Using Decision Making to Improve Energy Efficiency of Buildings

Omid Ardakanian and **S. Keshav**
David R. Cheriton School of Computer Science
University of Waterloo
Waterloo ON N2L 3G1, Canada
{oardakan,keshav}@uwaterloo.ca

Abstract

Building heating and cooling accounts for nearly 18-24% of all energy usage. Building energy management can reduce both its operating costs and its carbon footprint. We propose the use of a factored Partially Observable Markov Decision Process (POMDP) to model and efficiently control a commercial building heating system. We use supervised learning and ground truth data to capture the parameters of the POMDP.

Introduction

Energy consumption of residential and commercial buildings currently accounts for about 30-40% of global energy use (UNEP 2007). About 60% of this energy is used for space heating (or cooling), 18% for water heating, and only 3% is used for lighting. In high- and middle-income countries, energy is mostly generated from fossil fuel, directly contributing to global climate change. Our work, therefore, uses decision making techniques to reduce inefficiencies in space heating and cooling. Specifically, we use decision making to control the temperature of a building according to the activity pattern of its occupants. We believe that planning is useful in optimizing the thermal control problem because it takes time to warm up a house.

Related Work

Energy resource monitoring and activity recognition have been extensively studied in the literature. The SpotLight system (Kim et al. 2008) monitors a user's energy usage profile using wireless sensors, assuming that user proximity is the cause of measured energy usage. The ACme (Jiang et al. 2009) system monitors household power consumption in real time at a power outlet to help occupants understand their electricity usage pattern. It does not monitor user activity. The HydroSense (Froehlich et al. 2009) system, similarly, monitors home water usage. Finally, the ViridiScope (Kim et al. 2009) project monitors power usage indirectly by sensing signals emitted by an electrical appliance. Our work differs from these projects in two respects. First, we study space heating and cooling, which account the bulk of energy usage in most buildings. Second, we use an automated decision making system to optimize thermal control.

ICAPS 2010 POMDP Practitioners Workshop, May 12, 2010, Toronto, Canada.

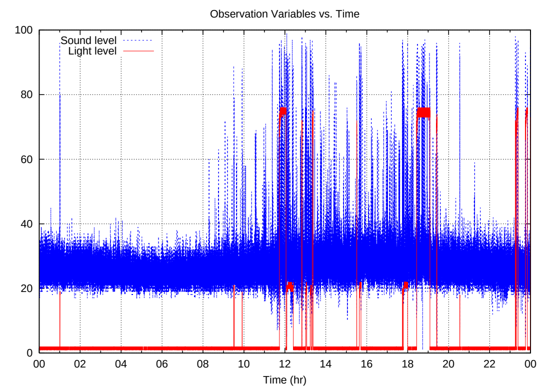

Figure 1: Sound and light measurement for a day

Overview

We propose a fully autonomous decision-making system to optimally control a building heating system. Our main goal is to not expend energy on heating or cooling a space in the building if there is no occupant currently present; although, in the initial model we only tackle the space heating problem. We say that a room in a building is *active* if it has an occupant. To preserve privacy, we do not directly observe if a room is active. Instead, we use measurements of sound and light levels to infer activity. For instance, Figure 1 shows the sound and light level we measured in a typical room over the period from midnight to midnight on a typical day. We see that sound levels increase perceptibly at around 9am and drop off at 7:30pm, resuming at around 11:30pm. Light levels also show a similar trend. We verified that these correspond to room activity by students and by custodial staff. This suggests that a control system based on these measurements would make good decisions.

We define a *policy* to be the sequence of decisions that control the temperature of a room in a building based on the control system's belief about current and past activity patterns. We use a POMDP to find optimal thermal control policies of a single room. This is a substitute for the traditional thermostat devices, which are oblivious to activity, and therefore likely to be ineffective in practice.

Model

Figure 2 is the model of the system. The state variables are the activity in a room (S^A), its temperature (S^T), and the current half-hour interval of time within a week (S^C). To keep the model tractable, temperature and clock state variables are discretized. In our current model, the domains of these state variables are defined as follows:
$Dom(S^A) = \{Active, Inactive\}$
$Dom(S^T) = \{15, 15.5, 16, ..., 29, 29.5, 30\}$
$Dom(S^C) = \{0:30, 1:00, 1:30, ..., 167:30, 168:00\}$
We plan to investigate the sensitivity of our results to the underlying discretization in future work.

The belief of being in a certain activity state is updated based on the value of two $observation\ variables$: the sound and light levels, denoted O^S and O^L respectively. Although our sensors can measure sound and light levels with high precision, we find it sufficient to define only three values for these variables: *high, normal and low*.

We define only two possible actions A: *blocking a heating or cooling vent* and *unblocking the vent*. Blocking a vent decreases energy consumption and is assumed to change the temperature by a value δ (corresponding to heat loss or gain) in one time step. Symmetrically, unblocking the vent increases energy consumption and is assumed to change the temperature by δ' in one time step. Given fixed parameters such as the size of the room, the difference between indoor and outdoor temperatures, and the capacity of the heating system, we can compute δ and δ' from analytic models. To this end, we use the monthly average outdoor temperature to compute δ and δ' for possible values of the indoor temperature. These values can also be experimentally determined, and we intend to do so in future work.

Finally the reward function R is defined as a function both of the comfort of the occupants and of the energy consumed. If a room is active, then the reward is $|T_{preferred} - T| \times C$, where $T_{preferred}$ is the preferred temperature of the room, T is the deterministically measured temperature, and C is a normalizing constant. If a room is inactive, the reward is $|T - T_{setpoint}| \times C'$ where $T_{setpoint}$ is the minimum temperature of the room and C' is another normalizing constant. Since heating increases the temperature of the room, optimizing the energy consumed for heating is equivalent to finding the optimized temperature profile.

Learning Model Parameters

To learn the standard POMDP transition and the observation functions, we have deployed 24 WeatherDuck v2 sensors that measure the temperature, humidity, light level, sound level, and air flow around them. Each sensor attaches to the serial port of one of a set of 40 Linux-based embedded systems deployed in offices, labs and public areas in three floors of the Davis Center at the University of Waterloo (Ahmed and Ismail 2009). We poll each sensor every two seconds and collect data from the sensors into a central server weekly. To get ground truth data, we record the activity inside one of the monitored rooms using a log sheet that is filled out voluntarily by room occupants. We plan to determine model parameters from the measurements and the ground truth.

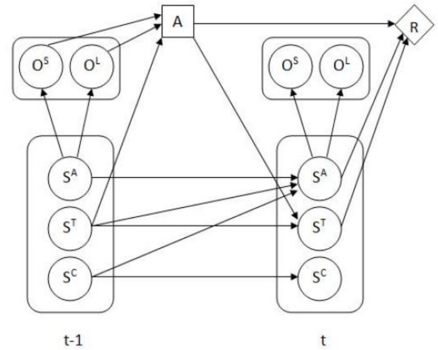

Figure 2: POMDP model

Roadmap

Here is an outline of our future work. First, in the parameter estimation phase, we plan to use ground truth to determine model parameters. We will also experimentally determine δ and δ'. Second, in the optimization phase, we plan to use the Symbolic Perseus package (Poupart 2005) for solving factored POMDPs to determine optimal policies. Third, in the evaluation phase, we plan to compute the averaged total discounted reward for one or more of these policies corresponding to users' satisfaction and the energy saved using that policy. Finally, based on the intuition gained from this work, we plan to refine the model, for instance, analyzing the sensitivity of the results to the choice of discretization, the length of a time step and the setpoint temperature.

Conclusion

We present the use of POMDPs to model and solve the real world problem of measurement-based building thermal control. We believe that the use of a sophisticated decision-making approach coupled with a large-scale sensor deployment and field measurements will allow us to make significant gains in energy efficiency.

References

Ahmed, N., and Ismail, U. 2009. Designing a High Performance WLAN Testbed for Centralized Control. *Testbeds and Research Infrastructures for the Development of Networks and Communities, International Conference on* 1–6.

Froehlich, J. E.; Larson, E.; Campbell, T.; Haggerty, C.; Fogarty, J.; and Patel, S. N. 2009. HydroSense: Infrastructure-Mediated Single-Point Sensing of Whole-Home Water Activity. In *Ubicomp '09: Proceedings of the 11th international conference on Ubiquitous computing*, 235–244. New York, NY, USA: ACM.

Jiang, X.; Dawson-Haggerty, S.; Dutta, P.; and Culler, D. 2009. Design and Implementation of a High-Fidelity AC Metering Network. In *IPSN '09: Proceedings of the 2009 International Conference on Information Processing in Sensor Networks*, 253–264. Washington, DC, USA: IEEE Computer Society.

Kim, Y.; Charbiwala, Z. M.; Singhania, A.; Schmid, T.; and Srivastava, M. B. 2008. SPOTLIGHT: Personal Natural Resource Consumption Profiler. In *HotEmNets 2008*.

Kim, Y.; Schmid, T.; Charbiwala, Z. M.; and Srivastava, M. B. 2009. ViridiScope: Design and Implementation of a Fine Grained Power Monitoring System for Homes. In *Ubicomp '09: Proceedings of the 11th international conference on Ubiquitous computing*, 245–254. New York, NY, USA: ACM.

Poupart, P. 2005. *Exploiting structure to efficiently solve large scale partially observable markov decision processes*. Ph.D. Dissertation, Toronto, Ont., Canada, Canada.

UNEP. 2007. *Building and Climate Change: Status, Challenges and Opportunities*. United Nations Environment Program.