

Intelligent Light Control using Sensor Networks

Vipul Singhvi
Civil Engineering Dept.
Carnegie Mellon University

Andreas Krause
Computer Science Dept.
Carnegie Mellon University

Carlos Guestrin
Center for Automated
Learning & Discovery
Carnegie Mellon University

James H. Garrett Jr.
Civil Engineering Dept.
Carnegie Mellon University

H. Scott Matthews
Civil Engineering Dept.
Carnegie Mellon University

ABSTRACT

Increasing user comfort and reducing operation costs have always been two primary objectives of building operations and control strategies. Current building control strategies are unable to incorporate occupant level comfort and meet the operation goals simultaneously. In this paper, we present a novel utility-based building control strategy that optimizes the tradeoff between meeting user comfort and reduction in operation cost by reducing energy usage. We present an implementation of the proposed approach as an intelligent lighting control strategy that significantly reduces energy cost. Our approach is based on a principled, decision theoretic formulation of the control task. We demonstrate the use of mobile wireless sensor networks to optimize the trade-off between fulfilling different occupants' light preferences and minimizing energy consumption. We further extend our approach to optimally exploit external light sources for additional energy savings, a process called *daylight harvesting*. Also we demonstrate that an active sensing approach can maximize the mobile sensor network's lifetime by sensing only during most informative situations. We provide efficient algorithms for solving the underlying complex optimization problems, and extensively evaluate our proposed approach in a proof-of-concept testbed using MICA2 motes and dimmable lamps. Our results indicate a significant improvement in user utility and reduced energy expenditure.

Categories and Subject Descriptors: C.3 [Special-purpose and application-based systems]: Real-time and embedded systems; G.3 [Probability and Statistics]: Experimental design, Markov processes; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence - *Intelligent Agents*; J.2 [Physical Science and Engineering]: Engineering; J.5 [Arts and Humanities]: Architecture

General Terms: Algorithms, Management, Measurement, Performance, Design, Experimentation, Verification.

Keywords: Sensor Networks, Light Control, Intelligent Buildings, Active Sensing.

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SenSys'05, November 2-4, 2005, San Diego, California, USA.
Copyright 2005 ACM 1-59593-054-X/05/0011 ...\$5.00.

1. INTRODUCTION

Increasing user comfort and reducing energy costs have always been two primary objectives of intelligent buildings [35, 15]. Typically the trade-off between meeting user preferences for indoor environmental condition and reduction in energy usage leads to a difficult optimization problem. This optimization task becomes more complex as occupants have different preferences over the state of the indoor environment and these preferences change with time. For example, people prefer different light levels, even when performing the same task. In addition, the state and usage of the indoor environment changes over time, e.g., due to changing levels of sunlight and tasks being performed by users.

1.1 Current energy centric strategies

The present generation of intelligent buildings aims to utilize computer technology to autonomously govern the building environment and try to optimize user comfort, energy consumption, safety and monitoring-functions [39]. Significant progress has been made in integrating advanced computer science concepts in developing building control strategies [7, 23]. The current control strategies can actively adapt to the changing outdoor or indoor environment of the building, which amounts to significant energy cost savings. However, these control strategies are often unable to adapt to the requirements at the occupant level due to limited knowledge about occupants' preferences and the state of their immediate environment. This leads to uncomfortable work environments, sick building syndrome and various other work related diseases [44]. In most commercial buildings, the cost of salaries and benefits exceeds energy, maintenance, and annualized construction costs or rent, by approximately a factor of 100 [43, 34]. Estimates for the United States show potential annual savings and productivity gains from improved indoor environments are about \$10 to \$30 billion from reduced sick building syndrome symptoms, and \$20 to \$160 billion from direct improvements in occupant performance that are unrelated to health [14]. These facts highlight significant problems in selecting building control strategies that focus only on energy concerns, while not addressing occupant preferences.

1.2 Current occupant centric strategies

There are various standards available that recommend certain levels for various indoor parameters [3], including lighting, heating, etc. However, they fail to capture the personal preferences that an individual might have and thus

provide a gross approximation of occupant needs. Addressing the desire of increasing occupant productivity, many new buildings are now equipped with *Personal Environment Modules* (PEMs) [4]. The goal of providing PEMs is to enable occupants to control their own environment. This often can lead to conflicts as occupants have different preferences for various environmental parameters such as light levels. Studies have shown that people tend not to use the personal control modules if they perceive conflict with the other occupants sharing their workspace [30]. This kind of behavior is common in commercial buildings with open space plans. Also, providing occupants with the entire control leaves few options for the building operator whose goal is to efficiently manage energy usage. As a result the overall benefits in building operation are decreased.

1.3 Our approach: An optimal trade-off

There is thus a trade-off in benefits, choosing between an energy efficient and a user-centered control strategy. Typically these two strategies are seen as two ends of the available control strategy spectrum. These opposing interests indicate that there is a need for a bilateral or user mediated central control strategy, which fits in the middle of the control strategy spectrum. Such a control strategy would integrate occupants' preferences, the state of their immediate indoor and outdoor environment and the operation goals, in terms of reduction in energy usage.

The challenges to develop such a balanced control strategy are threefold: First, we need to identify the preferences of individual occupants in indoor environments continuously, as preferences change over time and as the occupants move in the building. The second challenge is to gather information about the immediate indoor and outdoor environment of the occupants. The third challenge is to optimize the trade-off between meeting occupants preferences and reducing energy usage.

To address these challenges, we present a building control strategy based on a principled, decision theoretic formulation of the complex optimization task. Our control strategy integrates individual occupants' preferences and the real state of indoor and outdoor environment through a network of wireless sensors. The control strategies quantify the trade-off between meeting occupant preferences and the corresponding energy utilization. We present an implementation of the control strategy as an intelligent lighting system. Our intelligent lighting system adapts to occupant preferences, trading-off energy consumption with user satisfaction. Additionally, we optimally exploit external light sources in a process called *daylight harvesting*, to significantly reduce energy usage. We address the issue of sustaining long-term deployment of wireless sensors by developing active sensing algorithms that exploit the correlation in the natural light intensity and significantly increase the lifetime of the sensors.

1.4 Contributions

In this paper, we provide the following contributions to the complex problem described above.

- We provide a principled, decision theoretic basis for intelligent light management using mobile wireless sensor networks.
- We present an efficient coordinated illumination algorithm for optimizing the trade-off between meeting

user preferences and reducing energy consumption.

- Additionally, we extend our method to optimally exploit external light sources. This *daylight harvesting* strategy significantly reduces energy usage.
- We propose an *active sensing* strategy which exploits spatial and temporal correlation in natural light intensity to schedule sensing. Our results indicate a drastic decrease in the amount of sensing necessary per day. This leads to a significant increase in battery lifetime of the deployed sensors.
- Furthermore, we extend our methods to adapt to the dynamics of moving occupants of the building. This extension uses probabilistic models for movement which naturally fit into our decision theoretic framework.
- All our proposed novel methods are extensively evaluated in an actual deployed testbed. In addition to the empirical evaluation provided in this paper, we created videos demonstrating the performance of our system. These videos are accessible at the website: www.cs.cmu.edu/~gustrin/Publications/SenSys05.

2. COORDINATED ILLUMINATION

Lighting is a major factor affecting both occupants' comfort and energy costs in a building. With a total area of 12 billion square feet, U.S. office buildings use over 86 billion kWh for lighting each year. At the 2003 U.S. average energy cost of \$0.08 per kWh, the potential savings from implementing energy efficient lighting in 50% of office building is more than \$2.1 billion each year [10]. However, at the same time, energy efficient lighting is typically associated with reduced lighting which can affect the productivity of the occupants. Reduced productivity and costs incurred due to loss of work can outweigh energy savings.

Many commercial buildings have installed intelligent light controls that enable occupants to adjust levels according to their preferences. However, research has shown that occupants avoid using such controls if they perceive conflict with other occupants [30]. This is the case in many commercial offices with open floor plans. Achieving occupants' light preferences and reducing energy usage requires solving a multi-criterion optimization problem for which the complexity grows exponentially with the number of light fixtures.

Fortunately, we can exploit architectural design principle to solve the optimization problems efficiently and optimally. The key concept we will utilize is *zoning*: the process of dividing the floor plan in smaller zones which have separate light controls. However, in open floor plans it is difficult to separate the effect of a light source in one zone from the adjacent zones. This leads to a situation where a single light source affects multiple zones, making it difficult to find a mutually acceptable setting across adjacent zones.

In this section, we present a coordinated illumination approach, which exploits the zoning concept, in lighting system design, to solve the described multi-criterion trade-off problem optimally and efficiently. We first formalize the control strategy in a decision theoretic way. We then propose an efficient algorithm for solving the arising optimization problem, and evaluate the proposed method using our testbed.

2.1 Problem definition

Typically, building occupants prefer environmental conditions in which they are more productive than the conditions

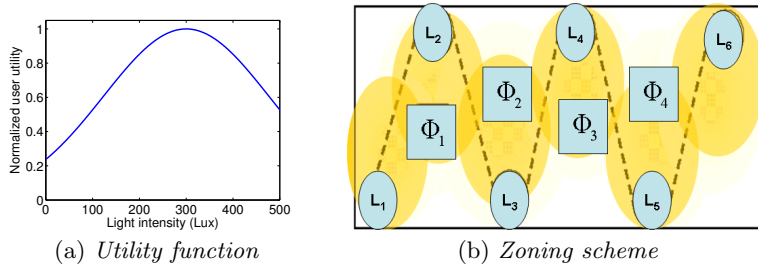


Figure 1: Utility function, zoning/locality concept.

in which they feel uncomfortable. We capture this preferences ordering by representing them as utility functions, which can be different for each occupant. Higher utility function value means higher satisfaction for the occupant.

In our formulation, each building occupant i has a utility function $\Phi_i(x_i, \mathbf{a})$, representing their preference at a given location x_i for a given light setting \mathbf{a} , where $\mathbf{a} = (a_1, \dots, a_m)$ is a vector specifying settings for each lamp. For example, a particular setting \mathbf{a} will lead to a particular light intensity at position x_i , Φ_i will indicate how happy the i th user is with this light intensity. Specifying this utility function, thus, requires us to learn a mapping from lamp settings \mathbf{a} to Lux levels at each location x_i . We learn such mappings from data using standard techniques, as described in Section 2.3. More generally, the utility function may depend on other considerations that affect user comfort, such as glare, evenness of light, and so on.

A building has multiple occupants with varying preferences and hence with varying utility functions Φ_1, \dots, Φ_n . When considering only user satisfaction, our objective is to find the joint setting \mathbf{a} of the lamps that maximizes the sum of the values of the utility functions of the occupants that are being affected by those lamp settings:

$$\operatorname{argmax}_{\mathbf{a}} \sum_{i=1}^n \Phi_i(\mathbf{a}, x_i).$$

In some applications, the preferences of some users may be weighed higher than others.

This simple formulation, however, does not consider energy costs. To address this issue, we define the building operator utility function Ψ , which decreases monotonically with the amount of energy expended. We must thus trade-off operator utility Ψ with user utilities $\sum_{i=1}^n \Phi_i$. A common technique for solving such multi-criterion optimization problems is *scalarization* [8]: A new objective value is computed as a linear combination of the individual objective functions. In this setting, the scalarized total utility is defined as

$$U(\mathbf{a}, \mathbf{x}) = \sum_{i=1}^n \Phi_i(\mathbf{a}, x_i) + \gamma \Psi(\mathbf{a}),$$

where γ is the scalarization, or trade-off, parameter regulating the penalty for high energy expenditure, and $\mathbf{x} = (x_1, \dots, x_n)$ is the vector of occupant locations. For a fixed vector of locations \mathbf{x} , the optimal deterministic control strategy is to maximize

$$\mathbf{a}^* = \operatorname{argmax}_{\mathbf{a}} U(\mathbf{a}, \mathbf{x}). \quad (1)$$

We initially address the static case where \mathbf{x} is assumed to be known and constant. This assumption is relaxed in

Section 5. We furthermore assume that the lamps can be actuated within a discrete set of d distinct settings. Hence, the maximum utility in (1) is only over a finite set of d^m values. Note that this quantity is exponential in the number m of lamps, and enumeration of all possible settings is infeasible – in our experimental setup, we deployed $m = 10$ lamps, with $d = 10$ settings each, which is already beyond scope of enumeration. Even if the total utility function is concave, we are dealing with a discrete optimization problem, and any continuous solution would only be an approximation.

2.2 An efficient algorithm

A simple heuristic for optimizing light settings – corresponding to greedy interactions of occupants with building light systems – would be to, for some fixed order, iterate over the building occupants and have each of them adjust the light settings to maximize their own comfort. Unfortunately, since the preferences of the occupants are not considered simultaneously, this greedy approach is suboptimal.

Fortunately, we are not required to rely on such heuristics: we can use *locality*, or zoning, in our problem structure to solve the discrete optimization problem efficiently and exactly. The key observation is, due to the use of the zoning concept in lighting design, a lamp only influences occupants in the small area of space around it. For example, in Fig. 1(b) L_6 only affects the occupant represented by his utility function Φ_4 , whereas L_3 has a larger influence and affects occupants 1, 2 and 3 represented by their utility function Φ_1 , Φ_2 and Φ_3 respectively. Hence, we utilize the fact that the occupants' utility functions Φ_i only depend on a smaller subset of lamps, which drastically reduces the complexity of the problem.

Additionally, the operator's utility Ψ is also local, represented as a sum of costs of actuating the individual lamps, i.e.,

$$\Psi(\mathbf{a}) = \sum_{j=1}^m \Psi_j(a_j),$$

Using locality, we rewrite (1) as

$$\mathbf{a}^*(\mathbf{x}) = \operatorname{argmax}_{\mathbf{a}} \sum_{i=1}^n \Phi_{i,\mathbf{x}}(a_{i_1}, \dots, a_{i_k}) + \gamma \sum_{j=1}^m \Psi_j(a_j). \quad (2)$$

where i_1, \dots, i_k refer to the lamps that influence the utility of occupant i . To make this notation more explicit, we write $\Phi_{i,\mathbf{x}}(a_{i_1}, \dots, a_{i_k})$ instead of $\Phi_i(\mathbf{a}, x_i)$ to indicate the dependencies of Φ_i on the light level x_i at location i .

We thus observe that (2) requires the maximization of a sum of *local factors*, each of which depends only on a small number of variables, i.e., on the setting of only a small sub-

Input: Local factors F_1, \dots, F_k , elimination ordering $\{i_1, \dots, i_m\}$
Output: Optimal action selection \mathbf{a}
begin
 for $j = 1$ **to** m **do**
 select all factors depending on A_{i_j} , let $D_j \cup \text{dom } A_{i_j}$ be the union of their domains and let G_j be their sum;
 foreach $\mathbf{d} \in D_j$ **do**
 compute $E_j(\mathbf{d}) = \max_{a_{i_j}} G_j(a_{i_j}, \mathbf{d})$;
 compute $\hat{E}_j(\mathbf{d}) = \text{argmax}_{a_{i_j}} G_j(a_{i_j}, \mathbf{d})$;
 end
 replace all factors depending on A_{i_j} by E_j .
 end
 for $j = m$ **downto** 1 **do**
 Let $a_{i_j} := \hat{E}_j(a_{i_{j+1}}, \dots, a_{i_m})$.
 end
end

Algorithm 1: Optimal action selection.

set of the lamps. We exploit this locality by using the coordination methodology developed in [6, 19] to obtain the optimal setting for the lights. The coordination methodology maximizes the sum of the utility function of the occupants by minimally violating the occupant preferences represented by individual utility functions. Consider the following example where we want to maximize $F(y_1, y_2, y_3, y_4) = F_1(y_1, y_2) + F_2(y_1, y_3) + F_3(y_3, y_4) + F_4(y_2, y_4)$. F_1 and F_2 do not depend on the value of y_4 , hence we can “push the maximization with respect to y_4 through the sum” and write

$$\max_{\mathbf{y}} F(\mathbf{y}) = \max_{y_1, y_2, y_3} (F_1(y_1, y_2) + F_2(y_1, y_3) + \max_{y_4} F_3(y_3, y_4) + F_4(y_2, y_4)).$$

For any value of y_2, y_3 , the maximization over y_4 can be carried out independently of the value of y_1 . Letting $E(y_2, y_3) := \max_{y_4} F_3(y_3, y_4) + F_4(y_2, y_4)$, we find

$$\max_{\mathbf{y}} F(\mathbf{y}) = \max_{y_1, y_2, y_3} F_1(y_1, y_2) + F_2(y_1, y_3) + E(y_2, y_3),$$

from which y_4 has been “eliminated”. This process is repeated, eliminating the variables y_i one at a time. The values of y_1, y_2, y_3, y_4 maximizing F can be recovered by computing and plugging in the partial solutions in reverse order. Note that for this example, only a maximization over three variables had to be carried out by enumeration, not over four. In the general case, the maximal domain size of the local factors (as F_1, F_2) and intermediate factors (as E in the above example), called the *treewidth*, depends on the order in which variables are eliminated. Picking the optimal elimination ordering is **NP**-complete [2]. Fortunately, there is a large literature [5, 25, 26] on heuristics for finding good elimination orderings, of which we expect many to apply to the locality exhibited by a typical building.

Fig. 1(b) indicates the structure for our testbed. Here, the *treewidth* is three. Hence the exact optimization can be carried out in order 10^3 steps instead of 10^{10} required by naive enumeration. The algorithm is presented as Algorithm 1.

2.3 Evaluation of our algorithm

We created a testbed to emulate the real situation at a smaller scale to test our light control strategy. Our testbed




Figure 2: Testbed characteristics.

consists of twelve MICA2 motes [42] and ten 60 Watt table lamps, arranged on a 146 inches by 30 inches table. The lamps are actuated by the X10 system [45], which wirelessly communicates with a single desktop PC, and uses power lines for controlling the lamps. Each lamp can be actuated to produce ten different light intensities. The lamps are arranged in a triangular pattern which corresponds to the zoning concept as introduced in Section 2. Each of the seven triangular zones is affected by three lamps, and each lamp affects up to three regions. The motes are distributed over the different zones, and communicate with the base station over the wireless network.

An important aspect of the evaluation is the reliable prediction of the light levels in each zone, for each combination of lamp settings. We used TinyDB [1] to query the sensor values for various lamp settings. To calibrate sensors (convert the raw ADC to Lux), an external light meter [29] was used to measure ground truth Lux levels at various locations. Piecewise linear regression on a log-scale was then used to calibrate the sensors. Fig. 2(b) shows how the ADC responses varied over the different motes.

Once the sensors are calibrated, we need to determine how light settings affect light intensities in each zone. Note that each zone is affected by multiple lamps, thus, a naive approach would require us to search through all possible combinations of settings for all lamps for each location. Fortunately, since light intensities in Lux are additive, not all possible combinations of lamp settings have to be considered. We, thus, learn how each lamp affects a region, and the combined effect is simply the sum of the intensity generated by each lamp.

In order to determine power consumption, we measured the amount of power required by the lamp for each setting using a power meter. Power consumption scales approximately linearly with the ten discrete lamp settings.

Since we want to optimize user comfort, we needed to select a functional form for the user utility functions. We chose simple Gaussian kernels with varying centers and bandwidths for the different occupants, each represented by a mote. Fig. 2(a) presents the variety of utility functions we chose. The bell shape of the Gaussian kernel corresponds to the assumption that each user has a particular light intensity which is most comfortable; exceeding or falling short of the appropriate value will decrease the comfort. The choice in different bandwidths corresponds to the idea that some occupants can work well under a larger variety of conditions than others. We understand that the problem of eliciting appropriate utility functions is a difficult problem of its own. However, we believe that our particular choices of utility functions represents a variety of different preferences, which we expect to at least qualitatively correspond to real occu-

pants’ needs. Our optimization method does not make any assumption about the form of these utility functions.

Since the decision theoretic objective function proposed in Section 2.1 depends on the trade-off parameter γ , we performed an initial experiment to determine an appropriate value for our setup. Selecting an appropriate trade-off parameter is in general a difficult problem, and depends on the units in which energy and user comfort are measured. It is general practice to look for “knees” in the trade-off curves, values for which a small increase in either quantity leads to a drastic decrease of the other. Using this guideline, Fig. 3(c) shows that γ can be increased to a value of about 0.3 without any decrease of user comfort, decreasing energy cost by approximately 25%. A value of $\gamma = 0.4$ leads to a small reduction in comfort (about 7%), but achieves a significant reduction (about 33%) in energy cost. These graphs also indicate that the predicted utilities based on our regression model closely corresponds to the measured values.

We compared our optimal coordination strategy with the simple greedy heuristic mentioned in Section 2.1, in terms of total utility achieved for varying parameters of γ . Fig. 3(a) illustrates a setting in which occupants with conflicting utility functions were collocated in the same zones, while Fig. 3(b) shows the result for a setting in which sensors with conflicting utility functions are placed in adjacent zones. In both settings, the optimal, coordinated algorithm achieves significantly higher utility than the simple greedy heuristic.

2.4 Description of Video 1

We documented the experiment for optimizing the trade-off parameter γ in Video 1. As seen in the video, the occupants of each zone have different desire light intensities. In this video, γ gradually increases from 0 to 1.4. The video shows how the lamp settings adjust in order to maximize the total utility, dimming more and more as energy consumption is increasingly penalized.

3. ADAPTIVE, CLOSED-LOOP CONTROL

Electric lighting consumes 30% to 50% of the energy used in a typical commercial building [28]. Utilizing daylight offers a natural possibility for reducing the electric energy usage. In addition to energy savings, some evidence indicates that exposure to daylight reduces stress and stimulates occupants in the building [22]. An interesting study [20] shows that introduction of daylighting in retail buildings (shopping malls, grocery stores) stimulates buying behavior of the customers, increasing the sales.

The concept of utilizing sunlight for illumination is called *daylight harvesting*. Daylight harvesting is a difficult strategy to adopt in a building as daylight is highly variable and affects different locations of a building in different ways at different times. A post-occupancy evaluation of daylit buildings found that over 90% of those with offices near windows had the right amount of sunlight as opposed to 61% of people with interior offices [28]. For any daylight harvesting strategy, it is essential to have a pervasive network of sensors to measure the varying effect of daylight on different parts of the building. While construction costs of commercial buildings are typically around \$150/ft² and total utility costs are about \$1/ft² per year, the salaries and overhead for office workers are estimated at \$150/ft² per year. Thus, it is important that a daylight harvesting strategy either enhances, or at the very least, does not reduce worker productivity.

3.1 Closed-loop coordinated illumination

The approach described in Section 2.2 is *open-loop*, since we assume that the light intensities at each position are uniquely determined by the lamp settings, and no sensing is needed. For daylight harvesting, however, sun light will affect the light levels at each location, and, thus, the lamp settings must adapt to changing weather conditions and time of the day in a *closed-loop* fashion. In such settings, sensing is required at each time step to determine the effects of external light sources on the environment.

To incorporate daylight in our approach, we assume a mapping $\Theta : \mathcal{T} \times \mathcal{X} \rightarrow \mathbb{R}^+$, such that $\Theta(t, x)$ represents the sunlight intensity at time t and point in space x . The function Θ incorporates sensor readings measuring light levels. In general, we may not be able to measure sunlight levels directly, but the combination of the sun light with the lamps. Fortunately, since light levels are additive, we obtain sunlight levels by simply subtracting the intensity caused by the lamps, which can be determined as described in Section 2.2, from the total light.

We modify the local utility functions Φ_i to take daylight into account in our optimization process, and write $\Phi_{i,x_i}(\mathbf{a}, y)$, where $y = \Theta(t, x_i)$. The optimization problem then becomes to maximize

$$\mathbf{a}^*(t, \mathbf{x}) = \underset{\mathbf{a}}{\operatorname{argmax}} U(\mathbf{a}, t, \mathbf{x}), \quad (3)$$

where

$$U(\mathbf{a}, t, \mathbf{x}) = \sum_{i=1}^n \Phi_{i,x_i}(\mathbf{a}, \Theta(t, x_i)) + \gamma \sum_{j=1}^m \Psi_j(a_j).$$

The algorithm presented in Section 2.2 applies to this dynamic setting exactly as to the static setting – the only difference is that the local utility functions Φ_i depend on the sensed sun intensities.

3.2 Evaluation in our testbed

We extended the testbed described in Section 2.3 by five lamps, which we used to simulate daylight. The setup is displayed in Fig. 4(a). The sun lamps are also controlled by X10 and can be set to ten different intensities. To simulate realistic relative variations in sun intensity, we used sun trajectories for Pittsburgh, PA from the National Solar Radiation Data Base provided by the Renewable Resource Data Center [36]. We modulated the trajectories using piecewise linear functions to achieve a spatial gradient similar to the effect where the sun rises in the East and sets in the West. Fig. 4(b) shows 30 trajectories from the database for June 1990, and Fig. 4(c) shows the corresponding measurements from the testbed for the most central sensor.

To evaluate the accuracy of our closed-loop adaptive controller as sun intensities change, we ran experiments over several simulated days. The sun intensity was estimated by subtracting the light intensity produced by the lamps from the measured total light intensity. After estimating the amount of sunlight at each time step, the lamps were actuated accordingly, and the total utility was computed from a light measurement following the actuation.

Fig. 5(a) shows the average measured and predicted user utility as well as the power consumption, for a very small value of the trade-off parameter, $\gamma = 0.01$. It can be seen that the measured utility corresponds closely to the predicted utility. Furthermore, there are significant power sav-

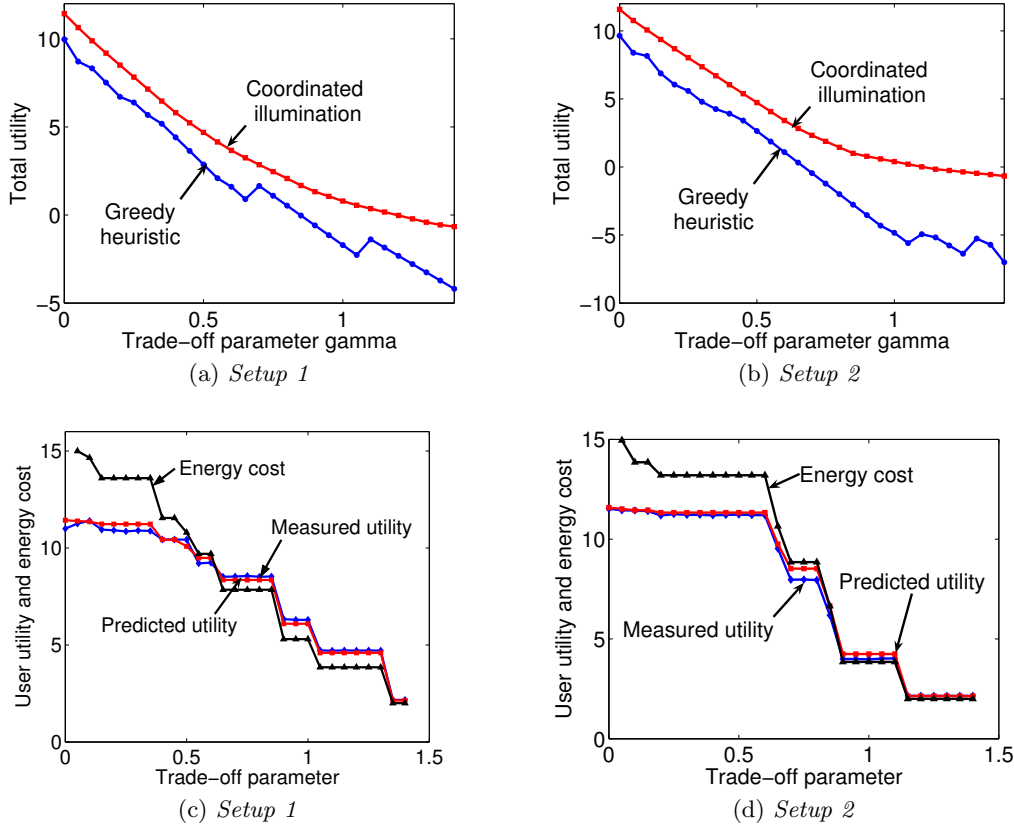


Figure 3: Comparison of greedy heuristic with coordinated illumination.

ings during the hours where sun is shining. In total, about 15 percent of energy could be saved on average compared to the amount of energy required if sunlight is not taken into account. Fig. 5(b) shows the results of the same experiment for the setting $\gamma = 0.4$ which was determined to be appropriate in Section 2. In this more energy-conscious setting, about 55 percent of the energy cost could be saved using daylight harvesting. This experiment showed the same absolute savings in energy during sunshine, but on a much lower base level: in total, about 70 percent energy could be saved in direct comparison to the setting of $\gamma = 0.01$.

The slightly reduced user utility during daytime is due to the particular selection of utility functions: Two sensors were assigned a very narrow preference (*c.f.* curve 2 in Fig. 2(a)) for a light intensity exceeded by sunlight (even with all lamps switched off). This corresponds to the assumption that sunlight can sometimes slightly blind building occupants and make them less productive. An easy extension of our model could include the actuation of blinds, dimming sunlight and avoiding this reduction in user utility. Overall, this experiment shows that, consistently over several values of the trade-off parameter, daylight harvesting achieves a significant reduction in energy usage.

3.3 Description of Video 2

We demonstrate our closed-loop control strategy in Video 2. In this video, four complete day-night cycles were recorded, and played back in increased pace. It can be seen how the table lamps adjust to varying sun intensities. A special notice

can be given to the rightmost lamp, which is often switched on even when the sun is shining, since the rightmost region requires particularly high light intensity due to our choice of utility functions.

4. ACTIVE SENSING

To achieve the benefits of the adaptive control strategy described in Section 3, we continuously require knowledge of the daylight intensity in all areas of the building where occupants are located. In our mobile sensor network vision, all building occupants are equipped with sensors, and hence this information is easily accessible. Unfortunately, continuous sensing and communication costs quickly deplete batteries and hence diminish the lifetime of the sensor network. To overcome this problem, we propose an *active sensing* strategy which schedules the sensors to sense at times that maximize the *expected utility*. Our approach builds on a recent sensor network tasking approach that uses probabilistic models to select a subset of sensors to observe [13].

4.1 Basic foundation for sensor scheduling

The key observation is that the distribution of daylight is highly correlated, both spatially across the building, and over time. This implies that sensors can sense *cooperatively*, requiring drastically fewer observations each. To simplify the discussion, we assume a discretization of space $\mathcal{X} = \{\xi_1, \dots, \xi_r\}$ and time $\mathcal{T} = \{\tau_1, \dots, \tau_l\}$. We model the spatial and temporal variation in sun intensity Θ defined in Section 3 as a stochastic process. This means that




Figure 4: Closed loop experimental setup.

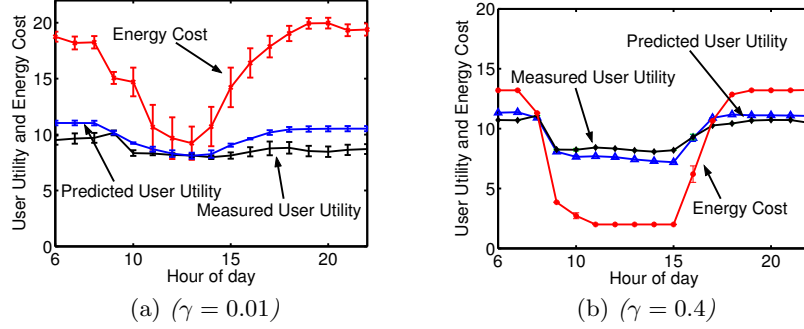


Figure 5: Adaptive, Closed-Loop control.

we consider the sunlight at each time τ_i and point x_i as a random variable $\Theta(\tau_i, \xi_i)$, and specify a joint (prior) distribution over light intensity

$$P[\Theta(\tau_1, \xi_1) = \theta_{1,1}, \dots, \Theta(\tau_l, \xi_r) = \theta_{l,r}]. \quad (4)$$

This joint distribution models stochastic temporal and spatial relationships among the random variables. These dependencies are very strong – during a dark day, most sensors are expected to measure lower values than during brighter days, mostly independently of the sensors’ positions. Also, at noon it will generally be brighter than in the late afternoon or even at night. These correlations allow us to observe only a subset of the sensors at each time step.

After observing a subset of sensors, we obtain a probability distribution over light levels at unsensed locations. Since sunlight levels are now uncertain, instead of working with local utility functions $\Phi_{i,\mathbf{x}}$ as deterministic quantities (as introduced in Section 3), these utilities now become random variables. In such cases, we address users’ *expected utility*:

$$\mathbb{E}[\Phi_{i,\mathbf{x}}(\mathbf{a}, \Theta(x_i, t))] = \sum_{\theta} P(\Theta(t, x_i) = \theta) \Phi_{i,\mathbf{x}}(\mathbf{a}, \theta), \quad (5)$$

and the optimization problem for light settings becomes the maximization of expected utility

$$\mathbf{a}^*(t, \mathbf{x}) = \underset{\mathbf{a}}{\operatorname{argmax}} EU(\mathbf{a}, t, \mathbf{x}), \quad (6)$$

where

$$EU(\mathbf{a}, t, \mathbf{x}) = \sum_{i=1}^n \mathbb{E}[\Phi_{i,\mathbf{x}}(\mathbf{a}, \Theta(x_i, t))] + \gamma \sum_{j=1}^m \Psi_j(a_j).$$

Note that this linear decomposition of expected utilities allows us to use the coordination mechanism introduced in

Section 2 to find optimal light settings, assuming the expectations can be computed efficiently. The computation of expected utilities requires probabilistic inference in the stochastic process Θ , i.e., the computation of $P(\Theta(t, x_i) = \theta)$, the probability that the sunlight $\Theta(t, x_i)$ at position x_i at time t has light level θ . These probabilities can often be computed efficiently. Even if no efficient exact inference is possible, there are efficient approximate inference algorithms available [9].

The prior distribution respresented in Eq. (4) does not take observed sensor values into account. If we observe a particular light intensity $\Theta(\xi, \tau) = \hat{\theta}$ at some time τ and position ξ , we can use Bayes rule to compute the *posterior distribution*:

$$P[\Theta(\tau_1, \xi_1) = \theta_{1,1}, \dots, \Theta(\tau_l, \xi_r) = \theta_{l,r} \mid \Theta(\xi, \tau) = \hat{\theta}],$$

representing the distribution over possible sunlight levels at each location at each time point, given the measurements. More generally, after observing a value \mathbf{o} for some sensors O , we compute the posterior distribution $P(\Theta(t, x_i) = \theta \mid O = \mathbf{o})$. Using this posterior distribution we can obtain the *conditional expected utility*:

$$\begin{aligned} EU(\mathbf{a}, t, \mathbf{x} \mid O = \mathbf{o}) = \\ \sum_{i=1}^n \sum_{\theta} P(\Theta(t, x_i) = \theta \mid O = \mathbf{o}) \Phi_{i,\mathbf{x}}(\mathbf{a}, \theta) \\ + \gamma \sum_{j=1}^m \Psi_j(a_j). \end{aligned}$$

The posterior distribution reflects our updated belief about the sun intensity distribution after the observations have been made by the sensors. This implies that updating our belief also changes the optimal lamp settings to choose in (6). Intuitively, if we are certain about the amount of natural light at a particular location, we can allocate the arti-

cial light more economically. More generally, we would like to schedule sensing times that lead to more effective decisions. This concept is formalized by the *maximum expected utility* framework: We want to find the set of observations O^* that yields best decisions (light settings) in expectation after making these observations. Specifically, we want to maximize

$$O^* = \operatorname{argmax}_O \mathcal{J}(O), \quad (7)$$

where

$$\mathcal{J}(O) = \sum_{\mathbf{o}} P(O = \mathbf{o}) \left(\sum_{t \in \mathcal{T}} \max_{\mathbf{a}} EU(\mathbf{a}, t, \mathbf{x} \mid O^{(1:t)} = \mathbf{o}^{(1:t)}) \right),$$

O is a set of observations, and $O^{(1:t)} = \mathbf{o}^{(1:t)}$ is the subset of observations up to time t . Note that we are dealing with a dynamic system, thus, when computing the expected utility at time t , we only consider measurements made up to this time. The quantity \mathcal{J} is the expected value of the expected utility, where the first expectation is computed over all possible outcomes of the measurements O , and the second expectation is taken over the sunlight intensity probability distribution. For example, if in Fig. 1(b), the sensor Φ_2 makes an observation, we will be more certain about the light intensity at the location of Φ_1 at the same time step, than if we had observed sensor Φ_4 three time steps back.

The more certain we are about the light intensity, the more accurately we can coordinate the lamp settings, and the more comfort the occupants will have. We are more certain about the sunlight level when more observations are made. In contrast, more observations will use up the sensors batteries faster. A natural approach is to constrain the number of observations made by sensors, for example, by providing a certain budget B of energy each sensor can spend for sensing during each day. Since in our proposed system, the light sensors are mobile and collocated with the building occupants, and the sensing and communication cost is assumed to be uniform over the day, we define our active sensing goal as that of specifying a schedule of B sensing times for each sensor.

4.2 Algorithms for sensor tasking

In the following, we adopt a Markov assumption for our stochastic sunlight intensity process, which states, that the state of the process at time t is conditionally independent of all states at times t' for $t' < t - 1$, given the state at time $t - 1$. This is a very natural assumption for sensor networks [13].

Unfortunately, even in Markovian settings, the optimization problem (7) is intractable [27], even for the case of two sensors. In [27], it is however shown that for one single sensor, the problem of finding the optimal B observations can be solved efficiently. We first review this approach. For a single sensor, the light intensity process forms a Markov chain over time. Their algorithm uses a dynamic programming approach, leveraging the fact, that the objective function in (7) decomposes along time, i.e., the total utility is the sum of all sub-chains separated by the times of observation. Fig. 6(a) illustrates this idea. In this graph, the sensor is scheduled to only observe at times t_1 and t_3 , indicated as $S_1^{(1)}$ and $S_3^{(1)}$. The sunlight intensity distribution at times t_4 and t_5 only depend on the observation at time t_3 , and not on the observation at time t_1 .

Using this dynamic programming approach, we can compute the expected utilities for all possible sub-chains – and hence for the total Markov chain – efficiently. If no observation is made, the following formula computes the maximum expected utility for the sub-chain from t_a to t_b as

$$\mathcal{J}_{a:b}^{(i)}(0) = \sum_{j=a+1}^{b-1} EU(\mathbf{a}, t_j, \mathbf{x} \mid \Theta(t_a, x_i)).$$

If k observations can be made, the maximum expected utility achievable for the sub-chain starting at time t_a and ending at time t_b can be recursively computed by

$$\mathcal{J}_{a:b}^{(i)}(k) = \max_{j:a < j < b} \{ EU(\mathbf{a}, t_j, \mathbf{x} \mid \Theta(t_j, x_i)) + \mathcal{J}_{a:j}^{(i)}(0) + \mathcal{J}_{j:b}^{(i)}(k-1) \}.$$

Finally, the maximum expected utility schedule from time 1 to t_{max} can be obtained by $\max_O \mathcal{J}^{(i)}(O) = \mathcal{J}_{0:t_{max}+1}^{(i)}$. We refer the reader to [27] for details and proofs.

Input: Budget B
Output: Optimal selection O of observation times
begin
 for $0 \leq a < b \leq t_{max} + 1$ **do**
 $rew := 0$;
 for $j = a + 1$ **to** $b - 1$ **do**
 $rew := rew + EU(\mathbf{a}, t_j, \mathbf{x} \mid \Theta(t_a, x_i))$;
 end
 $\mathcal{J}_{a:b}^{(i)}(0) := rew$;
 end
 for $k = 1$ **to** B **do**
 for $0 \leq a < b \leq n + 1$ **do**
 for $j = a + 1$ **to** $b - 1$ **do**
 $sel(j) := EU(\mathbf{a}, t_j, \mathbf{x} \mid \Theta(t_j, x_i)) + \mathcal{J}_{a:j}^{(i)}(0) + \mathcal{J}_{j:b}^{(i)}(k-1)$;
 end
 $\mathcal{J}_{a:b}^{(i)}(k) = \max_{a < j < b} sel(j)$;
 $\mathcal{J}_{a:b}^{(i)}(k) = \operatorname{argmax}_{a < j < b} sel(j)$;
 end
 end
 $a := 0$; $b := t_{max} + 1$; $k := B$; $O := \emptyset$;
 for $k = B$ **downto** 1 **do**
 $a := \widehat{\mathcal{J}}_{a:b}^{(i)}(k)$;
 $O := O \cup \{t_a\}$;
 end
end

Algorithm 2: Optimal observation selection for a single sensor. t_{max} denotes the maximum number of time steps considered for scheduling.

In the case of multiple sensors, this decomposition does not hold anymore, since influence can flow across chains. Fig. 6(b) visualizes this problem – there, the sun intensity distribution for sensor (2) depends on all three observations $S_1^{(1)}$ and $S_4^{(1)}$ from sensor (1) and $S_2^{(2)}$ from sensor (2). Exact computation of EU is in general not possible, because computing the expectation requires the summation over a number of terms exponential in the number of observations.

We address this complexity issue using an (approximate) extension of the decomposition approach used for single chains.




Figure 6: Scheduling sensors in dynamic processes.

More specifically, for sensor selection purposes, we only allow a sensor to take into account the most recent observations. Intuitively, this appears to be a reasonable approximation, especially if the potential scheduling times in \mathcal{T} are reasonably far apart. Formally, we replace the set of observations up to time t , $O^{(1:t)}$ by a subset $O'^{(1:t)}$, and approximate $\mathcal{J}(O)$ by $\mathcal{J}(O')$. In Fig. 6(b) for example, the total expected utility at time t_4 would be computed using only observations from time t_3 for sensor one, and time t_2 for sensor two, ignoring influence originating from observation $S_1^{(1)}$ and flowing through the chains as indicated by the dashed arrow. The following proposition proves that this approximation, while not exact, leads to the maximization of a lower bound:

PROPOSITION 1.

$$\mathcal{J}(O) \geq \mathcal{J}(O')$$

for all $O' \subseteq O$.

PROOF. This bound follows from the fact that maximization over \mathbf{a} is convex, and an application of Jensen's inequality. Using an induction argument, we simply need to show that

$$\mathcal{J}(O) \geq \mathcal{J}(\emptyset).$$

$$\begin{aligned} \mathcal{J}(O) &= \sum_{\mathbf{o}} P(O = \mathbf{o}) \left(\sum_{t \in \mathcal{T}} \max_{\mathbf{a}} EU(\mathbf{a}, t, \mathbf{x} \mid O^{(1:t)} = o^{(1:t)}) \right) \\ &\geq \sum_{t \in \mathcal{T}} \max_{\mathbf{a}} \left(\sum_{\mathbf{o}} P(O = \mathbf{o}) EU(\mathbf{a}, t, \mathbf{x} \mid O^{(1:t)} = o^{(1:t)}) \right) \\ &= \sum_{t \in \mathcal{T}} \max_{\mathbf{a}} EU(\mathbf{a}, t, \mathbf{x}) = \mathcal{J}(\emptyset) \end{aligned}$$

□

Proposition 1 proves that conditioning only on the most recent observations can only decrease our objective function, hence maximizing this simplified objective implies maximizing a lower bound on the true objective.

We propose the following heuristic for maximizing the lower bound on the expected utility. The basic algorithm will be a coordinate ascent approach, scheduling observations for each sensor individually. The maximization for the individual sensors uses the same dynamic programming trick introduced above. Since all sensors maximize the same global objective, the local optimization is guaranteed to monotonically increase the global objective, and hence must converge (to a local optima) after a finite number of steps. The procedure is formalized in Algorithm 3.

Several remarks are required in order to implement Algorithm 3. First of all, in a naive implementation, the computation $EU(\mathbf{a}, t_j, \mathbf{x} \mid O^{(1:t)})$ requires time exponential in

Input: Budget B

Output: Selection $O^{(i)}$ of observation times for each sensor

begin

while not converged do

 Select $1 \leq i \leq n$ at random;

 Use Algorithm 2 to select observations O_i for sensor i , but conditioning on current sensor scheduling O for remaining sensors (i.e. computing $EU(\mathbf{a}, t_j, \mathbf{x} \mid O^{(1:t)})$ instead of $EU(\mathbf{a}, t_j, \mathbf{x} \mid \Theta(t, x_i))$);

 Compute difference in total expected utility;

end

end

Algorithm 3: Multi-Sensor scheduling.

the number of chains. This can be avoided using a sampling approximation, using Hoeffding's inequality to derive polynomial bounds on sample complexity. Details are omitted here due to lack of space. In practice, a small number of samples appears to provide reasonable performance. Secondly, inference itself becomes intractable with an increasing number of sensors. Approximate inference algorithms such as the algorithm proposed by Boyen and Koller [9] provide a viable way around this problem.

Although we cannot in general provide performance guarantees for the procedure, we are building on an algorithm that provides an optimal schedule for each sensor in isolation, which should benefit from observations provided by the remaining sensors. Algorithm 3 will always converge, and always compute a lower bound on the expected total utility. Considering the intractability of the general problem ($\mathbf{NP}^{\mathbf{PP}}$ -completeness [27]), these properties are reassuring. In our experiments, the coordinated sensor scheduling performed very well, as discussed in Section 4.3. It is also possible to generalize this algorithm to compute *conditional plans*, for which the selection of observations depends on the observations already made. For details please refer to [27].

4.3 Experimental evaluation in our testbed

To evaluate our active sensing strategy, we first learned a probabilistic model for predicting sun intensities. We used 75 trajectories from the day-night cycles using our sunlight lamps as presented in Fig. 4(c) as training data, and 3 separate trajectories as testing data. From this training data, we learned a discrete, non-stationary dynamical Bayesian network for three sensors, with nodes for one measurement every hour per sensor, starting at 6 am till 22 am. The discretization was done using eight 50 Lux bins and the transition probabilities were computed using a Gaussian approximation to avoid overfitting. We used data from the two




Figure 7: Active sensing results.

sensors at both far ends (East and West) of our testbed and the central sensor (South) to capture changes in the sun intensities. We then used our sensor scheduling Algorithm 3 to select informative times for querying each sensor. Using this probabilistic model, we executed day-night cycles from the test set, and continuously measured light intensities from the sun lamps and the table lamps. However, we only provided the sensor measurements taken at the times selected by Algorithm 3 to the prediction routine, and actuated the lamps to maximize the expected utility as defined in Section 4.1. We used the trade-off parameter $\gamma = 0.4$ as identified in Section 2.

Fig. 7(a) presents the mean total utility achieved using no observations, one observation or ten observations per sensor each day. It can be seen that even a single observation increases the total utility close to the level achieved by continuous sensing. Fig. 7(b) shows the mean energy consumption required for the same experiment. Here, the single sensor observation strategy comes even closer to the power savings achieved for continuous sensing. Fig. 7(c) shows the mean user utility and energy savings achieved, for a number of observations varying from no observations to continuous sensing. These results imply that using the predictive model and our active sensing strategy, even a very small number of observations achieves results approximately as good as the results achieved by continuous sensing. Since the sensor network battery lifetime is in general inversely proportional to the amount of power expended for sensing and communication, we conclude that our active sensing strategy promises to lead to drastic increases in sensor network lifetime, deployment permanence and reduced maintenance cost.

5. PREDICTIVE LIGHT CONTROL

Thus far we have assumed that occupants are static. We now generalize our approach to dynamic settings, where occupants may move throughout the building. In such settings, the adaptive light control has to behave proactively in order to be acceptable for the user, e.g., a hallway has to be lit even before people enter it.

To simplify the discussion, we will again assume a discretization of the space \mathcal{X} . This assumption is also motivated by the idea that buildings usually have only a finite number of places of interest, such as office cubicles, cafeteria tables, hallway segments etc. We assume full observability, i.e., that our tracking system is able to exactly locate the discrete position of every building occupant. In our scenario, where the occupants are equipped with sensor badges, it is

possible to achieve relatively accurate localization using, for example, RFID tags [33]. For our application, we consider tracking as a black-box problem for which many existing techniques can be utilized.

We assume that the occupants' movements can be formalized using a probabilistic model and that people move independently of each other. The adjacency of the locations implies a graph structure \mathcal{G} with nodes \mathcal{X} and edges $\mathcal{E} \subset \mathcal{X} \times \mathcal{X}$. This means that two locations are connected by an edge if they are physically next to each other. In Fig. 1(b), the locations Φ_i and Φ_{i+1} for $1 \leq i < 4$ would hence be connected by an edge each.

Using this graph formalization, each building occupant performs a random walk on this graph, with certain transition probabilities. These random walks are naturally modeled by a Markov chain. Instead of using the locations \mathcal{X} themselves as states in the Markov chain, we perform our random walk on the edges \mathcal{E} of \mathcal{G} . For a location x , the graph will always contain the "loop" (x, x) , which corresponds to the action of remaining at a particular location. This random walk on the edges is equivalent to a second order Markov process on the locations \mathcal{X} , and allows us to capture the intuitive modeling requirement that a person is more likely to move if they moved immediately before.

Since we assume full observability, at each time step t we have a distribution $P(x_i^{(t+1)} = \cdot | x_i^{(t)}, x_i^{(t-1)})$ which we can use for predicting occupant i 's position at time $t+1$. We can easily integrate this model into our maximum expected utility framework, by computing the expected utility as

$$EU(\mathbf{a}, t, \mathbf{x}^{(t)}, \mathbf{x}^{(t-1)}) = \gamma \sum_{j=1}^m \Psi_j(a_j) + \sum_{i=1}^n \sum_x P(x_i^{(t+1)} = x_i) \mathbb{E}[\Phi_{i,x_i}(\mathbf{a}, \Theta(x_i, t))].$$

Note that this computation is tractable; the complexity only increases by a factor proportional to the number of outgoing edges at position $x_i^{(t)}$. We can also extend this simple model to handle occupants entering or exiting the building.

A more general approach would be to model the positions of the building occupants as a partially observable stochastic process. This formalization allows, for example, uncertainty about the tracking process, but is often intractable (c.f. [37]). Since we do not need extremely precise localization of the occupants in the building automation setting, the full observability assumption is often realistic.

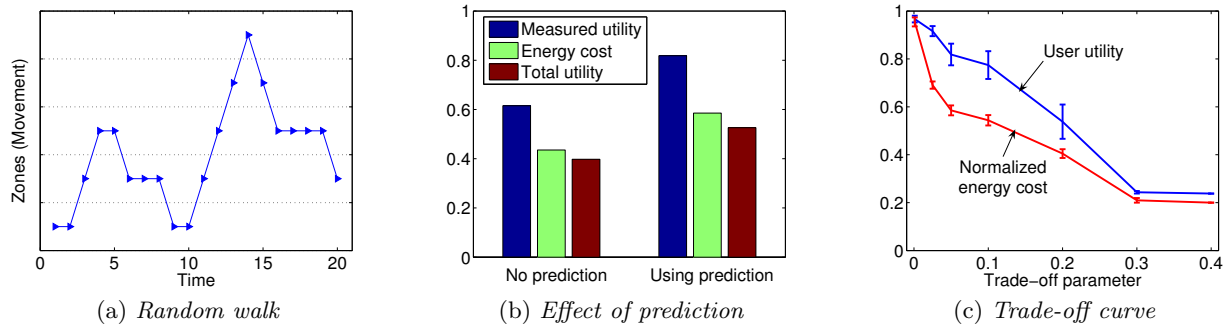


Figure 8: Results using predictive model.

5.1 Evaluation in our testbed

To evaluate our predictive light control strategy, we simulated movements of a single occupant following a random walk. The random walk was based on a second order Markov process: If the simulated occupant moved in the last transition, we assumed that the probability of continuing the movement in the same direction was 0.4; with probability 0.5 it would stop moving and with probability 0.1 it would move in the opposite direction. If the occupant did not move in the last step, the probability of remaining static was set to 0.6, and with probability 0.2 it would start to move in either direction. The discrete locations the occupant would move among were the centers of the sequence of triangular regions defined in our testbed as indicated in Fig. 1(b). We provided the transition probabilities to the coordinated illumination algorithm, and then measured the utility at the actual next position using the MICA2 motes. Our experiments consisted of a fixed 20 step random trajectory generated from our movement model.

We first compared the predictive model with the simple non-predictive version. Fig. 8(b) shows a drastic increase in mean user and total utility, along with an increase in energy consumption used for lighting up areas adjacent to the occupants current position. In spite of the increase in energy consumption, the total utility increased by more than 25%.

We also performed another experiment to explore the trade-off curve between user utility and energy consumption for varying values of γ . In interpreting these results, it has to be noted that the trade-off parameter γ is not normalized by the number of occupants, hence the values presented here cannot be directly compared to the values in Section 2. Fig. 8(c) shows a similar dynamics as observed in Fig. 3(c) in Section 2 – Increased user comfort incurs increased energy cost. It is again interesting to look for “knees” in the trade-off curve, the most prominent of which appears for $\gamma = 0.025$. For this setting, there is almost no reduction in user utility but a large decrease in cost. This experiment also indicates that the utility trade-off can be very different for a dynamic setting compared to a static situation.

During the experiments, it was also possible to observe interesting qualitative behavior of the light management strategy. For very low values of the trade-off parameter ($\gamma \approx 0.025$), the system chose to favor occupants utility by strongly illuminating all possible adjacent regions which the occupant might enter. For medium values ($\gamma \approx 0.1$), the system chose to save energy by not illuminating regions which the occu-

pant just left, since the modeled probability of an immediate return was only 0.1. For large values of $\gamma > 0.3$, the light management strategy simply chose to completely switch off all lights, in order to maximize the expected total utility.

5.2 Description of Video 3

We documented the movement experiment using our predictive model in Video 3. In this video, the lights adjust, following the trajectory of the simulated building occupant. The actual trajectory is presented for reference in Fig. 8(a).

6. CONCLUSIONS

Simultaneously optimizing user comfort and reducing energy usage in building automation is of enormous economic importance. Implementing an effective control strategy poses a complex optimization problem. In this paper, we presented a decision theoretic formalization of this optimization problem, and provided efficient algorithms for optimally trading off these two main conflicting goals. Our algorithms presented in Section 2 utilize building structure for efficient and effective coordinated illumination.

The intelligent lighting implementation presented in our paper requires the nodes to send the data to the central controller. The proposed approach though tested on a small scale with a central controller can be further extended for a large scale implementation using the distributed architecture presented in [32, 18]. The authors present a distributed architecture for solving inference problem similar to the ones formulated in this paper. The distributed architecture is able to handle the inference and control problems that is robust to unreliable communication and node failures. In this architecture, the nodes of the sensor network assemble themselves into a junction tree and exchange messages between neighbors to solve the inference problem efficiently and exactly.

Our proposed system provides a novel use of wireless sensors to capture the real state of the indoor and outdoor environment and use that information to intelligently control lights to meet user preferences. The closed-loop control strategy developed in Section 3 adapts to the presence of natural light sources and provides a significant reduction in electric light usage by harvesting daylight.

A major challenge in deploying wireless sensors in a building is caused by high maintenance cost due to quickly draining battery life. Our active sensing algorithms presented in Section 4 significantly reduce the sensing required to provide sufficient information to the control strategy, and hence

drastically increases the life of the wireless sensor network. This novel approach towards reducing the maintenance effort makes mobile, wireless sensor networks a more attractive solution even for large scale, long-term deployments.

As presented in the paper, utility theory provides a rational basis of decision making in building control. However, it also highlights various challenges which need to be addressed in further studies. The preferences of occupants as modeled by their utility functions change over time and also with the task the occupants perform. It is impractical to expect users to continuously and reliably provide their preferences to the system. Utility elicitation in complex environments is generally considered a difficult problem, and in the competitive setting of user preferences it also involves game theoretic aspects [17]. There are many elicitation techniques available which often disagree even when applied to the same person [16]. There are however promising results for using decision theoretic techniques for reliable utility elicitation [11, 12]. We expect these techniques to be a natural fit to our proposed mobile sensor network solution. In our vision, the sensor network is not only used for providing information to the building control strategy, but also to learn about the preferences of individual occupants based on their interaction with the environment. We envision a system, in which any interactions of the occupants with the building controls automatically affect their respective utility functions. The mobile sensor information can then be used to learn not only a fixed set of utility functions, but also a predictive model of how the utility functions change over time.

Building automation is an intriguing application of sensor networks. We furthermore believe that the methods proposed in this paper, as well as our proof-of-concept study, will significantly advance the applicability of sensor networks to such complex real world control problems.

Acknowledgements. This work was supported by NSF Grant No. CNS-0509383 and a gift from Intel Corporation.

7. REFERENCES

- [1] A Declarative Database of Sensor Networks. Tiny DB. <http://telegraph.cs.berkeley.edu/tinydb/>.
- [2] S. Arnborg, D. G. Corneil, and A. Proskurowski. Complexity of finding embedding in a k-tree. *SIAM J. Alg. Disc. Meth.*, 8(2):277–284, 1987.
- [3] ASHRAE, Atlanta. *ASHRAE Standards*, 90-1980a edition, 1980.
- [4] F. Bauman, A. Baughman, G. Carter, and E. Arens. A field study of pem (personal environmental module) performance in bank of americas san francisco office buildings. Technical report, University of California, Berkeley, 1997.
- [5] A. Becker and D. Geiger. A sufficiently fast algorithm for finding close to optimal junction trees. In *UAI*, 1996.
- [6] U. Bertele and F. Brioschi. *Nonserial Dynamic Programming*. Academic Press, New York, 1972.
- [7] M. Boman, P. Davidsson, and H. Younes. Artificial decision making under uncertainty in intelligent buildings. In *UAI*, 1999.
- [8] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [9] X. Boyen and D. Koller. Tractable inference for complex stochastic processes. In *UAI*, 1998.
- [10] CBEC Survey. Energy Information Administration. <http://www.eia.doe.gov/emeu/cbecs/>, 2003.
- [11] U. Chajewska and D. Koller. Utilities as random variables: Density estimation and structure discovery. In *UAI*, 2000.
- [12] U. Chajewska and D. Koller. Learning an agent's utility function by observing behavior. In *UAI*, 2001.
- [13] A. Deshpande, C. Guestrin, S. Madden, J. Hellerstein, and W. Hong. Model-driven data acquisition in sensor networks. In *VLDB*, 2004.
- [14] W. Fisk. Health and productivity gains from better indoor environments and their relationships with building energy efficiency. Technical Report 25:537-66, LBNL, 2000.
- [15] B. M. Flax. Intelligent buildings. *IEEE Communication Magazine*, April 1991.
- [16] D. Fromberg and R. Kane. Methodology for measuring health-state preferences -ii: Scaling methods. In *Journal of Clinical Epidemiology*, 1989.
- [17] D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, 1991.
- [18] C. Guestrin. *Planning Under Uncertainty in Complex Structured Environments*. PhD thesis, Stanford University, 2003.
- [19] C. Guestrin, D. Koller, and R. Parr. Multiagent planning with factored mdps. In *NIPS*, 2001.
- [20] L. Heschong, R. L. Wright, and S. Okura. Daylighting impact on retail sales performance. *Journal of Illuminating Engineering Society*, 2002.
- [21] H. Hirsh. Roomservice ai-style. In *IEEE Intelligent Systems*, volume 14, pages 8–19, 1999.
- [22] F. Hollich. *The Influence of Occular Light Perception of Metabolism on Man and Animal*. New-York: Springer Verlag, 1979.
- [23] A. Huberman and S. Clearwater. A multi-agent for controlling building environment. In *Proceedings of 1st International Conference on Multi-agent Systems*, 1995.
- [24] Intelligent Systems. Standards-Based Wireless Networking Alternatives. <http://www.sensorsmag.com/>, December 2003.
- [25] F. Jensen and F. Jensen. Optimal junction trees. In *UAI*, 1994.
- [26] U. Kjærulff. Reduction of computational complexity in bayesian networks through removal of weak dependencies. In *UAI*, 1994.
- [27] A. Krause and C. Guestrin. Optimal nonmyopic value of information in graphical models - efficient algorithms and theoretical limits. In *Proc. of IJCAI*, 2005.
- [28] J. McHugh, P. J. Burns, and D. C. Hittle. The energy impact of daylighting. *ASHRAE Journal*, 1998.
- [29] Minolta Light Meter. Minolta USA.
- [30] T. Moore, D. J. Carter, and A. I. Slater. Conflict and control: The use of locally addressable lighting in open space office plan. In *Proc. Of the Chartered Institute of Building Service Engineers*, 2000.
- [31] J. Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, New Jersey, 1947.
- [32] M. Paskin, C. Guestrin, and J. McFadden. A robust architecture for distributed inference in sensor networks. In *Proceedings of 4th IPSN*, April, 2005.
- [33] M. Philipose, K. P. Fishkin, and D. Hahnel. Mapping and localization with rfid technology. Technical report, Intel Research, 2003.
- [34] R. Evans et al. The long term costs of owning and using buildings. The Royal Academy of Engineering, 1998.
- [35] M. R. Finley Jr., A. Karakura, and R. Nbogni. Survey of intelligent building concepts. *IEEE Communication Magazine*, April 1991.
- [36] Renewable Resource Data Center. National solar radiation data base. http://rredc.nrel.gov/solar/old_data/nsrdb/.
- [37] J. Rintanen. Complexity of planning with partial observability. In *ICAPS*, 2004.
- [38] U. Rutishauser, J. Joller, and R. Douglas. Control and learning of ambience by intelligent building. In *IEEE transactions on Systems, Man and Cybernetics - Part A: Systems and Humans*, volume 35, pages 121–132, Jan 2005.
- [39] S. Sharples, V. Callaghan, and G. Clarke. A multi-agent architecture for intelligent building sensing and control. *International Sensor Review Journal*, 1999.
- [40] V. Singhvi, M. Bigrigg, H. S. Matthews, and J. Garrett. Continuous commissioning using embedded sensors. In *Proc. of ICCCE*, 2005.
- [41] H. Varian. *Microeconomic Analysis*. ISBN 0-393-95735-7.
- [42] Wireless Sensor Platform. MICA 2. <http://www.xbow.com/>.
- [43] J. Woods. Cost avoidance and productivity in owning and operating buildings. *OM*, 4(4):753–70, 1989.
- [44] D. Wyon. Indoor environmental effects on productivity. In *Proc. of Indoor Air*, 1992.
- [45] X10 Building Automation Protocol. X10.