Wireless Sensor Networks for Commercial Lighting Control: Decision Making with Multi-agent Systems

Jaspal S. Sandhu, Alice M. Agogino, Adrian K. Agogino

Department of Mechanical Engineering
University of California, Berkeley
Berkeley, CA 94720
{jaspal, aagogino}@me.berkeley.edu

University of California, Santa Cruz
NASA Ames Research Center
Moffett Field, CA 94035
adrian@email.arc.nasa.gov

Introduction

The application of wireless sensor networks to commercial lighting control provides a practical application that can benefit directly from artificial intelligence techniques. This application requires decision making in the face of uncertainty, with needs for system self-configuration and learning. Such a system is particularly well-suited to the evaluation of multi-agent techniques involving distributed learning.

Two-thirds of electricity generated in the US is for commercial buildings, and lighting consumes 40% of this. An additional 45% energy savings are possible through the use of occupant and light sensors (Yozell-Epstein 2003). The goal in this domain is to leverage wireless sensor networks (WSN) to create an intelligent, economical solution for reducing energy costs - and overall societal energy usage - while improving individual lighting comfort levels.

Much of the prior work in intelligent lighting control involves building control systems that focus on HVAC (heating, ventilation, and air-conditioning), security or other aspects of building management. Several groups have examined the use of multi-agent systems (MAS) for building control; however, this prior work varies significantly from the presented WSN problem.

Boman develops an MAS for decision making under uncertainty for intelligent buildings (Boman, Davidsson, and Younes 1999), though this approach requires complex agents. The system allocates one agent per room and the agents make use of pronouncers (centralized decision support), where decision trees and influence diagrams are used for decision-making. The primary application of this project is heating – with some basic light switching functionality – utilizing an active badge system for occupancy awareness. Also notable is that the decision making is based on static, pre-programmed plans; there are no provisions for learning.

Sharples and Hagras describe an MAS approach for intelligent residential buildings, focusing on adaptability and learning (Sharples, Callaghan, and Clarke 1999; Hagras et al. 2003). This approach utilizes a behavior-based approach drawn from robotics research (Brooks' subsumption architecture). Like Boman, the logical unit of the MAS is a room; these largely independent, room-based agents are not highly dependent on agent interactivity for decision making.

Agogino demonstrates the efficacy of influence diagrams (ID) for decision making in WSN-based lighting control. Combined with sensor fusion and validation techniques (Agogino, Granderson, and Qiu 2002), this model enables a system that can handle sensor uncertainty; in particular, it can provide functionality for dynamic electricity pricing, individual user preferences, natural light contribution, and sensed illuminance.

Several potential benefits can be gained from the use of MAS in this setting. The ID model – as is the case with many other WSN applications – depends on centralized base stations (or access points) for decision support or intensive data processing; a distributed agent-based approach, however, does not require such base stations. An MAS based on simple agent programs thus has the potential to support a highly deployable and scalable system. Additionally, distributed learning provides a robust means for system self-configuration and adaptation.

System Design

The proposed system consists of wireless sensor nodes located throughout the physical environment for purposes of sensing (light, temperature, and occupancy), actuation, and communication. Multiple sensors per node may be necessary for practical deployment; since a particular node may not need to use all sensors – or because it may simply act as a communication relay – dynamic resource allocation may be needed. All actuation will occur in ceiling-mounted, dimmable lighting ballasts.

Primary design requirements are the inclusion of individual user preferences and the ability for the user to override the intelligent system. The most desirable automatic daylighting systems control overhead lighting

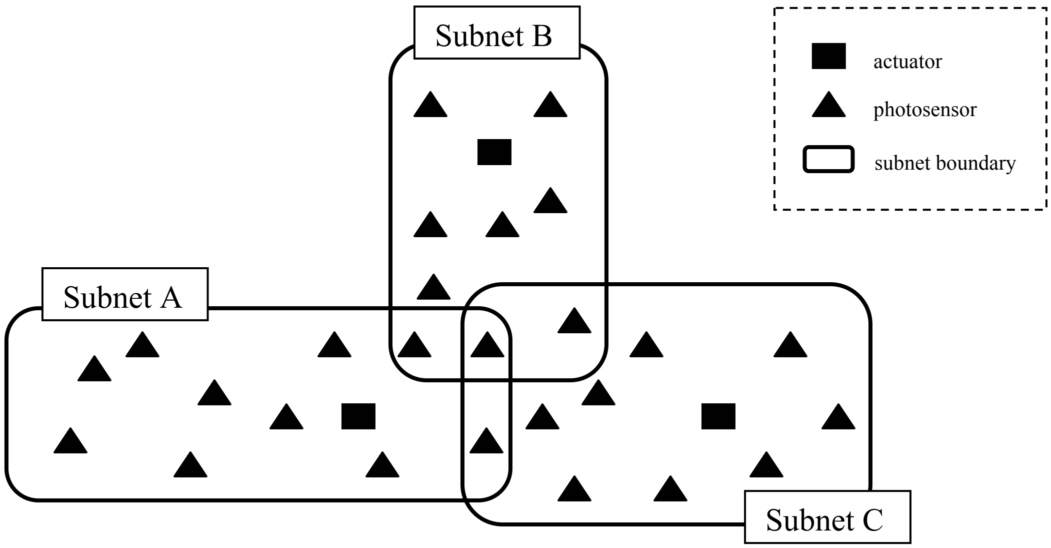

Figure 1. A subnet may rely on sensors that simultaneously belong to neighboring subnets. Each subnet is characterized by an MAS controlling a dimmable lighting ballast.

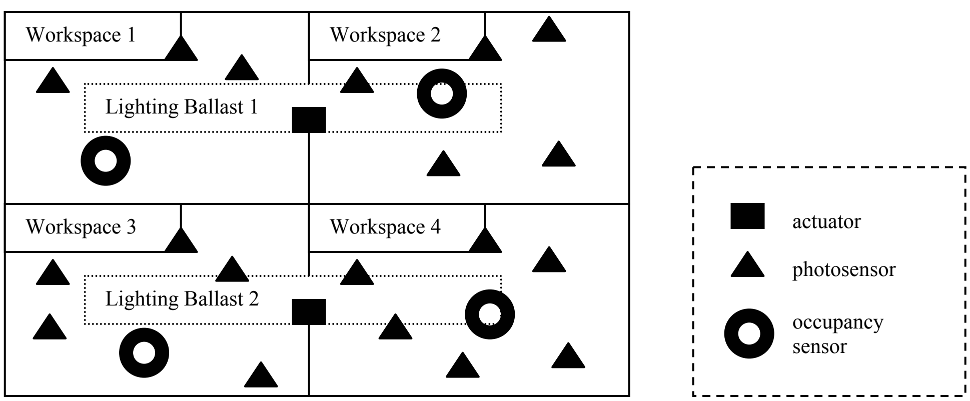

Figure 2. Workspace-based configuration for two subnets within a commercial office space. Occupancy sensors may, for example, take the form of an accelerometer attached to the user's chair. Neither node quantity nor node location is pre-defined in this configuration.

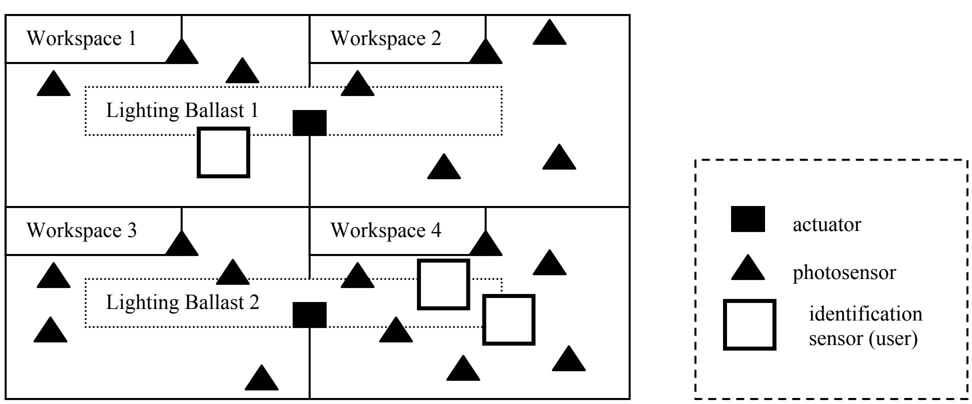

Figure 3. User-based configuration for the same physical space as in Figure 2. Identification sensors carried by the users decouple personal preferences from the workspace. Consequently, this model can handle multiple users within a single workspace, as illustrated.

but allow users to manually adjust desktop lighting (Yozell-Epstein 2003). In order to maintain a practical system it will be necessary to encode user preferences into the system and provide methods for modifying these preferences.

The overall system for a building will functionally be decomposed into many smaller pseudo-static[1] subnets since only local sensing affects local lighting actuation (Figure 1). With a single agent per node, these subnets still present multi-agent coordination problems. Within this framework, single nodes may belong to multiple adjacent subnets. While much sensor network literature predicts future networks on the order of hundreds or thousands of nodes, practical solutions to the presented problem can be accomplished with tens of nodes per subnet. At the same time that this scale makes the problem presently tractable, it also provides barriers to successful use of probabilistic techniques.

Research shows that different users not only have different preferences from one another, but also have widely varying lighting preferences for different tasks (Granderson et al. 2004). Two models are considered for achieving user preferences: workspace-based (Figure 2) and user-based (Figure 3) configurations.

If a single workspace maps to an individual user, the workspace-based preference model is acceptable. This model requires occupancy sensors to determine the presence of an individual, but does not reliably distinguish between users because of technical limitations of the sensors.

The more flexible user-based model entails the use of sensor nodes for user identification. While some researchers have integrated badges into their building control systems (Boman, Davidsson, and Younes 1999), others have indicated a need for less invasive approaches (Brooks 1997). Much as individuals in office environments carry RFID (radio frequency identification) access control badges, they could carry small-scale sensor nodes embedded in their identification cards. These nodes would be used to identify users to local sensors, and would contain lighting preference information. In this configuration, photosensors would still be located at the work surface; it would not be necessary to have occupancy sensors, as the identification nodes would assume their function. This system is more amenable to the addition, deletion, or modification of user preferences, since all the relevant information is stored on the identification nodes.

Adoption of the user-based configuration is highly dependent on privacy concerns. It remains unclear as to whether users will be willing to carry nodes to support the function of the intelligent lighting system. Unwillingness to participate presents an insurmountable barrier to the selection of this model.

[1] Mobility is not a feature of this application, but the system should be able to handle perturbations to node locations. This motivates system self-configuration and location awareness.

Agent Architecture

The primary goal of an MAS-based approach is to emulate the success of the ID model in a distributed manner. In particular, the interaction among the agents must emulate sensor validation and fusion techniques. Additionally, the decision making process must account for factors such as user preferences and variable electricity pricing.

There are many challenges to the design and implementation of a successful MAS for this application. Many of the stated challenges are more generally applicable to designing MAS solutions for WSN problems. Simple agents are necessary because of the limited memory and processing associated with each sensor node. Limited radio communication among the nodes is necessary to conserve power. Location awareness and reconfiguration are necessary aspects of a robust system. The system must be able to handle latency and time asynchronicity gracefully, due to communication constraints.[2]

Agent interaction is an essential aspect of this architecture; because of the communication and power constraints of sensor networks, agent interaction must be highly efficient. Multiple agents will contribute to the control of a given lighting actuator. In continuous domains such as this, control can be achieved by averaging agent actions or taking the median of their actions. Additionally, confidence values can be used to attenuate the global effects of aberrant local actions. When it is only necessary for the actuator to take on a fixed number of values[3], control can be achieved by voting on what action to take. These methods allow a solution to be formed based on information from multiple sensors in disparate locations. They also add redundancy and noise reduction allowing the system to overcome faulty sensors.

Many have used online learning techniques in automated building control systems, though the solutions tend to require significant computation and consequently centralized support (Mozer et al. 1995; Sharples, Callaghan, and Clarke 1999; Chang and Mahdavi 2001; Hagras et al. 2003). In order to avoid the need for centralization, this system must be able to learn in a distributed manner; depending on the information available to the agents, supervised and reinforcement learning are the two major classes of learning that apply to this environment.

For supervised learning, training data could be obtained in a user-interactive setting. The system could record sensor values (e.g. – user presence, sensed illuminance) and actions (e.g. – user `bob` sets light `#7` to intensity

[2] Latency is also problematic in the case of online learning, as delayed human response affects learning feedback.

[3] A discrete value implementation would utilize bi-level switching, where independent switching of bulbs in a ballast results in three effective lighting states: off, low, and high. The primary advantage of bi-level switching over continuously dimmable ballasts is cost.

0.62). The sensor readings would be the inputs to the supervised learner and the user actions would be the target values. Using the sensor readings and target values as training data, the supervised learner would learn to map the sensor values to the target values. After the agents are trained, they would control the individual lights based on the system sensor readings. One drawback is the real-time flexibility of a supervised learning system, as an agent encountering a state that it has not previously observed may take unreasonable actions depending of the generalization capabilities of the learning system. External rules may have to be included to bound the actions that the real-time learning system can take.

The supervised learning method assumes that the agents know the action of the user. The action could be ascertained if there is direct feedback about the user's actions, such as the lights communicating their precise settings. The task of the supervised learner here is to minimize the difference between its action, a_L, and the action of the user, a_U. If there is no such information (e.g. – because of communication limitations), the user's actions could be computed indirectly from the sensor readings if there is a precise model relating the user's action to sensor inputs. For instance, with such a model the agents could compute the light intensity set by the user based on the light readings of the agent's sensors. In this case the supervised learner would try to minimize the difference between its action, a_L, and the model of the user's action based on the sensor readings, s: $a_U(s)$. However, when such a model does not exists or is too difficult to compute in a complicated real-world environment, model-free reinforcement learning methods can be used to train the agents (Kaelbing et al. 1996). Here an agent would use the user's location and light readings as the state space for the reinforcement learner and would attempt take actions that lead to the appropriate light settings. In other words, the agent would base its learning on its perceptions relating to its actions. The reinforcement of the reinforcement learner would be based on how well the agent's actions emulated the users actions based on its light readings. For instance the absolute value of the difference between the light reading recorded after the agent's action, and the light reading previously recorded after the user's action in the same state would be used as negative reinforcement (punishment). In this case the task of reinforcement learning is to take actions a_L that minimize the difference between next state that would have resulted from the user's action $s_{t+1}(s_t, a_u)$ and the next state resulting from the agent's action, $s_{t+1}(s_t, a_a)$.The reinforcement learner would eventually learn to emulate the actions of the user, using only information from the sensors.

Even if direct information about the user's actions is known, reinforcement learning may still be useful. For instance the user overriding the ceiling-mounted light settings or adjusting the manual desktop lamps could be used as reinforcement. Also, if an agent turns on a light and the user turns it off, the agent would receive a negative reinforcement. If the person does not change anything the agent would receive a positive reinforcement (reward). In this domain of dimmable ballasts, continuous reward functions are appropriate. For example, if the user adjusts the brightness of a light slightly, the agent would only receive a small negative reward. An example agent reward function is given below.

$$R = -E(s, a) - U_A(a) - U_M(a) + P(s, a) - T(a)$$

Here a is the agent's action; s is the current state perceived by the agent; E is a function of the energy consumption resulting from the agent action; U_A is a function of the user overriding the automated lighting system; U_M is a function of the user adjusting the manual desktop lighting; P is a function of the agent action matching the user preferences; and T is a function of the radio transmissions used by agent. Negative values are used in cases where a greater function value results in a greater penalty (e.g. – an agent voting to set a light to level 0.6 results in a greater penalty than when the agent votes to set the level to 0.4). This agent reward function directly reflects the overall objectives of this project, to minimize energy costs while improving user preferences. Based on these design objectives, the non-obvious term in the reward function concerns radio transmission. Radio transmission takes up the majority of the active power consumption in WSN (Levis et al. 2004). As such, minimizing radio transmissions is critical to extending system life.

Another possible reward system involves direct human evaluation. The user could periodically rate all the actions that the agents take, perhaps via a terminal-based survey. Such human-machine interactive techniques may be appropriate in this case since user satisfaction is a critical component of the global system utility.

In some cases there can be coordination problems when multiple agents try to learn in the same environment, especially when training data about the user's preferences are incomplete. Often an agent greedily maximizing its own reinforcement in one area can take actions that reduce the reinforcement of other agents close by. However if an agent has some ability to communicate to nearby agents, several methods exists for coordinating reinforcement learning without centralized control. In Fredslund and Mataric this distributed reinforcement is accomplished by creating global policies based on the propagation of local information to allow robot collaboration (Fredslund and Mataric 2002). The theory of collectives (Wolpert, Wheeler, and Tumer 1999) can also be used as a general framework for distributed reinforcement learning, where each agent's reinforcement values are adjusted to encourage global coordination. In addition if there are relatively few sensors, "team game" approaches can often be used facilitate global coordination among agents (Crites and Barto 1996).

Future Work

Successful design of an MAS for WSN-based commercial lighting control requires us to address issues in several interrelated areas including systems design, decision making, machine learning, topology management, message routing, and HCI (human-computer interaction). Future planned work includes comprehensive assessment of the user-based and workspace-based system designs from the user perspective, with a particular emphasis on privacy issues. Beyond this, a simulation and visualization environment to support MAS experimentation will be developed. These tools will characterize both the state of the controlled environment (lighting conditions, sensor locations, occupancy) and the state of agent interactivity. Simulation for development of an appropriate agent system is a necessary predecessor to physical experimentation as sources of uncertainty can be better tracked and controlled in a virtual environment. Although simulation is a necessary step in the development of MAS for the WSN commercial lighting control problem, the eventual goal of this work is a physical implementation and evaluation of this system.

Acknowledgements

The authors would like to thank Jessica Granderson and Yao-Jung Wen for their closely related work and for constructive discussions regarding this paper. Jaspal Sandhu was partially funded by the University of California MICRO Program and General Electric, Corporate Research and Development.

References

Agogino, A.M. 2004. *MEMS 'Smart Dust Motes' for Designing, Monitoring and Enabling Efficient Lighting,* Final report to UC Project MICRO (2002-03).

Agogino, A.M.; Granderson, J.; and Qiu, S. 2002. *Intelligent Sensor Validation and Fusion with Distributed 'MEMS Dust' Sensors*, AAAS 2002 Spring Symposium on Information Refinement and Revision for Decision Making: Modeling for Diagnostics, Prognostics, and Prediction, 51-58, AAAI Press.

Boman, M.; Davidsson, P.; and Younes, H.L. 1999. *Artificial Decision Making Under Uncertainty in Intelligent Buildings*, Fifteenth Conference on Uncertainty in Artificial Intelligence, 65-70, Morgan Kaufmann.

Brooks, R.A. 1997. *The Intelligent Room Project*, Second International Cognitive Technology Conference.

Chang, S.; and Mahdavi, A. 2001. *A Hybrid System for Daylight Responsive Lighting Control*, Seventh International Building Simulation (IBPSA) Conference, 2:661-668.

Crites R.H.; and Barto A.G. 1996. *Improving elevator performance using reinforcement learning*. In D.S. Touretzky, M.C. Mozer, and M.E. Hasselmo, eds., *Advances in Neural Information Processing Systems 8*, 1017–1023, MIT Press: Cambridge.

Fredslund J.; and Mataric M.J. 2002. *Robots in formation using local information*. In Proceedings, 7th International Conference on Intelligent Autonomous Systems (IAS-7), 100–107.

Granderson, J.; Agogino, A.M.; Wen, Y.; and Goebel, K. 2004. *Towards Demand-Responsive Intelligent Daylighting with Wireless Sensing and Actuation*. Submitted to IESNA.

Hagras, H.; Callaghan, V.; Colley, M.; and Clarke, G. 2003. *A Hierarchical Fuzzy Genetic Multi-Agent Architecture for Intelligent Buildings Learning, Adaptation and Control*, The International Journal of Information Sciences, 150:33-54.

Kaelbing L.P.; Littman M.L.; and Moore A.W. 1996. *Reinforcement learning: A survey*. Journal of Artificial Intelligence Research, 4:237–285.

Levis P.; Madden S.; Gay D.; Polastre J.; Szewczyk R.; Woo A.; Brewer E.; and Culler D. 2004. *The Emergence of Networking Abstractions and Techniques in TinyOS*, First USENIX/ACM Symposium on Networked Systems Design and Implementation.

Mozer, M.C.; Dodier, R.H.; Anderson, M.; Vidmar, L.; Cruickshank III, R.F.; and Miller, D. 1995. *The Neural Network House: An Overview*. In Current Trends in Connectionism, Erlbaum: Hillsdale, NJ.

Sharples, S.; Callaghan, V.; and Clarke, G. 1999. *A Multi-Agent Architecture for Intelligent Building Sensing and Control*, International Sensor Review Journal, 19(2).

Wolpert, D.H.; Wheeler, K.; and Tumer, K. 1999. *General Principles of Learning-based Multi-agent Systems*. In Proceedings of the Third International Conference of Autonomous Agents, 77–83.

Yozell-Epstein, R. 2003. *Economic, Energy and User Needs Analysis of Intelligent Lighting System*, Master's Report, University of California, Berkeley.