A Power-Efficient Self-Calibrating Smart Lighting System

Yerbol Aussat, Ansis Rosmanis, and Srinivasan Keshav
University of Waterloo

Abstract

Lighting load accounts for a significant portion of overall energy consumption in office buildings. To reduce this load, we have designed and built a smart self-calibrating lighting control system that minimizes power consumption that automatically responds to changes in daylight and occupancy, while simultaneously providing personalized lighting comfort to each occupant. The system measures illuminance and occupancy from sensors located at each work station. Using an unobtrusive self-calibration process, it estimates the relationship between the dimming level of each bulb and the illuminance at each work station. Subsequently, an adaptive control algorithm maintains the desired illuminance at work surfaces despite environmental fluctuations by periodically recalculating the power-efficient and comfort-preserving dimming level for each bulb. Based on a realistic deployment of our system, we find that our system quickly responds to changes in occupancy, daylight and user preferences. We also show, through extensive simulations using 7 months of collected daylight and occupancy data, that our system reduces energy consumption by about 40% compared to conventional LED lighting systems.

Keywords: Building lighting, Efficiency, Control

1. Introduction

Artificial lighting accounts for about 17% of overall electrical load in commercial office buildings in the United States [1]. It is important to reduce this load and its associated carbon footprint, ideally without compromising the comfort of building occupants. The primary approach to reducing lighting load is to replace incandescent and compact fluorescent bulbs[1] with energy-efficient LED bulbs [2]. Recently, commercially-available bulbs such as the Philips Hue- and the Sylvania SMART+ allow the luminous output of individual LED bulbs to be controlled using software. This allows a further reduction in lighting load by adapting the level of illumination to occupancy changes and availability of daylight (also known as *daylight harvesting*), and is the focus of our work.

Lighting systems capable of occupancy detection and daylight harvesting have been studied over the past decade or so [3, 4, 5, 6, 7, 8]. Although these systems demonstrate that a reduction in energy of up to 61% is feasible [8], they have not been widely deployed for a variety of reasons, including difficulty of installation, the need for manual calibration and re-calibration, a high up-front cost, and the complexity of integration with other building systems [9, 10]. Moreover,

[1] One or more bulbs in a single lighting fixture is termed a 'luminaire' in the literature. Our work deals with luminaires with single bulbs, so we use the term 'bulb' exclusively in the sequel.

despite a study by Newsham et al. [11] that suggests that the illuminance on a work surface is the main determinant of occupant comfort, most existing systems measure illuminance not at a work surface, but at the lighting fixtures [3, 8, 5, 12] or on walls [13], deducing the illuminance on the desktop. This computation is necessarily flawed, reducing user comfort [14, 15, 16].

We present a power-efficient smart lighting control system that is both self-calibrating and easy to deploy. We deploy low-cost light and occupancy sensors adjacent to each work station and allow a personalized lighting level at each station to be chosen by a user (or the set of users sharing a common space). We then periodically compute a *nearly optimal* dimming level of each bulb using a linear program, whose objective is to minimize power consumption subject to user comfort requirements. To account for modeling and calibration errors, we use a feedback control algorithm that converges dimming levels despite these errors.

We have implemented a prototype of our system and evaluate it in a realistic setting (please see Figure 1 for a schematic of the system, and Figures 5 and 6 for images of the system we deployed and evaluated). We demonstrate that it reacts to changes in daylight conditions in under 2 seconds, and to changes in occupancy in about 350 milliseconds. Extensive simulations suggest that, in our experimental setting, it utilizes 40% less power compared to a conventional lighting system[2].

This paper represents four years of work in design, analysis, implementation, and performance tuning of our system. Our main contributions are:

- We have designed a power-efficient lighting system that exploits daylight and occupancy information to minimize energy consumption while satisfying the individual lighting preferences of all occupants.
- We have implemented the system with off-the-shelf components and deployed it in a realistic environment.
- We have evaluated the performance of our system using both laboratory experiments and simulations based on 7 months of collected daylight and occupancy data. These demonstrate the system's functionality, responsiveness, and ability to reduce energy consumption compared to existing systems.

The rest of this paper is structured as follows. In Sections 3 and 4 we formulate a mathematical bulb model and discuss the design of our smart lighting system. System implementation is outlined in Section 5. In Section 6 we evaluate the system's performance. Section 2 presents an overview of prior work and Section 7 concludes.

2. Related Work

2.1. Surveys on lighting control systems

Several recent papers have surveyed work in lighting control systems in office buildings [17, 10, 18, 19, 9, 20, 21]. These surveys describe systems that incorporate fixture-based light sensors [3, 22, 23, 12, 24, 25, 15, 26] and closed-loop control systems [15, 14, 25, 12, 27, 28]. Given the existence of multiple recent surveys of this area, in the remainder of this section we

[2]Throughout this paper, "conventional lighting systems" refers to LED-based systems that do not have occupancy and illuminance information and cannot control dimming levels of individual bulbs.

present only the most-closely related work in that it uses wireless illuminance sensors near work surfaces without associating bulbs one-to-one with work areas. The advantage of this design choice is that, due to the explicit knowledge of illuminance near the work surface, such systems are usually capable of finding optimum—or nearly-optimum—dimming levels, achieving target illuminances at all work surfaces. However, as we discuss, prior work makes some assumptions that precludes practical deployment.

2.2. Closely-related work

Caicedo *et al.* [4] placed additional wireless light sensors at each work surface to provide periodic feedback to bulb-based sensors. However, their system responds to changes in environmental illuminance very slowly (in around 100 seconds) which is sub-optimal. Moreover, because light sensors placed on a work surface are subject to occlusion, this method does not work as well as a design that places light sensors *just above* the work surface, as we do.

Borile *et al.* [29, 5] proposed a data-driven approach for determining the linear mapping between measurement points on the ceiling and points of interest at the work surface. This approach requires collecting sensor measurements from both work plane-based and bulb-based sensors during the daytime when the office is not occupied. Then, the collected data is used to learn the daylight mapping. One practical limitation of the proposed method is that system recalibration is slow, requiring the collection of a new training data set. Also, the proposed method does not guarantee good performance under different weather conditions and varying amounts of daylight.

Miki *et al.* [30] present a distributed lighting control strategy that utilizes infra-red communication between neighbouring bulbs. This is not supported by commercially-available bulbs today.

Similar to our work, Wen and Agogino [31] propose an energy-efficient linear optimization-based lighting control. However, to generate an illuminance model and determine artificial light distribution in the office, the authors use the RADIANCE [32] image rendering program, which is based on backward ray tracing. This requires explicit knowledge of several office parameters, such as office dimensions, internal surface reflectance, locations and geometries of furniture and other objects, as well as bulb parameters and location, precluding practical deployment.

Yeh *et al.* [33] and Pan *et al.* [34] present a system for daylight harvesting, assuming that locations of office occupants are known and that they carry wireless light sensors on their mobile phones. Lighting control strategies based on linear programming and sequential quadratic programming algorithms were proposed to satisfy individual illumination requirements of the users, based on their activities. Again, this approach suffers from potential occlusion of sensors. Moreover, office occupants may not wish to have sensing software installed on their personal devices.

Ravi *et al.*![35] a surveillance camera infrastructure as the sole sensing substrate to control smart lighting power levels. They show that their system can sense per-desk lighting levels accurately. However, the deployment of cameras in workplaces has privacy implications that preclude widespread deployment.

In summary, we are not aware of an alternative system that meets our design criteria of power-optimality, personalized lighting comfort, robustness to estimation errors, fast response time, and plug-and-play deployment. Moreover, only a handful of other systems have been implemented in a realistic setting, perhaps due to the complexity in the design, calibration, and installation of daylight-linked control systems [9].

2.3. Occupant Preferences for Lighting

Newsham *et al.* [11] investigated how well various metrics correlate with occupant satisfaction with office lighting. They observed that illuminance *measured on the work surface* was the best predictor of whether participants were satisfied with their lighting level. Illuminance measured at the ceiling was a substantially worse predictor. Recently, luminance-based metrics [36] have been proposed as better indicators of lighting comfort than horizontal illuminance. Although this line of work indicates the drawbacks of horizontal illuminance, we have stayed with the simpler metric because there does not appear to be expert consensus on the best luminance-based metric.

In other work, Lashina *et al.* and Newsham *et al.* studied occupants in an open-plan office laboratory [37, 38, 39] and demonstrated that occupants whose preferred light levels were met had significant improvements in mood, productivity, and comfort. Galasiu *et al.* [40] found that user acceptance becomes higher when users are provided with at least partial control of their lighting system. They also found that occupants strongly prefer daylight to artificial lighting.

To sum up, a well-designed office lighting system measures illuminance on (or near) work surfaces, allows users to control their lighting conditions, and works synergistically with natural lighting. These results inform our design.

3. System Design

3.1. Problem Formulation

Consider an office that has N work stations, such as the one illustrated in Fig. 1. We assume that it is lit both by daylight and M controllable LED bulbs, and we only have control over the latter.[3] The office can be multi-occupancy or single-occupancy with several work stations belonging to the same person. In addition, occupants can adjust illumination levels at their work stations, allowing personalization. The sensors installed at each work station communicate occupancy status and illuminance level to the central controller, which, in turn, determines nearly optimal dimming levels for all individual bulbs, and sends them control signals.

The goal of a smart lighting system is to provide the desired level of illuminance at each of the work stations while minimizing energy consumption by fully exploiting the available daylight and avoiding any unnecessary over-illumination of work stations. The system should quickly respond to changes in daylight levels and occupancy. Finally, it should be plug-and-play and self-calibrating so that it can be easily deployed without manual calibration, regardless of the geometry and configuration of the room.

3.2. Mathematical Model

This section develops a mathematical model of PAR38 Philips Hue LED bulbs [42] used in our study. We chose this because it is widely used and provides a software API for dimming control. Although our analysis is specific to this choice, a similar approach can be used to model any bulb.

Let $A(t) = (A_{ij}(t))$ be the *illuminance gains matrix*, where each of its elements $A_{ij}(t)$ is the illuminance gain of sensor i from a fully-lit bulb j at time t. $A_{ij}(t)$ is time-dependent because

[3] Although we are aware of some systems that control daylight level using motorized blinds [41, 7, 8], this paper does not address this.

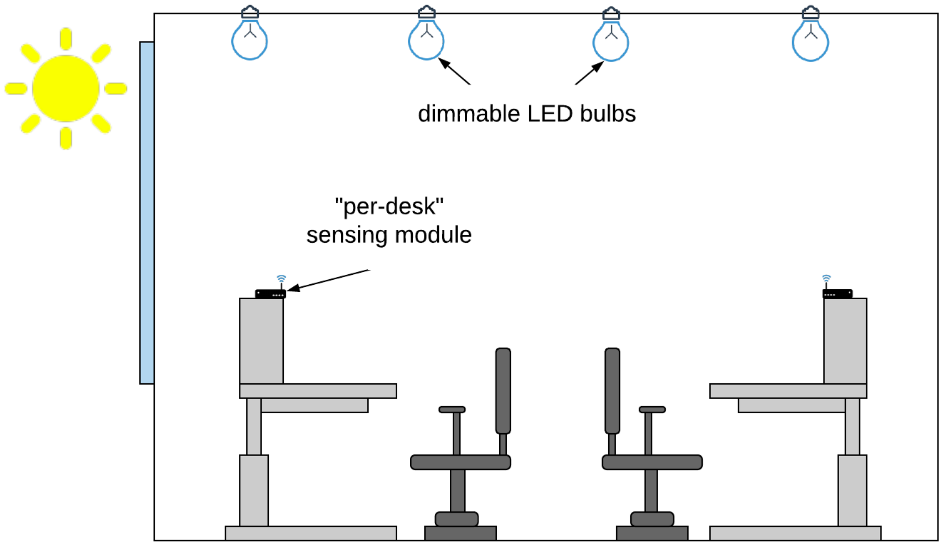

Figure 1: A typical shared office.

it can be affected by objects or people between sensor i and bulb j, slight accidental movements of sensors or bulbs as well as changes in the bulbs' temperatures [43]. Let $d_j \in [0, 1.0]$ be a *dimming level* of the jth bulb: if $L_{ij}(d_j, t)$ is the illuminance gain of sensor i from a bulb j whose dimming level is d_j at time t, then

$$d_j(t) = \frac{L_{ij}(d_j, t)}{L_{ij}(1.0, t)} = \frac{L_{ij}(d_j, t)}{A_{ij}(t)} \tag{1}$$

The Philips Hue API [44] does not allow us to control dimming levels directly. Instead, it only allows turning a bulb j on and off, as well as setting its *brightness control value* b_j to an integer between 0 and 255. We determined the mapping from b_j to d_j empirically, as shown in Fig. 2. An analytic relationship was determined by fitting a curve to the experimental data points as:

$$d_j(b_j) = \begin{cases} 2.55 \cdot 10^{-5} \cdot b_j^{1.90} + 0.047 & \text{if j is on and } 0 \leq b \leq 255 \\ 0 & \text{if bulb j is off} \end{cases} \tag{2}$$

Note that when the brightness control value is 0, a bulb's dimming level is 0.047 (4.7%). The dimming level becomes 0 only when a bulb is completely turned off.

Power consumption model: We experimentally studied the relationship between dimming level and power for a PAR38 Philips Hue bulb. It is nearly linear when the bulb is on, but with a clear discontinuity when the bulb is off (see the blue dashed line in Fig. 3). This is because the bulb has a standby power draw of 1.18 W due to its use of the Zigbee communication protocol. The power vs. dimming relationship can be represented using the best fit line $9.97d_j + 2.47$ as:

$$P_j(d_j) = \begin{cases} 1.18 & \text{if } d_j = 0.0 \\ 9.97d_j + 2.47 & \text{if } 0.0 < d_j \leq 1.0 \end{cases} \tag{3}$$

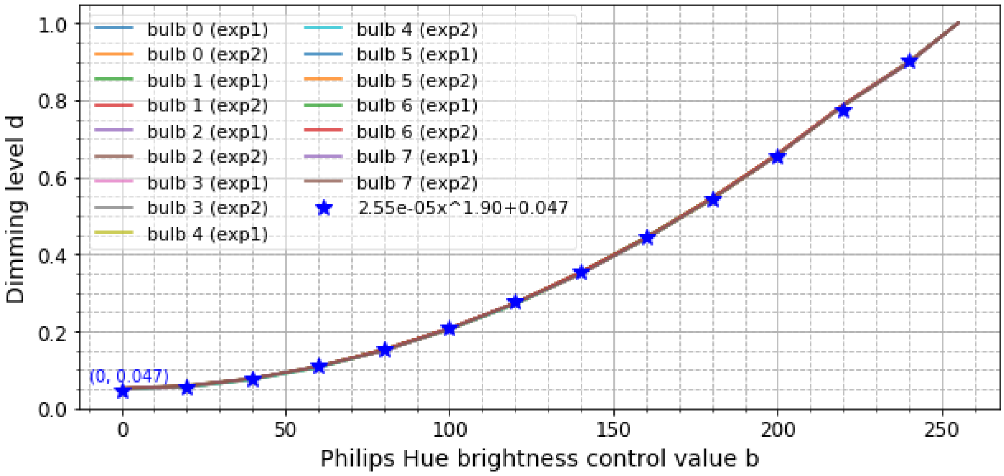

Figure 2: Empirically obtained relationship between dimming level d_j and brightness control value b_j of a PAR38 Philips Hue bulb, and the corresponding best-fit curve.

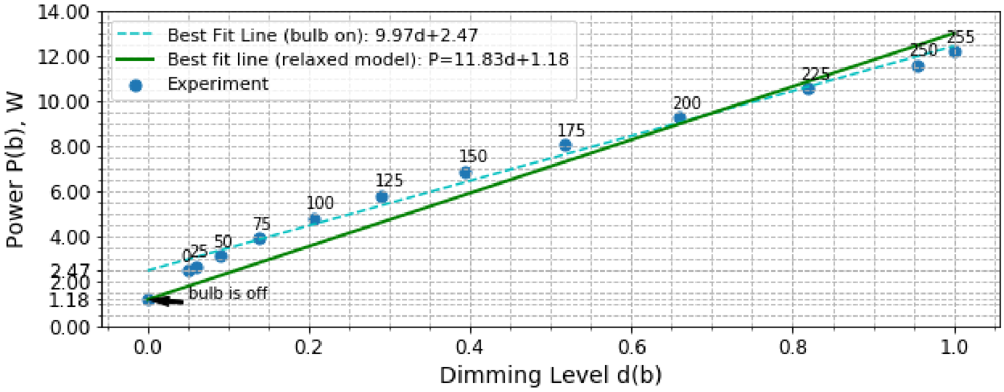

Figure 3: Empirically obtained relationship between power $P_j(b_j)$ and dimming level $d_j(b_j)$ of a PAR38 Philips Hue bulb. Each experimental data point is labeled with a corresponding Philips Hue brightness control value b_j.

Unfortunately, in the power optimization step, which is discussed in Section 4, the discontinuity at zero would result in a mixed-integer optimization program, which is computationally expensive to solve. Hence, we relax this relationship to :

$$P_j(d_j) = 11.83d_j + 1.18 \tag{4}$$

represented by the solid green line in Fig. 3. Note that the effect of this linearization is slight over-estimation (up to $0.75W$) of power consumption for high dimming levels, and under-estimation (up to 1.25 W) at lower dimming levels.

The electrical power consumption of LED bulbs depends slightly on their temperature [43]. We find, for example, that after 1 hour of operation, the bulbs consume about 3% less power to provide the same illuminance. For simplicity, we use a steady-state estimate of the power vs. dimming level relationship to model the bulbs. This can introduce an error of up to about 3% in our results.

3.3. Determining the Illuminance Gains Matrix

We now discuss how to determine the illuminance gains matrix $A(t)$, whose elements $A_{ij}(t)$ represent the illuminance from bulb j on sensor i at time t. Throughout this paper, we refer to the process of determining matrix $A(t)$ as the *calibration process*. To begin with, let $R_i(\vec{d}, t)$ be the illuminance at sensor i at time t when the bulb dimming levels are represented by the dimming vector $\vec{d} = [d_1, ..., d_M]$, where M is the total number of bulbs. From the additivity of light [15, 3]:

$$R_i(\vec{d}, t) = E_i(t) + \sum_{j=1}^{M} L_{ij}(d_j, t) = E_i(t) + \sum_{j=1}^{M} d_j A_{ij}(t) \tag{5}$$

where $E_i(t)$ is the time-dependent illuminance gain of sensor i from the environment, which comprises of daylight and other environmental light sources out of our control.

One straightforward way to obtain $A(t)$ is to first measure environmental illuminance gains by reading sensors while all bulbs are off. Next, we could sequentially turn on one bulb at a time, record new illuminance readings, and subtract respective environmental illuminance gains from these readings. Even though this calibration process allows us to estimate the matrix $A(t)$, it is obtrusive, and thus cannot be performed when the system is in use. However, re-calibration is necessary whenever a significant change in the illuminance gains matrix occurs, such as when furniture is moved.[4] To address this, we developed an unobtrusive calibration method, based on the observation that, while a human eye is insensitive to minor lighting changes [11], photosensors are capable of detecting them accurately.

Suppose that, just before calibration, the dimming vector is $\vec{d}$. Let $\vec{\hat{d}}^{(j)}$ be a dimming vector with all the entries equal to $\vec{d}$ except:

$$\hat{d}_j^{(j)} = \begin{cases} d_j + S & \text{if } d_j < B \\ d_j - S & \text{if } d_j \geq B \end{cases} \tag{6}$$

[4]Re-calibration is not necessary if the illuminance gains matrix does not change significantly because the feedback control algorithm described in Section 4.2 ensures that comfort is maintained despite smaller errors in the matrix. In particular, re-calibration is not needed if there is a change in daylighting level.

where, in our experiment, we use a *dimming step* $S = 0.1$ and the pivot point $B = 0.65$.

We first record illuminance readings $R_i(\vec{d}, t)$. Immediately after, sequentially for every bulb j we change the dimming level of bulb j to $\hat{d}_j^{(j)}$, record illuminance readings $R_i(\vec{\hat{d}}^{(j)}, t')$ for all sensors i, and then restore the brightness of the bulb to its original value d_j, and proceed to the next bulb. We ensure that t and t' are at most a few seconds apart, and we assume that neither daylight nor the illuminance gains matrix change much on such a timescale. Hence, considering that $A_{ij}(t) \approx A_{ij}(t')$ and $E_i(t) \approx E_i(t')$ and using Eq. 5 we get

$$A_{ij}(t) = \frac{|R_i(\vec{d}^{(j)}, t') - R_i(\vec{d}, t)|}{S} \tag{7}$$

for all sensors i and bulbs j. Note that this calibration procedure allows estimating the illuminance gains matrix without requiring explicit knowledge of office geometry and locations of bulbs and photosensors, thereby contributing to the system's plug-and-play design.

In our prototype implementation (Section 6), in an office with 8 bulbs, this calibration process takes about 10s. Depending on the number of bulbs and the nature of the environment, we suggest that a re-calibration period of 10min-1hr.

3.4. Estimating Environmental Illuminance E

Because we obtain columns of the illuminance gains matrix one after another in quick succession, we assume that its values, as well as environmental illuminance gains, do not change drastically during this procedure. Thus, knowing total illuminance values $R_i(\vec{d}, t)$ on all sensors i, dimming level settings d_j on all bulbs j, and illuminance gains $A_{ij}(t)$, we can estimate the environmental illuminance gains $E_i(t)$ from Eq. 5 as:

$$\vec{E}(t) = \vec{R}(\vec{d}, t) - A(t)\vec{d} \tag{8}$$

3.5. Estimation Error for A and E

$A_{ij}(t)$ is a time-varying ground-truth illuminance contribution on sensor i from bulb j. However, the proposed calibration process gives us only an empirical estimate of a snapshot of matrix A at the calibration time, which we denote as $\widetilde{A} \in \mathbb{R}^{N\times M}$. $\widetilde{A}$ is not a function of time and is subject to estimation errors.

Let $\epsilon(t) \in \mathbb{R}^{N\times M}$ be the time-dependent *estimation error* that captures both calibration errors, caused by imperfect sensor measurements and environmental fluctuations, and estimation errors, caused by the time-varying nature of $A(t)$. It is given by:

$$\epsilon(t) = \widetilde{A} - A(t) \tag{9}$$

Note that Eq. 8 uses the true matrix $A(t)$ to estimate the environmental illuminance contribution on sensors. However, in practice, only the estimated matrix $\widetilde{A}$ is available through the calibration process. By combining Eq.'s 8 and 9, we get

$$\vec{R}(\vec{d}, t) = \vec{E}(t) + (\widetilde{A} - \epsilon(t))\vec{d} \tag{10}$$

Since the error term $\epsilon(t)$ in Eq.10 is unknown, the environmental contribution $\vec{E}(t)$ cannot be calculated directly. We can only estimate it as:

$$\vec{\widetilde{E}}(t) = \vec{E}(t) - \epsilon(t)\vec{d} = \vec{R}(\vec{d}, t) - \widetilde{A}\vec{d} \tag{11}$$

where $\vec{\tilde{\tilde{E}}}(t)$ is an estimate of the environmental illumination at time t, and $-\epsilon(t)\vec{d}$ is the associated estimation error in the environmental illumination. The effect of the estimation error on the system performance is further discussed in Sections 4.2 and 6.5.

4. Optimal and Adaptive Control

4.1. Optimization Program

Let the target illuminance on sensor i be h_i, where this target can be appropriately set according to user preferences and occupancy status of the corresponding workspace. We assume that the installed light capacity is such that, with the maximum power level at all bulbs, this target can be met; otherwise the program is trivially infeasible. The goal of the system is to minimize the total power consumed by M bulbs while providing (at least) the target illuminance levels at N work stations. Recall that the (relaxed) relationship between power and dimming level is linear (Eq. 4). Therefore, the optimization program reduces to minimizing the sum of dimming levels subject to illumination requirements.

Let $\mathcal{D} = \sum_{j=1}^{M} d_j$ denote the sum of the components of the dimming vector. Then, the optimization program is:

$$\begin{aligned} \underset{\vec{d}}{\text{minimize}} \quad & \mathcal{D} \\ \text{subject to} \quad & \vec{E}(t) + A(t)\vec{d} \geq \vec{h} \\ & \vec{0} \leq \vec{d} \leq \vec{1} \end{aligned} \tag{12}$$

To solve this linear program, the knowledge of $A(t)$ and $\vec{E}(t)$ is required. If we knew the real $A(t)$ and $\vec{E}(t)$ at every moment t, optimal dimming levels could be chosen at each optimization step. However, in reality, we can only estimate $A(t)$ and $\vec{E}(t)$, and these estimates are prone to error. Hence, to design a practical system, we re-express the optimization program in terms of $\widetilde{A}$, obtained from the calibration process and defined in Eq. 9, and the environmental illumination estimate $\vec{\tilde{\tilde{E}}}(t)$, defined in Eq. 11. We then use an iterative process, discussed next, to achieve the control objective despite estimation errors.

Let index k denote a variable's value at the moment of iteration k. The optimization program can then be written as

$$\begin{aligned} \underset{\vec{d_k}}{\text{minimize}} \quad & \mathcal{D}_k \\ \text{subject to} \quad & \vec{\tilde{\tilde{E}}}_k + \widetilde{A}\vec{d_k} \geq \vec{h} \\ & \vec{0} \leq \vec{d_k} \leq \vec{1} \end{aligned} \tag{13}$$

This linear program cannot be solved directly because the environmental illumination estimate $\vec{\tilde{\tilde{E}}}_k$ is unknown before iteration k. Recall from Eq. 11 that

$$\vec{\tilde{\tilde{E}}}_k = \vec{E}_k - \epsilon_k\vec{d_k} = \vec{R}_k - \widetilde{A}\vec{d_k} \tag{14}$$

Both terms on the right-hand side of Eq. 14 are unavailable to the optimizer since $\vec{d_k}$ is the unknown variable that we want to optimize and $\vec{R}_k$ can be measured only after dimming levels

$\vec{d}_k$ have been set on bulbs. We approximate the environmental illumination estimate $\tilde{\vec{E}}_k$ by $\tilde{\vec{E}}_{k-1} = \vec{R}_{k-1} - \widetilde{A}\vec{d}_{k-1}$, which can be readily calculated. Then we can rewrite the program 13 as

$$\begin{aligned} \underset{\vec{d}_k}{\text{minimize}} \quad & \mathcal{D}_k \\ \text{subject to} \quad & \widetilde{A} \cdot (\vec{d}_k - \vec{d}_{k-1}) + \vec{R}_{k-1} \geq \vec{h} \\ & \vec{0} \leq \vec{d}_k \leq \vec{1} \end{aligned} \tag{15}$$

Now, all of the terms, except for the unknown dimming vector $\vec{d}_k$, are readily available.

4.2. Adaptive Control

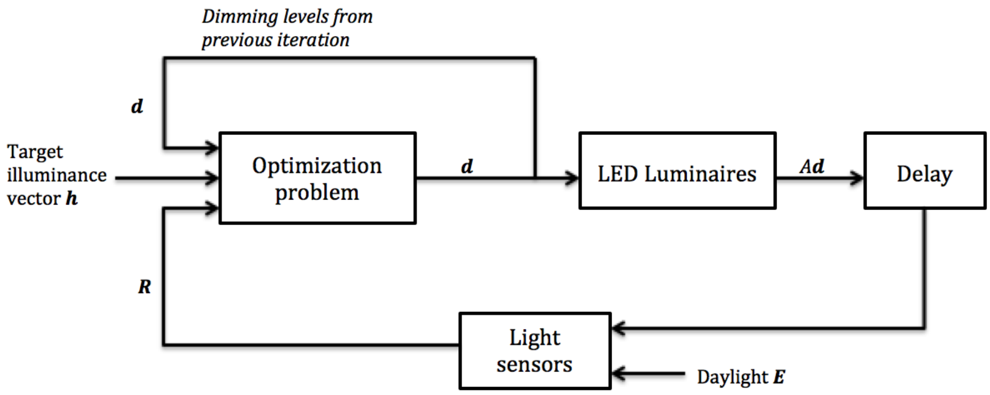

Figure 4: Control diagram for the smart lighting control.

In order to compute the optimal dimming levels in iteration k, our optimization program given in (15) uses the results of the previous iteration $k - 1$. This feedback loop forces the illuminance on the light sensors to converge to target set points, making the system robust to the model imperfections. Fig. 4 shows the corresponding control diagram. At each iteration k the controller repeatedly executes the following steps:

1. Using target illumination $\vec{h}$, the previous iteration's dimming vector $\vec{d}_{k-1}$ and illuminance readings $\vec{R}_{k-1}$, find $\vec{d}_k$ by solving the optimization program given in (15).

2. Set the computed dimming levels $\vec{d}_k$ on bulbs, and wait until the bulbs fully adapt to the new dimming levels.

3. Measure new illuminance $\vec{R}_k$ on the light sensors.

We now analyze the convergence of this approach. For simplicity, assume that the environmental illuminance gains and the estimation error matrix do not change significantly between two iterations of the control algorithm, i.e., $\epsilon_k = \epsilon_{k-1} = \epsilon$ and $\vec{E}_k = \vec{E}_{k-1} = \vec{E}$. In addition, assume that the solution to the optimization program (15) satisfies the illuminance constraints with equality, so that when in iteration k the controller computes the dimming vector $\vec{d}_k$, we get

$$\widetilde{A} \cdot (\vec{d}_k - \vec{d}_{k-1}) + \vec{R}_{k-1} = \vec{h}$$

or, considering Eq. 14,

$$\widetilde{A}\vec{d}_k - \epsilon\vec{d}_{k-1} + \vec{E} = \vec{h} \tag{16}$$

Then, after the bulbs fully adapt to the new dimming levels $\vec{d}_k$, sensors read new illuminance measurements $\vec{R}_k$:

$$A_k\vec{d}_k + \vec{E} = \vec{R}_k$$

where A_k is a true theoretical illuminance gains matrix. We can express A_k in terms of the empirically obtained matrix $\widetilde{A}$ using Eq. 9:

$$\widetilde{A}\vec{d}_k - \epsilon\vec{d}_k + \vec{E} = \vec{R}_k \tag{17}$$

Similarly to Eq. 16, in the next iteration:

$$\widetilde{A}\vec{d}_{k+1} - \epsilon\vec{d}_k + \vec{E} = \vec{h} \tag{18}$$

By subtracting Eq. 17 from Eq. 16 we get:

$$\epsilon(\vec{d}_k - \vec{d}_{k-1}) = \vec{h} - \vec{R}_k \tag{19}$$

On the other hand, subtracting Eq. 17 from Eq. 18 gives

$$\widetilde{A}(\vec{d}_{k+1} - \vec{d}_k) = \vec{h} - \vec{R}_k \tag{20}$$

Note that, since Eq.s 19 and 20 have the same right-hand side, they can be combined as:

$$\widetilde{A}(\vec{d}_{k+1} - \vec{d}_k) = \epsilon(\vec{d}_k - \vec{d}_{k-1}) \tag{21}$$

By introducing $\Delta\vec{d}_k = \vec{d}_k - \vec{d}_{k-1}$, we rewrite Eq. 21 as:

$$\widetilde{A}\Delta\vec{d}_{k+1} = \epsilon\Delta\vec{d}_k \tag{22}$$

and, therefore

$$\Delta\vec{d}_{k+1} = \widetilde{A}^{-1}\epsilon\Delta\vec{d}_k \tag{23}$$

where $\widetilde{A}^{-1}$ denotes the pseudo-inverse. Eq. 23 indicates the convergence condition of the control system. That is, whether the system converges depends on the product of matrices $\widetilde{A}^{-1}$ and ϵ. Intuitively, if the elements of the error matrix ϵ are smaller than the elements of the illuminance gains matrix $\widetilde{A}$, then the system converges. We evaluate the convergence of the system for various levels of estimation error in Section 6.5

Adaptation to changes in $\vec{E}$ and $\vec{h}$: Note that changes in environmental illumination $\vec{E}$ (e.g., daylight) are handled directly by the control loop in the following control epoch. After $\vec{E}$ changes, sensor readings $\vec{R}_k$ start to deviate from the target $\vec{h}$, and the optimizer computes new dimming levels to restore the target illuminance on the sensors. On the other hand, we designed the system to handle changes in target illuminance $\vec{h}$ (e.g., occupancy or users' illuminance preferences) differently for time-efficiency reasons. When $\vec{h}$ changes, we immediately terminate the ongoing control iteration and start the next one with the updated vector $\vec{h}$.

5. Implementation

As a proof-of-concept, we have implemented a prototype of the proposed smart lighting system. It consists of three principal components: stand-alone wireless sensing modules that measure occupancy and illuminance of each work station, dimmable LED bulbs, and a central controller that receives sensor inputs and sends control signals to each bulb. A high-level diagram of the system is shown in Fig. 5. For reasons of space, the details are elided; a full description can be found in the extended version of this paper [45].

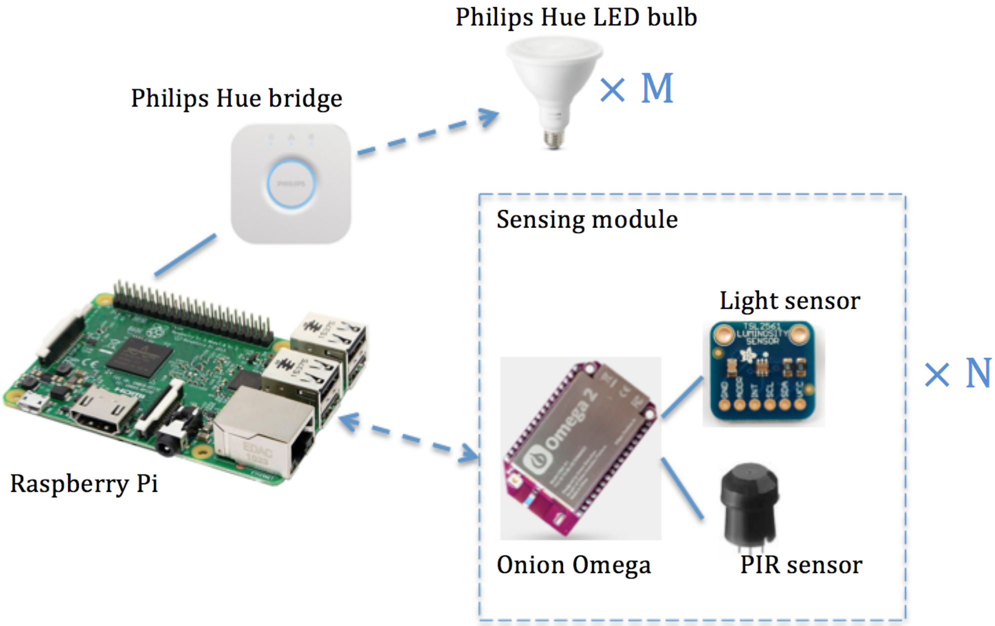

Figure 5: High-level diagram of the smart lighting system.

6. System Performance

This section investigates the performance of the smart lighting system, and in particular, how quickly and accurately it converges to the target illuminance levels. A short video that demonstrates the system can be found at https://youtu.be/G8RIXDEUX20.

6.1. Evaluation Testbed

We installed a 2.45 m. high ceiling above four desks in a secluded 3.0×3.8 m. work area (Fig. 6). The space is illuminated by eight dimmable 1300 lumen-rated LED bulbs evenly installed in the ceiling, that wirelessly communicate with a central control module. Each desk has a wireless sensing module deployed on the edge of its shelf, which periodically sends occupancy status and illuminance level to the control module via a local wireless network. Also, each occupant can express individual illuminance preference at their desk through that desk's sensing module. Users' illuminance preference, along with sensor readings, serve as inputs to the controller which implements our control algorithm described in Section 4.2.

Unfortunately, our testbed is based in an internal room that does not have a window. Therefore, it is not possible to carry out daylight harvesting. Instead, we emulate daylight by switching the laboratory's central light on and off. In other experiments, not described here, we also used an intense work lamp to simulate daylight to some extent. We realize that this does not capture fast time-scale variations in daylight. Nevertheless, please note that in the mathematical system model, the external artificial light and daylight are equivalent, and both of them are modeled as environmental illuminance gains $\vec{E}(t)$. Thus, switching the artificial lighting on and off is similar to quickly opening and closing blinds or the sun passing behind or emerging from a cloud.

System calibration: By using our automated calibration process, the following illuminance gains matrix $\widetilde{A} \in \mathbb{R}^{4\times 8}$ is obtained for our evaluation testbed with 4 light sensors and 8 bulbs:

	b0	b1	b2	b3	b4	b5	b6	b7
s0	176.19	5.46	13.66	329.15	2.73	6.83	25.95	13.66
s1	5.40	195.84	12.16	1.35	430.85	21.61	4.05	9.45
s2	198.54	2.70	6.75	14.86	4.05	5.40	301.19	17.56
s3	5.87	199.69	5.87	2.94	13.21	431.68	5.87	19.09

Note that the illuminance gains on all sensors from bulbs $b2$ and $b7$ are relatively small, due to these bulbs being located further away from the sensing modules. If they are removed from the system, the resulting illuminance gains matrix is 4×6 and achieves a maximum achievable illuminance (i.e., row sum) of:

s0	s1	s2	s3
546.31	659.11	526.75	659.27

This shows that six 1300 lumen-rated bulbs are sufficient to illuminate the office space (delivering at least 500 lux to all work stations), even when no environmental lighting is available. Therefore, in the evaluation of the reduction in energy consumption, presented in Section 6.6, we consider systems with six bulbs.

6.2. Changes in Environmental Illuminance

To evaluate the smart lighting system's responsiveness to changes in environmental illuminance, we set the system to maintain heterogeneous illuminance levels, namely, 300, 350, 450 and 500 lux, on the four sensing modules. Then, by turning on and off the laboratory's central lighting, we simulate the opening and closing of blinds. On Fig. 7, white and yellow regions correspond to the laboratory's main light being off and on, respectively. The top time series show the illuminance signals on the four sensing modules. The bottom time series show dimming levels on the bulbs with small filled circles corresponding to moments when the dimming levels are set.[5]

Abrupt spikes on the illuminance time series correspond to abrupt changes in environmental illuminance, i.e., turning on or turning off the laboratory's central light. We see that it takes the system about 2-4 seconds to fully adapt to the environmental changes and restore illuminance levels on all sensors. Recall that the control of bulbs' dimming levels is done via the iterative control algorithm from Section 4.2. Therefore, the system's timescales can be expressed in terms of the number of iterations. Based on our empirical measurements, the maximum time of one iteration is $\sim$1.65s. The bulk of this time is due to the need for the LED bulbs to fully adapt to

[5] Bulbs $b2$ and $b7$ were off throughout the experiment, so are not shown in this and subsequent figures.

Figure 6: Testbed implementation of the smart lighting system.

the new dimming level. Our measurements indicate that this typically requires less than 700ms, but, for additional robustness, we allocate 1.3s for this adaption. The linear program solver and network calls take most of the remaining 350ms.

From the dimming level time series (the bottom graph in Fig. 7) we see that the time it takes to set new dimming levels on bulbs after the environmental change occurs is within $\sim$2 seconds of the change in external lighting, and that only 1 iteration is required for the system to fully converge if we have an accurate estimated matrix $\widetilde{A}$. If the matrix $\widetilde{A}$ is poorly estimated, it would take more iterations for the system to converge, as discussed in Section 6.5.

It is worth noting that while this experiment considers significant abrupt changes in environmental illuminance, natural changes in daylight are much smoother and subtler. When the environmental lighting changes per iteration are lower than typical eye sensitivity, we found that the system adapts to these changes imperceptibly.

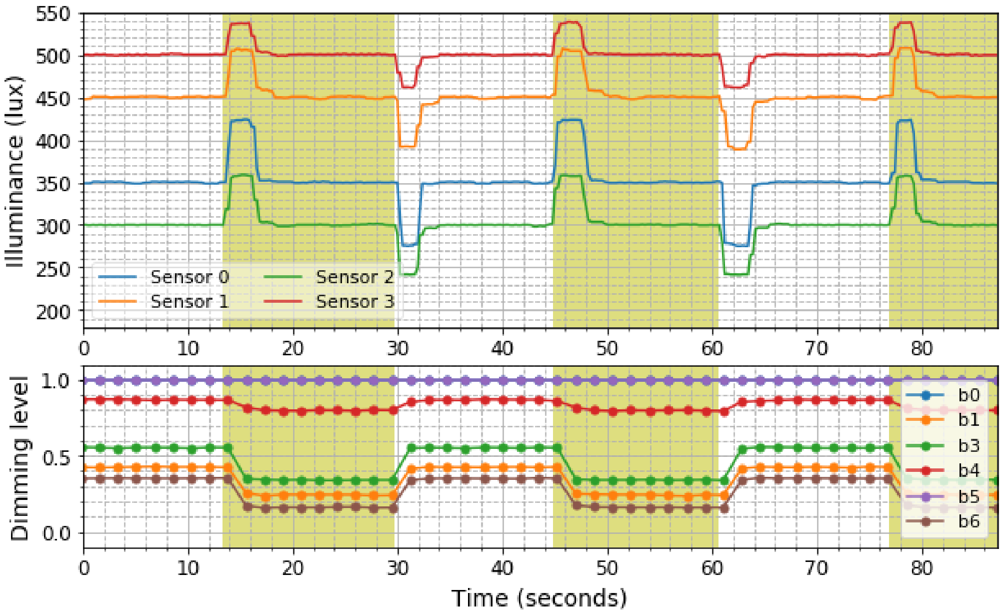

Figure 7: Response to changes in environmental illuminance. $b0$ overlaps with $b5$.

6.3. *Changes in Illuminance Preference*

To evaluate the system's responsiveness to changes in users' illuminance preferences, we require the system to maintain constant illuminance levels of 300, 350 and 500 lux on three of the desks, while on the fourth desk a new illuminance preference is set every 10-15 seconds by its occupant.

The results of this experiment are shown in Fig. 8. On the top plot, solid lines correspond to sensor readings, while the dashed line corresponds to the user's illuminance preference that has been changed several times throughout the experiment. The bottom plot shows dimming levels on the bulbs.

Note that the biggest changes in dimming levels are for bulbs that most affect a sensor, e.g., bulbs 1 and 4 for sensor 1. By comparing the top time series to the bottom ones, one can see

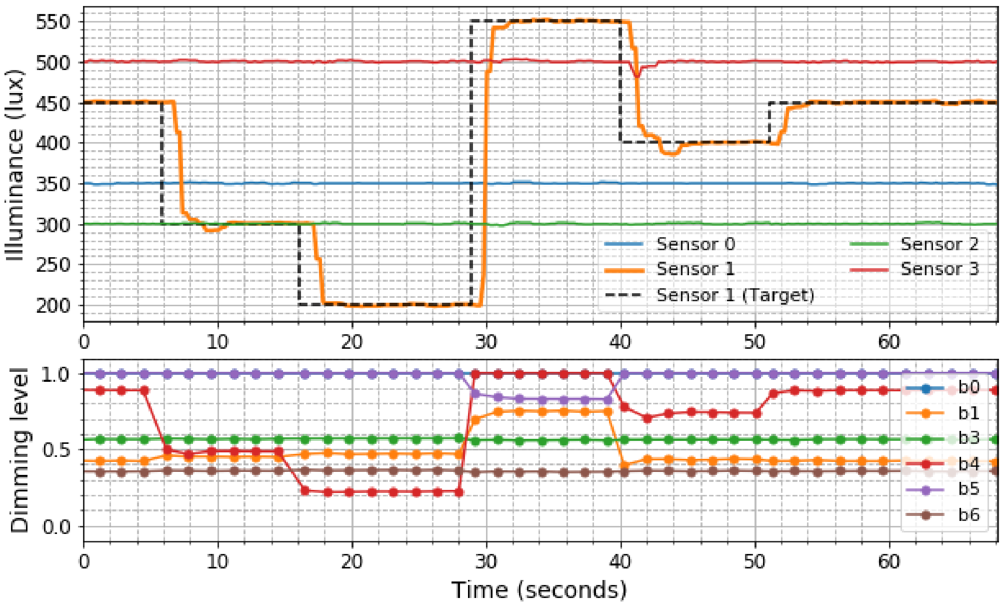

Figure 8: Response to changes in target illuminance. $b0$ partially overlaps with $b5$.

that the target illuminance changes are followed by the system's reaction almost immediately, within a fraction of a second. Recall from Section 4.2 that when a change in target illuminance is registered, the system immediately terminates the ongoing iteration of the control loop and restarts the optimizer with the new target illuminance settings. This results in a fast reaction time to change of about 350ms. After dimming levels are set, it takes the bulbs at most 1.3s. to fully adapt to these new dimming settings. Thus, provided that we have an accurately estimated matrix $\widetilde{A}$ (i.e., it takes 1 iteration for the system to converge), the maximum total time to adapt to the new target illuminance is about 1.65s.

6.4. Changes in Occupancy

We next test the system's response to sequential changes in occupancy of all 4 work stations. We simulate a scenario where users come to their work station one by one, stay for 50-60 seconds, and then leave. The target illuminance on occupied desks is set according to user preferences, which are chosen to be 300, 350, 450 and 500 lux. On the other hand, the target illuminance of unoccupied desks is 0 lux, as we assume that an unoccupied desk does not have to be illuminated.

The results of a typical experiment are shown in Fig. 9. The top four plots show real and target illuminances on the four light sensors, indicated by solid and dashed lines, respectively. The bottom plot shows the dynamically changing dimming levels of the bulbs. We find that the system succeeds at illuminating the occupied work stations according to the users' preferences, while almost not illuminating the unoccupied ones. By examining the experimental data, we find that the system reacts to changes in occupancy within 350ms, and fully adapts to them in $\sim$1.65s.[6]

[6]Responsiveness to changes in occupancy and users' preferences is expected to be the same since both of them are

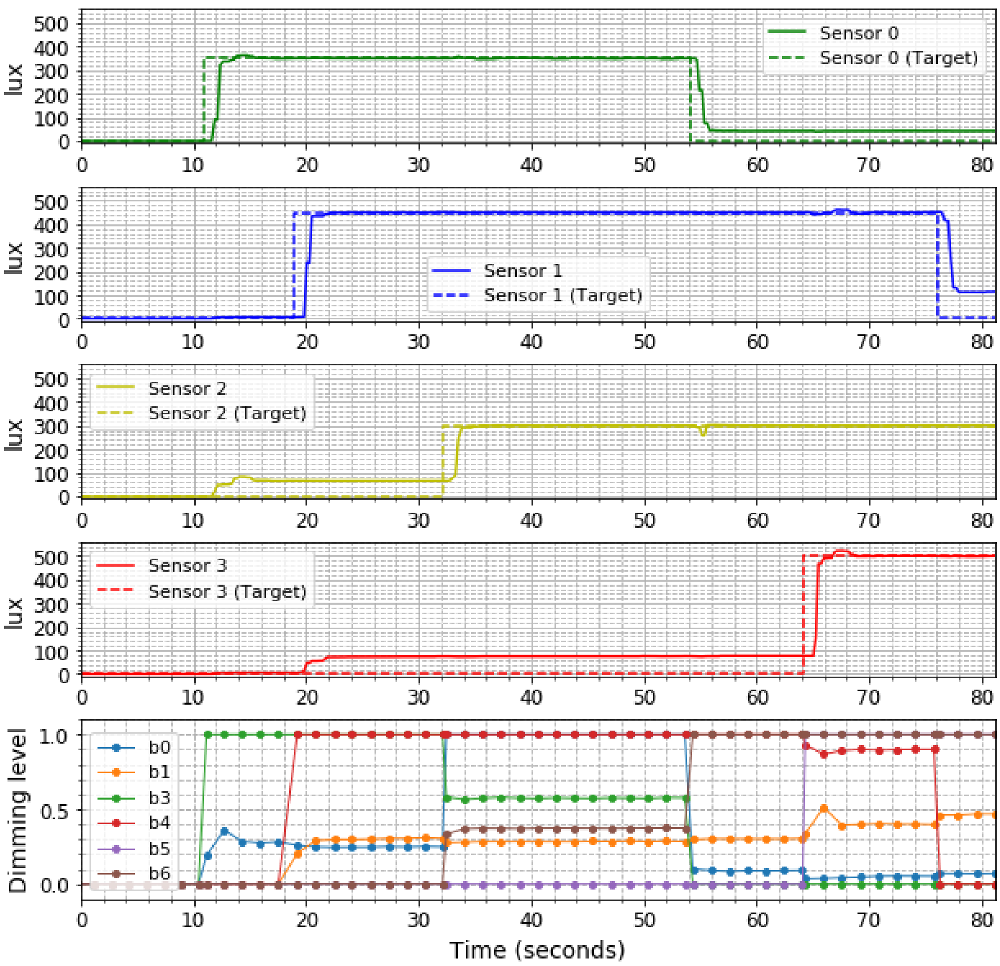

Figure 9: Response to changes in occupancy.

Note that sometimes it is impossible to achieve the exact target illuminance on a sensor due to physical limitations. For instance, in Fig. 9, the illuminance levels at unoccupied desks are sometimes higher than the target value of 0 lux. The unoccupied desk gets some unintended illumination, when a neighbouring desk's target illuminance value is high. However, the control algorithm always tries to minimize any over-illumination, due to the optimization program's objective function that minimizes overall power consumption.

6.5. Effect of Error in Matrix A on Performance

In practice, an estimated illuminance gains matrix $\widetilde{A}$ is subject to error due to several factors including inaccurate calibration, accident sensor movements, and decrease in bulbs' luminous flux with temperature. Recall that Eq. 23 shows the theoretical effect of this error on the convergence of the control algorithm. This section empirically investigates the performance of the system when the estimated illuminance gains matrix $\widetilde{A}$ has various degrees of inaccuracy. Specifically, we artificially introduce errors in the $\widetilde{A}$ matrix by partially covering sensors *during the calibration phase* so that they under-report the true illuminance[7]. The system then tries to maintain constant heterogeneous illuminance levels on the four light sensors, namely, 175 lux, 200 lux, 225 lux and 250 lux, despite changes in the environment due to the laboratory's central lighting being turned on and off.

Typical results of these experiments are shown in Fig. 10. White and yellow regions correspond to the laboratory's main light being off and on, respectively, causing a sharp change in the environment. We show the illuminance at the four sensing modules with a 0%, 30%, and 60% average error in the illuminance gains matrix $\widetilde{A}$. Note that the system rapidly converges even with 30% error. However, the system does not converge when $\widetilde{A} \approx \epsilon(t)$, for example, when the mean error is 60%.

6.6. Reduction in Energy Consumption

Lighting power consumption is a function of bulbs' dimming levels (plus the power required by sensors/microcomputers). Given a particular work station configuration and a desired workspace illumination, these dimming levels depend on the occupancy of work stations and the level of available daylight. To evaluate the energy saving potential of our work, we built a custom simulator that estimated the energy cost of lighting using different technology options. We first describe how we chose the work station configuration and lighting levels, then discuss how the occupancy and daylight availability were modeled.

The simulated work station configuration is the testbed described in Section 6 with six 1300-lumen bulbs. The simulations approximate the time-varying illuminance gains matrix $A(t)$ by the constant estimate $\widetilde{A} \in \mathbb{R}^{4\times6}$. Two target illuminance requirements were chosen: at least 450 lux for occupied work stations, and 0 lux for unoccupied work stations. With these assumptions, the only information required by the simulator to compute time-varying optimal dimming levels (and the corresponding power consumption) are occupancy and daylight signals ($\vec{o}(t) \in \mathbb{R}^{4\times1}$ and $\vec{E}(t) \in \mathbb{R}^{4\times1}$, respectively) from each sensing module. We obtained these from 7 months of measured illuminance and occupancy data, as discussed next.

---

expressed by the changes in target illuminance.

[7]Note that sensors are occluded only during the calibration phase. During the operation phase, they are uncovered (which causes the estimation error in matrix $\widetilde{A}$).

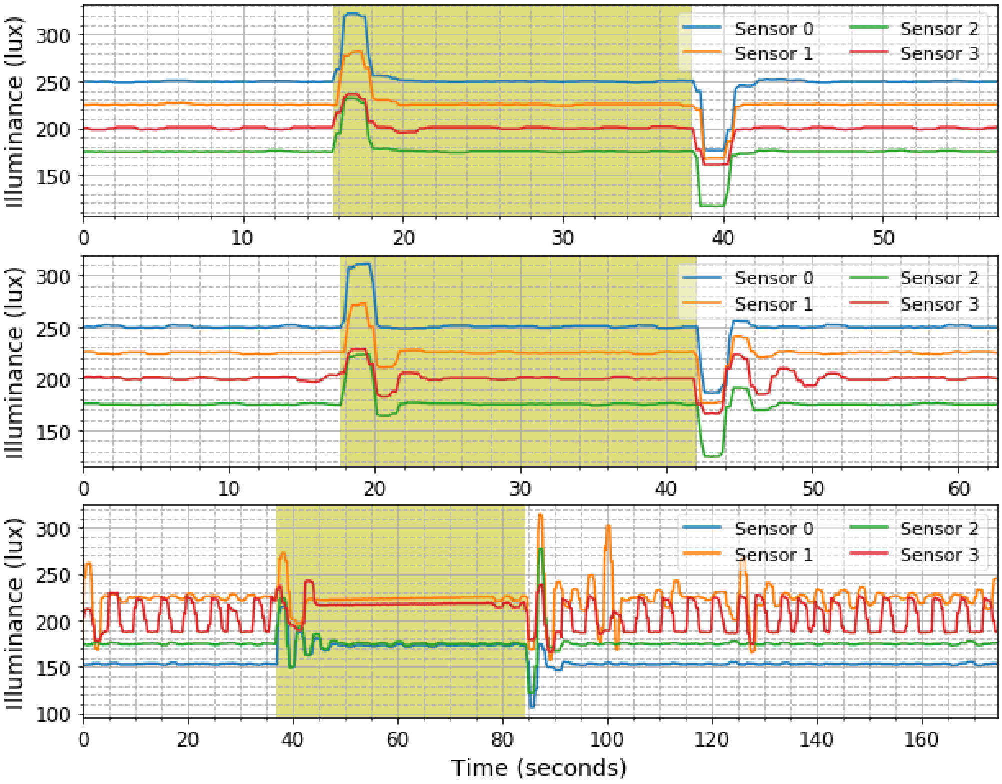

Figure 10: Impact of 0%, 30%, and 60% mean estimation error on convergence.

Recall that our experimental testbed has no windows. Therefore, the daylight illuminance data was collected from three other unoccupied offices, each with a large window, during the 7-month period from October 1, 2018, to May 1, 2019, in Waterloo, ON, Canada. Each of the offices was instrumented with two custom-built sensing modules, based on the Onion Omega microcontroller [46] augmented with a TSL 2561 light sensor, installed in two different locations. These sensing modules logged illuminance measurements every minute. Details on design, calibration, and management of this auxiliary sensing system can be found in the extended version of this paper [45].

Occupancy data used in the simulation was collected in the same building in the *SPOT* project [47]. The dataset contains 7-month long occupancy signals collected from 20 distinct work stations that belong to graduate students, faculty members or administrative staff, downsampled to 1 minute.

We collected two illuminance signals in each office, but there are four work stations in the testbed, so we added a zero-mean, $\sigma_{noise} = 0.05\mu_{signal}$ Gaussian noise to each daylight signal to mimic the daylight illuminance on two sensing modules from neighbouring work stations. The illuminance and occupancy signals are then randomly combined, resulting in 7-month long combined $(t, \vec{E}(t), \vec{o}(t))$ 1-minutely signals, as shown in Fig. 11.

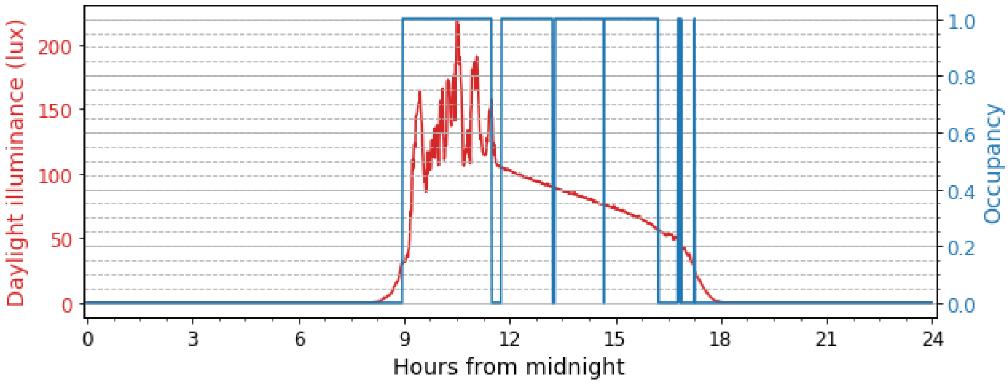

Figure 11: Example of daylight and occupancy signals.

We estimate the amount of energy consumed by bulbs in each 1-minute interval by computing the optimal dimming level vector $\vec{d}(t)$ and compare this to several other lighting system configurations, as described next. For existing systems based on incandescent[8] or fluorescent bulbs, we study two variants, one that is always on[9] (labelled *INC basic* and *CFL basic*), and one that turns on all bulbs if any work station is occupied (*INC occup-ol* and *CFL occup-ol*). For LED systems, we study four variants: with neither daylight harvesting nor occupancy awareness, i.e., simply replacing existing bulbs with LED bulbs (*LED basic*), with only daylight harvesting but not occupancy awareness (*LED-daylight*), with only office-level occupancy sensing but no daylight harvesting (*LED-occup-ol*) and with both daylight harvesting and per-desk occupancy sensing (*LED-daylight-occup-dl*). For a fair comparison, all systems use six 1300 lumen-rated

[8]Although these have been phased out in much of the world, we present this data as a point of reference.

[9]For all occupancy-unaware systems, we assume the light is turned on by the first person to arrive and turned off by the last to leave.

bulbs which consume 13 W for LED, 27 W for CFL [48] and 85 W for incandescent bulbs [49], respectively. In addition, note that systems with an occupancy detection and/or daylight harvesting consume 6.2W for sensing, computing and communicating or 0.1488kW h/day.

We conduct 90 simulations for each configuration using randomly combined daylight and occupancy traces. Fig. 12 shows the average daily energy consumption of different lighting systems, with 95% error bars. As expected, incandescent bulbs consume roughly 3 times more energy than the systems using CFL bulbs which, in turn, consume 2-3 times more energy than LED-based lighting systems. Office-level occupancy detection decreases the average daily energy consumption of the incandescent bulb-based and CFL-based systems by 15% and 8.5%, respectively.

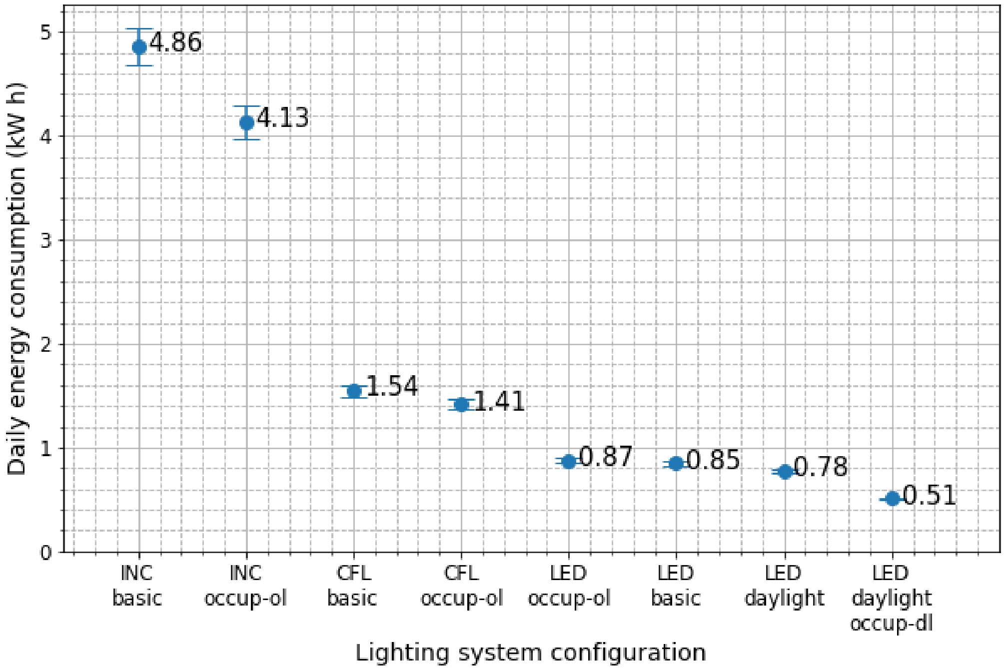

Figure 12: Daily energy consumption of different lighting system configurations.

Comparing the four LED-based systems, the average daily energy consumption of systems with either daylight harvesting or office-level occupancy detection alone is not significantly different from that of a basic LED lighting system (0.78 kWh vs. 0.87 kWh vs. 0.85 kWh, respectively). However, an LED lighting system that has *both* daylight harvesting and per-desk occupancy detection on average consumes 40% less thank the basic LED system (and about 3x less than the basic CFL system and 9x less than the basic incandescent system). Of course, the actual level of savings depend greatly on the degree of occupancy and the level of daylight available in the work stations. Interestingly, an LED lighting system with the office-level occupancy detection capability on average consumes slightly *more* energy than the basic one. This is because occupancy sensing requires an additional constant power of 6.2 W for the controller and wireless sensors, more than offsetting the reduction in bulb power consumption. Clearly, with very low occupancy, the standby losses would make smart lighting systems that do not have per-desk occupancy sensing cost ineffective. However, we are not able, at the present time, to specify the occupancy level at which this crossover would take place. In any case, note that our system, with the use of per-desk occupancy sensors, is cost-effective with the occupancy levels measured in our dataset. In work not presented here, we studied using wireline control and found that this, as expected, reduces the power cost, though at the expense of deployment complexity.

7. Conclusions

We present a power-efficient smart lighting control system that, in a realistic evaluation, reduces energy consumption by about 40% compared to a conventional system and is able to maintain heterogeneous illumination in the office, while quickly responding to dynamically changing illuminance preferences of the occupants. It is robust to errors and quickly adapts to changes in environmental illuminance and occupancy. System deployment is plug-and-play. Moreover, although we have not elaborated on this in the paper, new system components, such as additional bulbs and sensing modules, can be seamlessly connected and disconnected, even while the system is in use.

Our work has four limitations. First, we made some simplifying approximations, such as not modeling the effect of temperature on the bulbs' power consumption, and linearizing the power model. These can potentially cause small errors (up to 3%) in a bulb's power consumption estimates. Other errors, such as due to daylight fluctuations and sensor miscalibration, are likely to be larger sources of error in practice. Second, our system measures illuminance near work surfaces and not directly on them. However, in practice, this is likely not an issue because users can readjust the desired illuminance levels to compensate for sensor placement errors. Third, our analysis and implementation focuses on the Philips Hue bulbs, although our methods and software could be generalized to work with other software-controllable bulbs. Finally, we conducted our experimental evaluation in the laboratory setting without an external window. Hence, we were unable to accurately reproduce fluctuations in daylight. We note that our system is open source and we will also make our data traces publicly available.

In future work, in addition to controlling bulbs, we plan to consider controlling blinds to limit the daylight that enters the room, allowing us to optimize both the lower and the upper bound on the illuminance of the sensors. We would also like to implement a graphical user interface to intuitively set illuminance preferences.

References

[1] M. File, Commercial buildings energy consumption survey (cbecs), US Department of Energy: Washington, DC, USA (2012).

[2] C. de Bakker, M. Aries, H. Kort, A. Rosemann, Occupancy-based lighting control in open-plan office spaces: A state-of-the-art review, Building and Environment 112 (2017) 308 – 321. URL: http://www.sciencedirect.com/science/article/pii/S0360132316304735. doi:https://doi.org/10.1016/j.buildenv.2016.11.042.

[3] A. Pandharipande, D. Caicedo, Smart indoor lighting systems with luminaire-based sensing: A review of lighting control approaches, Energy and Buildings 104 (2015) 369 – 377. URL: http://www.sciencedirect.com/science/article/pii/S0378778815301511. doi:https://doi.org/10.1016/j.enbuild.2015.07.035.

[4] D. Caicedo, S. Li, A. Pandharipande, Smart lighting control with workspace and ceiling sensors, Lighting Research & Technology 49 (2017) 446–460. URL: https://doi.org/10.1177/1477153516629531. doi:10.1177/1477153516629531. arXiv:https://doi.org/10.1177/1477153516629531.

[5] S. Borile, A. Pandharipande, D. Caicedo, L. Schenato, A. Cenedese, A data-driven daylight estimation approach to lighting control, IEEE Access 5 (2017) 21461–21471. doi:10.1109/ACCESS.2017.2679807.

[6] Y.-J. Wen, A. Agogino, Control of wireless-networked lighting in open-plan offices, Lighting Research and Technologyy 43 (2011) 235 – 248.

[7] N. Kandasamy, G. Karunagaran, C. Spanos, K. Tseng, B.-H. Soong, Smart lighting system using ANN-IMC for personalized lighting control and daylight harvesting, Building and Environment 139 (2018) 170 – 180. URL: http://www.sciencedirect.com/science/article/pii/S0360132318302658. doi:https://doi.org/10.1016/j.buildenv.2018.05.005.

[8] B. Roisin, M. Bodart, A. Deneyer, P. D'Herdt, Lighting energy savings in offices using different control systems and their real consumption, Energy and Buildings 40 (2008) 514 – 523. URL: http://www.sciencedirect.com/science/article/pii/S037877880700134X. doi:https://doi.org/10.1016/j.enbuild.2007.04.006.

[9] L. Bellia, F. Fragliasso, E. Stefanizzi, Why are daylight-linked controls (DLCs) not so spread? A literature review, Building and Environment 106 (2016) 301 – 312. URL: http://www.sciencedirect.com/science/article/pii/S0360132316302475. doi:https://doi.org/10.1016/j.buildenv.2016.06.040.

[10] A. Pandharipande, G. Newsham, Lighting controls: Evolution and revolution, Lighting Research & Technology 50 (2018) 115–128. URL: https://doi.org/10.1177/1477153517731909. doi:10.1177/1477153517731909. arXiv:https://doi.org/10.1177/1477153517731909.

[11] G. Newsham, M. Aries, S. Mancini, G. Faye, Individual control of electric lighting in a daylit space, Lighting Research and Technology 40 (2008) 514 – 523. doi:https://doi-org.proxy.lib.uwaterloo.ca/10.1177/1477153507081560.

[12] S. Afshari, S. Mishra, A plug-and-play realization of decentralized feedback control for smart lighting systems, IEEE Trans. on Control Systems Technology 24 (2016) 1317–1327. doi:10.1109/TCST.2015.2487880.

[13] Y. K. Tan, T. P. Huynh, Z. Wang, Smart personal sensor network control for energy saving in DC grid powered LED lighting system, IEEE Trans. on Smart Grid 4 (2013) 669–676. doi:10.1109/TSG.2012.2219887.

[14] F. Rubinstein, G. Ward, R. Verderber, Improving the performance of photo-electrically controlled lighting systems, Journal of the Illuminating Engineering Society 18 (1989) 70–94. URL: https://doi.org/10.1080/00994480.1989.10748745. doi:10.1080/00994480.1989.10748745. arXiv:https://doi.org/10.1080/00994480.1989.10748745.

[15] F. Rubinstein, Photoelectric control of equi-illumination lighting systems, Energy and Buildings 6 (1984) 141 – 150. URL: http://www.sciencedirect.com/science/article/pii/0378778884900690. doi:https://doi.org/10.1016/0378-7788(84)90069-0.

[16] L. Doulos, A. Tsangrassoulis, F. Topalis, Multi-criteria decision analysis to select the optimum position and proper field of view of a photosensor, Energy Conversion and Management 86 (2014) 1069 – 1077. URL: http://www.sciencedirect.com/science/article/pii/S0196890414005573. doi:https://doi.org/10.1016/j.enconman.2014.06.032.

[17] M. H. et al, A review on lighting control technologies in commercial buildings, their performance and affecting factors, Renewable and Sustainable Energy Reviews 33 (2014) 268 – 279. URL: http://www.sciencedirect.com/science/article/pii/S1364032114001166. doi:https://doi.org/10.1016/j.rser.2014.01.090.

[18] A. Williams, B. Atkinson, K. Garbesi, E. Page, F. Rubinstein, Lighting controls in commercial buildings, LEUKOS 8 (2012) 161–180. doi:10.1582/LEUKOS.2012.08.03.001.

[19] I. Chew, D. Karunatilaka, C. Tan, V. Kalavally, Smart lighting: The way forward? Reviewing the past to shape the future, Energy and Buildings 149 (2017) 180 – 191. URL: http://www.sciencedirect.com/science/article/pii/S0378778817300440. doi:https://doi.org/10.1016/j.enbuild.2017.04.083.

[20] K. R. Wagiman, M. N. Abdullah, M. Y. Hassan, N. H. M. Radzi, T. C. Kwang, et al., Lighting system control techniques in commercial buildings: Current trends and future directions, Journal of Building Engineering 31 (2020) 101342.

[21] D. Plörer, S. Hammes, M. Hauer, V. van Karsbergen, R. Pfluger, Control strategies for daylight and artificial lighting in office buildings—a bibliometrically assisted review, Energies 14 (2021) 3852.

[22] A. Pandharipande, M. Rossi, D. Caicedo, L. Schenato, A. Cenedese, Centralized lighting control with luminaire-based occupancy and light sensing, in: 2015 IEEE 13th International Conf. on Industrial Informatics (INDIN), 2015, pp. 31–36. doi:10.1109/INDIN.2015.7281706.

[23] A. Pandharipande, D. Caicedo, X. Wang, Sensor-driven wireless lighting control: System solutions and services for intelligent buildings, IEEE Sensors Journal 14 (2014) 4207–4215. doi:10.1109/JSEN.2014.2351775.

[24] D. Caicedo, A. Pandharipande, Distributed illumination control with local sensing and actuation in networked lighting systems, IEEE Sensors Journal 13 (2013) 1092–1104. doi:10.1109/JSEN.2012.2228850.

[25] N. van de Meugheuvel et al, Distributed lighting control with daylight and occupancy adaptation, Energy and Buildings 75 (2014) 321 – 329. URL: http://www.sciencedirect.com/science/article/pii/

S0378778814001273. doi:https://doi.org/10.1016/j.enbuild.2014.02.016.

[26] D. Caicedo, A. Pandharipande, Daylight and occupancy adaptive lighting control system: An iterative optimization approach, Lighting Research & Technology 48 (2016) 661–675. URL: https://doi.org/10.1177/1477153515587148. doi:10.1177/1477153515587148. arXiv:https://doi.org/10.1177/1477153515587148.

[27] M. T. Koroglu, K. M. Passino, Illumination balancing algorithm for smart lights, IEEE Trans. on Control Systems Technology 22 (2014) 557–567. doi:10.1109/TCST.2013.2258399.

[28] G. Boscarino, M. Moallem, Daylighting control and simulation for LED-based energy-efficient lighting systems, IEEE Trans. on Industrial Informatics 12 (2016) 301–309. doi:10.1109/TII.2015.2509423.

[29] S. Borile, A. Pandharipande, D. Caicedo, A. Cenedese, L. Schenato, An identification approach to lighting control, in: 2016 European Control Conf. (ECC), 2016, pp. 637–642. doi:10.1109/ECC.2016.7810360.

[30] M. Miki, A. Amamiya, T. Hiroyasu, Distributed optimal control of lighting based on stochastic hill climbing method with variable neighborhood, in: 2007 IEEE International Conf. on Systems, Man and Cybernetics, 2007, pp. 1676–1680. doi:10.1109/ICSMC.2007.4413957.

[31] Y.-J. Wen, A. M. Agogino, Personalized dynamic design of networked lighting for energy-efficiency in open-plan offices, Energy and Buildings 43 (2011) 1919 – 1924. URL: http://www.sciencedirect.com/science/article/pii/S0378778811001459. doi:https://doi.org/10.1016/j.enbuild.2011.03.036.

[32] Radiance synthetic imaging system, https://floyd.lbl.gov/radiance/index.html, Accessed: 2019-07-19.

[33] L. Yeh, C. Lu, C. Kou, Y. Tseng, C. Yi, Autonomous light control by wireless sensor and actuator networks, IEEE Sensors Journal 10 (2010) 1029–1041. doi:10.1109/JSEN.2010.2042442.

[34] M. Pan, L. Yeh, Y. Chen, Y. Lin, Y. Tseng, A WSN-based intelligent light control system considering user activities and profiles, IEEE Sensors Journal 8 (2008) 1710–1721. doi:10.1109/JSEN.2008.2004294.

[35] A. Ravi, K. Galmath, S. Hu, A. Misra, Cs-light: camera sensing based occupancy-aware robust smart building lighting control, in: Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, 2021, pp. 61–70.

[36] K. V. D. Wymelenberg, M. Inanici, Evaluating a new suite of luminance-based design metrics for predicting human visual comfort in offices with daylight, LEUKOS 12 (2016) 113–138. URL: https://doi.org/10.1080/15502724.2015.1062392. doi:10.1080/15502724.2015.1062392. arXiv:https://doi.org/10.1080/15502724.2015.1062392.

[37] T. Lashina, S. van der Vleuten-Chraibi, M. Despenic, P. Shrubsole, A. Rosemann, E. van Loenen, A comparison of lighting control strategies for open offices, Building and Environment 149 (2019) 68 – 78. URL: http://www.sciencedirect.com/science/article/pii/S0360132318307492. doi:https://doi.org/10.1016/j.buildenv.2018.12.013.

[38] G. Newsham, J. Veitch, C. Arsenault, C. Duval, Effect of dimming control on office worker satisfaction and performance, Proceedings of the IESNA Annual Conf. (2004) 19–41.

[39] G. Newsham, J. Veitch, Lighting quality recommendations for VDT offices: a new method of derivation, Trans. of the Illuminating Engineering Society 33 (2001) 97–113. URL: https://doi.org/10.1177/136578280103300205. doi:10.1177/136578280103300205. arXiv:https://doi.org/10.1177/136578280103300205.

[40] A. D. Galasiu, J. Veitch, Occupant preferences and satisfaction with the luminous environment and control systems in daylit offices: a literature review, Energy and Buildings 38 (2006) 728 – 742. URL: http://www.sciencedirect.com/science/article/pii/S0378778806000624. doi:https://doi.org/10.1016/j.enbuild.2006.03.001.

[41] S. Matta, S. M. Mahmud, An intelligent light control system for power saving, in: IECON 2010 - 36th Annual Conf. on IEEE Industrial Electronics Society, 2010, pp. 3316–3321. doi:10.1109/IECON.2010.5675331.

[42] Philips Hue website, https://www2.meethue.com/en-us/p/hue-white-single-par38-outdoor/046677476816, Accessed: 2020-08-03.

[43] How are LEDs affected by heat?, https://www.lrc.rpi.edu/programs/nlpip/lightinganswers/led/heat.asp, Accessed: 2020-08-03.

[44] Python API for Philips Hue luminaires (github page), https://github.com/studioimaginaire/phue, Accessed: 2019-07-24.

[45] Aussat, Yerbol, A power-efficient smart lighting system: Modeling, implementation and evaluation, http://hdl.handle.net/10012/15008, Master's thesis, 2019.

[46] Onion website: Omega2, https://onion.io/omega2/, Accessed: 2020-08-03.

[47] A. Rabbani, S. Keshav, The SPOT* personal thermal comfort system, in: Proceedings of the 3rd ACM International Conf. on Systems for Energy-Efficient Built Environments, BuildSys '16, ACM, New York, NY, USA, 2016, pp. 75–84. URL: http://doi.acm.org/10.1145/2993422.2993578. doi:10.1145/

2993422.2993578.
[48] 27 W full spectrum CFL light bulb 5500 K, 1300 lumens (Amazon page), https://www.amazon.com/ALZO-Spectrum-Light-Lumens-Daylight/dp/B0018OS06S, Accessed: 2019-08-01.
[49] Energy guide: How do I read the energy label of a light bulb?, https://www.energuide.be/en/questions-answers/how-do-i-read-the-energy-label-of-a-light-bulb/50/, Accessed: 2019-08-01.