by Design. Air Conditioning Contractors of America.

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

Siroky, J., Oldewurtel, F., Cigler, J., & Privara, S. 2011. Experimental analysis of model predictive control for an energy efficient building heating system. Applied Energy, 88(9), 3079-3087.

Urieli D. and P. Stone. 2013. *A learning agent for heat-pump thermostat control.* In Proc. of the 2013 Int. Conf. on Autonomous Agents and Multi-agent Systems (AAMAS), pages 1093–100. Int. Found. for Autonomous Agents and Multiagent Systems, 2013.

W. Kleimingera, F. Matterna, and S. Santinib. Predicting household occupancy for smart heating control: A comparative performance analysis of STOA approaches. Technical report, ETH Zurich, 2013.