Pervasive and Mobile Computing 45 (2018) 35–54

Contents lists available at ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Counting via LED sensing: Inferring occupancy using lighting infrastructure

Yanbing Yang [a,b,*], Jun Luo [a], Jie Hao [c], Sinno Jialin Pan [a]

[a] School of Computer Science and Engineering, Nanyang Technological University, Singapore
[b] College of Computer Science, Sichuan University, China
[c] College of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, China

ARTICLE INFO

Keywords:
Occupancy inference
Visible Light Sensing
Statistical learning

ABSTRACT

As a key component of building management and security, occupancy inference through smart sensing has attracted a lot of research attention for nearly two decades. Nevertheless, existing solutions mostly rely on either pre-deployed infrastructures or user device participation, thus hampering their wide adoption. This paper presents CeilingSee, a dedicated occupancy inference system free of heavy infrastructure deployments and user involvements. Building upon existing LED lighting systems, CeilingSee converts part of the ceiling-mounted LED luminaires to act as sensors, sensing the variances in diffuse reflection caused by occupants. In realizing CeilingSee, we first re-design the LED driver to leverage LED's photoelectric effect so as to transform a light emitter to a light sensor. In order to produce accurate occupancy inference, we then engineer efficient learning algorithms to fuse sensing information gathered by multiple LED luminaires. We build a testbed covering a 30 m^2 office area; extensive experiments show that CeilingSee is able to achieve very high accuracy in occupancy inference.

1. Introduction

The awareness of (indoor) occupancy is crucial to many aspects of smart building management; these include, among others, controlling the HVAC and lighting systems for the sake of energy conservation, choosing the right information service based on congestion level, as well as safely evacuating people under life-threatening circumstances. In the past two decades, various smart sensing technologies have been dedicated[1] to infer occupancy for indoor facilities, and most of them require deploying certain sensing infrastructure with mainly three typical sensors: passive infra-red (PIR) [4,5], acoustic/ultrasonic [6–8], and camera [9,10]. Other solutions attempt to infer occupancy indirectly by monitoring the usage of existing services (e.g., Wi-Fi [11] and power grid [12]). Whereas the former method requires installations of extra infrastructure and hence incurs both high cost for building management and potential infringement of user privacy, the latter approach can hardly be accurate: what if some occupants simply do not use any services?

Our key observation is that, in any human occupied indoor spaces, lighting is a necessity while the resulting diffuse reflection can be "disturbed" by the presence of occupants. In the meantime, *Visible Light Sensing* (VLS), as a variance of heavily studied *Visible Light Communication* (VLC) [13–17], has started to show its potential in many sensing-intensive applications [18]. Therefore, a natural question is: *can we apply VLS to build an occupancy inference system that is free of*

---

* Corresponding author at: School of Computer Science and Engineering, Nanyang Technological University, Singapore.
E-mail address: yyang017@ntu.edu.sg (Y. Yang).

[1] Though occupancy information can be derived from an indoor localization system (e.g., [1–3]), no practical indoor localization system has been widely adopted so far. Moreover, relying on user location tracking to "count" occupancy is highly inefficient and may infringe privacy.

https://doi.org/10.1016/j.pmcj.2018.01.003
1574-1192/

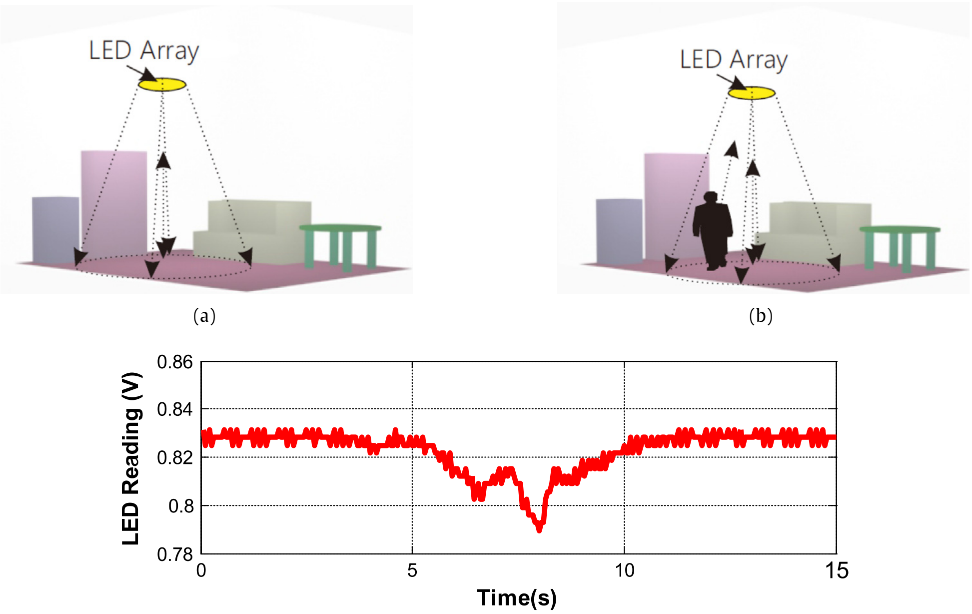

(a) (b) (c)

Fig. 1. Inferring occupancy by LED sensing. (a) The sensing coverage of an LED array. (b) An occupant enters the covered area. (c) The variance in LED readings due to the absence/presence of the occupant.

reliance on both heavy infrastructures and user involvements? In this paper, through a detailed exposition of our CeilingSee system, we intend to provide readers with a positive answer.

The first idea of CeilingSee is very intuitive: from the ceiling point of view, the diffuse reflection (consisting of main reflections from the floor and variously fixed furniture on the floor) is bounded to be affected by the presence of occupants. Therefore, sensing such perturbations could allow us to infer occupancy. While simply installing an array of light sensors on the ceiling could be a solution, it would introduce yet another infrastructure. Fortunately, the increasing popularity of LED lighting systems and the readily verifiable photoelectric effect of LED [19,14] have motivated our another novel idea: re-designing the driver of a *Commercial Off-The-Shelf* (COTS) LED could enable it to serve as both a light emitter and a light sensor. Consequently, CeilingSee simply leverages the existing LED lighting systems and borrows a fraction of LED chips to sense the variance in diffuse reflection. We illustrate these ideas in Fig. 1; it clearly demonstrates the potential and effectiveness of VLS-based occupancy inference.

The seemingly straightforward ideas of CeilingSee impose on us two major challenges.

Firstly, although conventional LED-to-LED communication [14] has already employed an LED as receiver (a special form of sensor), sensing the variance in diffuse reflection is much more challenging due to the very low SNR, hence it necessitates using the collective sensing ability of multiple LED chips. Existing LED receiver directly connects an LED chip to the I/O port of an MCU, thus relying on the controllable nature of the I/O port to toggle the states of the LED between *forward biased* (emitting or sending) and *reverse biased* (sensing or receiving). Unfortunately, this would not work when multiple LED chips are used together, as the voltage/current would exceed what an I/O can take (normally no more than 3.3V/20 mA) and one cannot afford to directly attach an MCU to each LED chip. As a result, we design a novel circuit for accommodating the collective photoelectric effects of an LED array.

Secondly, as the sensing coverage[2] of a single LED array (consisting of multiple "sensorized" LED chips) is limited, we have to use multiple arrays to cover a large indoor area, which happens to be in line with the lighting requirement. Moreover, CeilingSee needs to account for multiple occupants dispersed on the area, especially those not strictly under an LED array. Therefore, it is necessary that efficient inference algorithms are in place to utilize the collective sensing outcomes of all LED arrays. CeilingSee responds to this challenge by engineering a machine learning algorithm that maps the multi-dimensional sensing data to the demanded occupancy count.

To validate our design of CeilingSee, we build a testbed consisting of multiple LED arrays in order to cover a 30 m^2 office area. We implement the hardware part for controlling the LED arrays, as well as the software part for sensing data processing and hence occupancy inference. Our main contributions are as follows:

---

[2] It is also termed *Field of View* (FoV) in sensing nomenclature, so we use coverage and FoV interchangeably hereafter.