Multiagent Systems for Sustainable Energy Applications (Extended Abstract)

Jun-young Kwak
University of Southern California, Los Angeles, CA 90089
junyounk@usc.edu

ABSTRACT

Sustainable energy domains have become extremely important due to the significant growth in energy usage. Building multiagent systems for real world energy applications raises several research challenges regarding scalability, multiple competing objectives to be optimized, model uncertainty, and complexity in deploying the system.

Motivated by these challenges, my thesis proposes a new set of models and algorithms to conserve building energy. My thesis contributes to a very new area that requires considering large-scale multi-objective optimization as well as uncertainty over occupant preferences when negotiating energy reduction. My work has shown significant potential for energy savings by investigating effective and tailored methods in the multiagent system. The suggested methods have been verified in a validated simulation testbed and included a human subject study in the real-world as a trial study.

1. INTRODUCTION

The rapid growth in energy usage from commercial buildings in the U.S. has made the need for systems that aid in reducing energy consumption a top priority. Commercial buildings in the U.S. spent 18.5 QBtu in 2008, representing 46.2% of building energy consumption and 18.4% of U.S. energy consumption [1]. To that end, this work studies an innovative multiagent system to conserve building energy, specifically focusing on proposing new algorithms to be deployed at commercial buildings (e.g., Ralph & Goldy Lewis Hall (RGL) at the University of Southern California (Figure 1(a))).

The purpose of a sustainable energy system is to efficiently conserve energy in the real-world, which raises four major technical research challenges. First, there are inherently multiple competing objectives like limited energy supplies, demands to satisfy occupants' comfort levels, additional costs to maintain the system, etc. It makes the problem harder as we need to explicitly consider multi-objective optimization techniques. Second, in such a complex domain, precisely knowing the world model is very challenging. Furthermore, as human occupants are directly involved in the optimization procedures, we must understand human behavior models and simultaneously reason about such model uncertainty in the domain. In addition, we should address novel scenarios that require agents to negotiate with groups of building occupants to conserve energy; previous work has typically focused on agents' negotiation with individual occupants. Third, a large number of heterogeneous agents including human occupants, building components and building facility management system are involved in computing an optimal policy to save energy, which requires efficient and scalable multiagent systems and thus makes the problem challenging for existing techniques to be applied. Fourth, this system should be actually deployed and verified in real testbed buildings, which adds another layer of complexity.

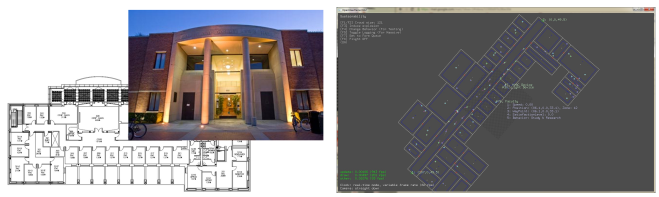

(a) Educational Building (RGL) (b) Simulation Testbed

Figure 1: The actual research testbed at the University of Southern California and our simulator

The desired goal of sustainable energy problems in real buildings is to achieve maximum energy savings without sacrificing the comfort level of occupants. Researchers have been developing multiagent systems to conserve energy for deployment in smart grids and in buildings [2, 6, 7, 9]. However, their work has been done with a particular focus on residential buildings and has not considered the combination of above research challenges in sustainable energy domains. To overcome weaknesses of prior work, I have developed a new set of models and algorithms to efficiently compute optimal policies for sustainable energy problems. In particular, my contributions are as follows: (i) optimized multiple competing objectives under uncertainty to compute well-balanced BM-MDP policies across objectives; (ii) considered model uncertainty and a novel scenario of group negotiations with human occupants to effectively conserve energy in commercial buildings, where human occupants do not have a direct financial incentive in saving energy and thus requires a fundamentally different mechanism to effectively motivate occupants; and (iii) developed efficient and scalable decision theoretic methods to solve large-scale real-world problems.

2. CONTRIBUTIONS

My work provides newer models and algorithms specifically designed to tackle technical research challenges to efficiently achieve energy savings in commercial buildings.

Multi-objective Optimization under Uncertainty: Limited availability of energy resources has necessitated the development of efficient measures of conserving energy. Along this line, I have developed a novel algorithm for generating optimal *Bounded parameter Multi-objective MDPs* (BM-MDPs) policies that explicitly considers multiple criteria optimization (energy and personal comfort) [3] based on an understanding of how to generate a policy in the presence of such multiple objectives that are not aggregated into one single value. The key principle, given the current domain of non-residential buildings is one of fairness; we wish to reduce energy usage, but we cannot sacrifice any one individual's comfort entirely in service of this goal. To meet this requirement, my work focused on minimizing the maximum regret instead of maximizing the reward value based on a min-max optimization technique to get a well-balanced solution.

Model Uncertainty and Human Behavior Study: In a complex domain such as ours, we only have an approximate agent model due to uncertainty over occupant preferences when negotiating energy reduction. In particular, the probabilities of occupants' acceptance of the agent's offer regarding energy savings, or the probabilities of other outcomes may not be precisely known. Indeed, precisely knowing the model is very challenging, and my work has built [3, 4, 5] to address such model uncertainty during group negotiations that include negotiating with groups of individuals to relocate meetings to smaller rooms to save energy, negotiating with multiple occupants of a shared office to reduce energy usage in the form of lights or HVACs, and others.

Scalability to Tackle Large-scale Real-world Problems: Saving energy in commercial buildings naturally involves in a large number of heterogeneous agents, and it causes a significant degree of uncertainty while reasoning about interactions among them. Although decision theoretic methods, specifically Distributed POMDPs (DEC-POMDPs), allow us to express uncertainty in outcomes of joint actions, uncertainty in observations, beliefs and rewards of different states, the key shortcoming was the computational burden to reason with DEC-POMDPs. However, over the past decade, there has been significant progress in understanding how to speed up DEC-POMDPs by orders of magnitudes, by exploiting domain and agent interaction structures and/or by introducing approximate algorithms relying on heuristic-based selective reasoning at the execution-time. TREMOR [8] and MODERN [4, 5] have been built on this understanding, exploiting sparsity of agent interactions, as well as developing heuristics to solve the given models, to allow a fundamental shift in level of scale-up that was not seen before.

3. PROBLEM DOMAIN & TESTBEDS

Jointly performed with the university facility management team, my research is based on actual occupant preferences and schedules, actual energy consumption and loss data, real sensors and hand-held devices, etc. Figure 1(a) shows one of the real testbed buildings (RGL) in which this work is to be deployed and the floor plan of the 3^{rd} floor. This campus building has three floors in total and is composed of classrooms, offices for faculty/staff, and conference rooms for meetings. Each floor has a large number of rooms and zones (a set of rooms that is controlled by specific equipment). The building includes building components such as HVAC (Heating, Ventilating, and Air Conditioning) systems, lighting systems, office electronic devices like computers and AV equipment, and human occupants are classified either permanent (faculty, staff, researchers, etc.) or temporary (students or faculty attending classes/meetings, etc.).

In this domain, there are two types of energy-related occupant behaviors that this work can influence to conserve energy use: individual and group behaviors. Individual behaviors only affect an environment where the individual is located, and group behaviors lead to changes in shared spaces and require negotiation with a group of occupants.

As an important first step in deploying this work in the actual building, I have constructed a realistic simulation testbed (Figure 1(b)) based on the open-source project OpenSteer (http://opensteer.sourceforge.net/) and validated the simulation testbed using real building energy and occupancy data. This validated simulation environment has been used to evaluate novel models and algorithms in terms of energy savings and occupants' comfort levels. Additionally, as a real-world test, I have provided results of a human subject study where this research is shown to lead human occupants to significantly reduce their energy consumption in real buildings [3].

4. FUTURE WORK

Thus far, my contributions have been in developing models and algorithms to conserve energy in commercial buildings. While these models and algorithms were verified in validated testbeds, the scope of real-world test was limited and the human behavior study was not seamlessly incorporated. To enhance usability and operability of my work, I plan to extend this work in three specific directions: (i) extend BM-MDPs to provide a well-fitted policy to real-world situations a) by introducing an additional parameter that can be specified by occupants based on their preferences and potentially learned during the operation in the real-world and b) by considering uncertainties in states and agent observations that can be captured by DEC-POMDPs; (ii) enhance the human behavior study by including standard metrics, which have been broadly investigated in social psychology studies, such as irritation and distraction factors that would primarily affect negotiation procedures to conserve energy; and (iii) develop the complete closed loop for the system by applying learning methods to adjust based on actual feedbacks from the buildings.

5. REFERENCES

[1] *Buildings Energy Data Book*. U.S. Dept. of Energy, 2010.

[2] S. Kamboj, W. Kempton, and K. S. Decker. Deploying power grid-integrated electric vehicles as a multi-agent system. In *AAMAS*, 2011.

[3] J. Kwak, P. Varakantham, R. Maheswaran, M. Tambe, F. Jazizadeh, G. Kavulya, L. Klein, B. Becerik-Gerber, T. Hayes, and W. Wood. SAVES: A sustainable multiagent application to conserve building energy considering occupants. In *AAMAS*, 2012.

[4] J. Kwak, R. Yang, Z. Yin, M. E. Taylor, and M. Tambe. Robust execution-time coordination in DEC-POMDPs under model uncertainty. In *AAMAS Workshop on Multiagent Sequential Decision Making in Uncertain Domains (MSDM)*, 2011.

[5] J. Kwak, R. Yang, Z. Yin, M. E. Taylor, and M. Tambe. Teamwork in distributed POMDPs: Execution-time coordination under model uncertainty. In *AAMAS (Extended Abstract)*, 2011.

[6] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Agent-based control for decentralised demand side management in the smart grid. In *AAMAS*, 2011.

[7] A. Rogers, S. Maleki, S. Ghosh, and N. Jennings. Adaptive home heating control through Gaussian process prediction and mathematical programming. In *AAMAS Workshop on Agent Technology for Energy Systems (ATES)*, 2011.

[8] P. Varakantham, J. Kwak, M. Taylor, J. Marecki, P. Scerri, and M. Tambe. Exploiting coordination locales in distributed POMDPs via social model shaping. In *International conference on automated planning and scheduling*, 2009.

[9] T. Voice, P. Vytelingum, S. Ramchurn, A. Rogers, and N. Jennings. Decentralised control of micro-storage in the smart grid. In *AAAI*, 2011.