Visible Light based Activity Sensing using Ceiling Photosensors

Mohamed Ibrahim†, Viet Nguyen†, Siddharth Rupavatharam, Minitha Jawahar,
Marco Gruteser, Richard Howard
WINLAB, Rutgers University
†The first two authors are co-primary student authors.
{mibrahim, vietnh, gruteser, reh}@winlab.rutgers.edu
{siddharth.r, minitha.jawahar}@rutgers.edu

ABSTRACT

This paper explores the feasibility of tracking motion and activities of humans using visible light. Shadows created by casting visible light on humans and objects are sensed using sensors that are embedded along with the light sources. Existing Visible Light Sensing (VLS) techniques require either light sensors deployed on the floor or a person carrying a device. Our approach seeks to measure light reflected off the floor to achieve an entirely device-free and light-source based system. We co-locate photosensors with LED light sources to observe the changes in light level occurring on the floor. Our system also uses a sensitive standard difference measurement technique to detect small voltage changes, together with a time division flickering scheme to differentiate between light nodes. We evaluate the feasibility of our system in detecting simple activities, and show that it can detect door opening events at 12% Equal Error Rate.

CCS Concepts

•**Networks** → *Wireless local area networks;*

Keywords

Visible light sensing; Ceiling photosensors; Activity sensing

1. INTRODUCTION

Activity sensing on a building at a wide scale can support a broad spectrum of applications for indoor environments. In smart homes, for example, it could enhance controls on lighting, heating, ventilation, and air conditioning based on sensed and predicted activities across rooms. Useful information ranges from basic occupancy and movement tracking, to activity inference (e.g., sleeping, cooking, eating, watching TV or media).

These kinds of activities can currently be detected by a number of dedicated sensing systems. One approach leverages sensing devices continuously and physically connected with the user (wearable sensing) like smart watches and smart phones (e.g., [3]), that users continuously carry, wear, and usually charge. Device-free solutions have relied on cameras [6] or wifi-based activity sensing [11]. These systems have privacy and security issues, even for RF-based systems, since RF signals penetrate walls. Wifi-based activity sensing systems have been found to have interference problems while coexisting with ISM band operating devices like microwaves and cordless phones.

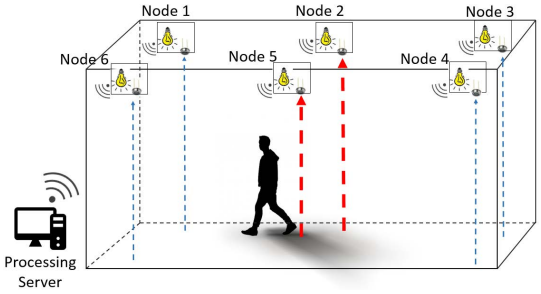

Figure 1: Conceptual Diagram. Node 2 and node 5 detect shadow caused by the person in the room, while other nodes do not. All data is wirelessly sent to the processing server.

Infrared (IR) sensing is one of the most prevalent techniques for motion sensing, starting with detecting motion by the blocking of transmission between an IR sender and receiver. Afterwards, passive or Pyroelectric infrared (PIR) sensors were introduced that detect the radiated IR energy from humans and animals [12]. PIR sensors can suffer from reliability problems whenever a sudden heat change happens, like a window or door opening. Moreover, PIR sensors require line-of-sight to the moving object, which limits the range of such sensors.

More recently, VLS has been considered for indoor motion and activity tracking. VLS is appealing because of the following properties: it can leverage existing lighting infrastructure as transmitters (with no building wiring overhead), it does not penetrate walls (preserves privacy and security and makes it easier to determine in which room an activity occurred), and unlike RF techniques it does not cause or suffer from radio interference. With nanometer wavelengths, it is highly sensitive to small motions and small objects compared to RF waves. Moreover, VLS is not affected by temperature changes.

Many current visible light localization techniques are active techniques that require the user to carry a sensor/device (smart phone/light sensor) and localize using standard trilateration leveraging three anchors (light sources) [10, 2, 4, 7, 14] or by sweeping a light beam [1]). Among passive (device-free) techniques [8, 13, 5], researchers

VLCS'16, October 03-07, 2016, New York City, NY, USA

DOI: http://dx.doi.org/10.1145/2981548.2981554

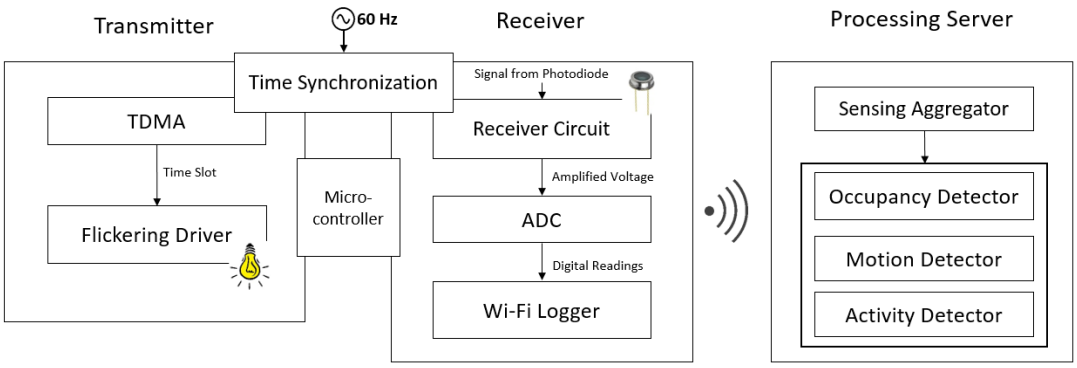

Figure 2: System Architecture.

have [13] used visible light to locate fingers within the workspace of mobile devices. Other work [8] demonstrate fine-grained gesture and human skeleton reconstruction using visible light sensing but require deploying photodiodes on the floor, which limits the scalability of such a motion sensing system.

Shadow-based activity detection. We propose a localization and activity sensing system that is device-free (users do not need to carry or be equipped with any device) and entirely set up on the ceiling (Fig. 1). The system is integrated seamlessly to the current lighting infrastructure, without the need of wiring photosensors on the floor. We observe that by monitoring the shadow cast on the floor by people, doors, and objects, we can infer their positions, assuming the positions of the light bulbs are known. These positions together with their timing information can also lead to activity sensing, such as walking, sitting, etc.

To realize this idea, we face several challenges. *First*, sensing light level changes caused by shadows on the floor is more difficult than sensing direct light from the LED lamps. Since the reflected material can absorb or diffuse the light, the light level visible to the photosensor is much lower. To combat this, we designed sensitive difference measurement circuit based on a Wheatstone bridge to detect small changes in light level caused by shadows. *Second*, each photosensor needs to identify which light bulb causes the observed shadow. We tackle this problem by introducing a TDMA-like scheme based on a common synchronization signal. By doing so, each lamp can be identified through its unique time slot where it blinks.

To demonstrate our idea, we implemented a prototype wherein one LED bulb is used to send light beacons while the existing indoor lighting remains on. This prototype shows the basis of our motion and activity detection algorithm. We also give an example of door opening detection using this prototype.

The remainder of the paper is organized as follows. Section 2 presents an overview of the system, while Section 3 demonstrates its implementation. Section 4 details the experimental setup and presents an evaluation of our current system. We conclude and give possible directions for future work in Section 5.

2. SYSTEM OVERVIEW

One of the main goals of our system is to improve deployability and maintenance by integrating seamlessly with the current lighting infrastructure. Therefore, we co-locate the photosensor and the sensing module with each lamp, which eventually should allow installing a system simply by installing new bulbs in lamp housings. The sensing module is powered by the lamps AC power source, eliminating the need for batteries. Bulbs can coordinate and report sensor information to a processing server or the cloud using Wi-Fi. As a result, our setup is completely on the ceiling; we do not rely on photosensors in other locations and need to install no new devices or wiring.

Our photosensors measure light reflected from the floor instead of direct light from the LED lamps, so a sensitive light measurement method is required. The distance between the ceiling and the floor is as much as 10 feet, further imposing sensitivity requirements on the sensing module. In the next section, we will describe how we apply a difference measurement technique based on a Wheatstone bridge to measure small changes in light intensity caused by shadows.

So that the photosensor can tell which lamp causes the shadow it is observing, we design a time-based identification method, in which each lamp only turns off during its own ***time slot***. This time division flickering scheme requires synchronization between all the light nodes. In the next section, we describe in detail how the synchronization can be achieved, and how each photosensor can identify the light source of each shadow from the timing information.

Fig. 2 presents the architecture of our sensing system. Our system consists of three main parts: transmitter, receiver and processing server. Note that both the transmitter and the receiver are present on the same ***light node***, and both are controlled by a single microcontroller.

2.1 Indirect shadow sensing

Since we want to sense the reflected light from the floor, the amount of light falling on the photosensor is relatively small compared to the direct light from the LED lamp. This imposes a high sensitivity requirement on our receiver circuit. To increase sensitivity, one can think of directly amplifying the absolute voltage measured from the photosensor. However, this approach limits the amount of signal gain we can achieve, since the absolute voltage range can be quite large, and the amplified signal should not saturate the ADC. To achieve more gain, we seek to amplify the **voltage difference** from a reference voltage instead of the absolute voltage value. To do so, we utilize a difference measurement technique based on a Wheatstone bridge circuit, and amplify the change in light level into a measurable voltage range.

As can be seen in Fig. 3, the Wheatstone bridge consists of two arms. One arm contains a photosensor connected with a resistor; the photosensor acts as a current source and produces voltage V_{pd}. The other arm is a voltage divider circuit, in which one of the two resistors is a digital potentiometer that can be controlled by a microcontroller. The voltage across this digital potentiometer is V_{pot}. The two voltages are then fed into an instrumentation amplifier, and the output V_o follows this formula:

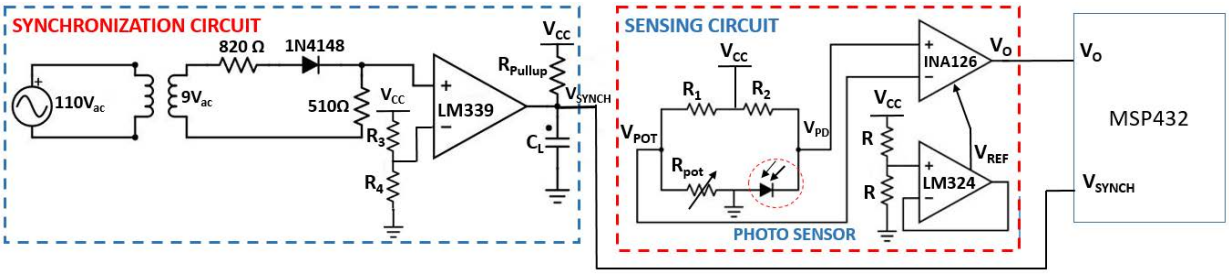

Figure 3: Synchronization and Sensing circuit.

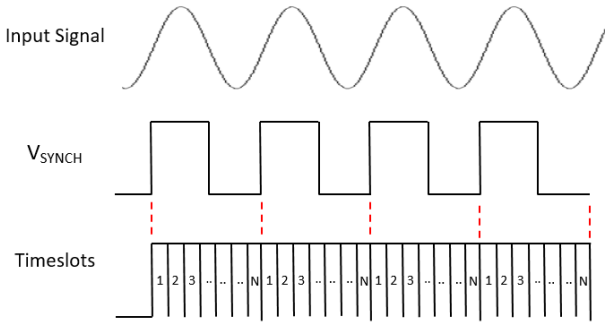

Figure 4: Synchronization Signal.

$$V_o = G \times (V_{pd} - V_{pot}) + V_{ref} \quad (1)$$

in which G denotes the gain of the amplifier and V_{ref} is a reference voltage. The first step is to calibrate the receiver to the ambient light level. This is done when there is no movement in the room and the light level is steady. The microcontroller runs a calibration step to adjust the potentiometer so that the value of V_{pot} comes close to V_{pd}. Once calibration is done, as shown in Eq. 1, with V_{pot} fixed, any small change in V_{pd} would be amplified and reflected in the value of V_o. *When more light shines on the photosensor, V_o decreases, and vice versa.* To preserve the detection performance in our system, the calibration step might have to be performed several times a day to adapt to different light levels.

The output of the amplifier V_o is read by the microcontroller through its built-in ADC. As indicated above, since the receiver is also synchronized with the mains AC power supply, each reading has its own timestamp with respect to the beginning of each cycle. The microcontroller then sends these readings together with their timestamps through Wi-Fi to the server.

2.2 Synchronization and signaling

Our system relies on the time coordination between nodes, so a good synchronization scheme is needed to synchronize all the nodes with the same clock. Fig. 4 demonstrates the different signals in the synchronization circuit shown in Fig. 3. The input signal used to synchronize all the nodes is a common sine wave, which can be obtained from the 60Hz AC mains power [9]. Using a comparator, the synchronization circuit converts this signal to a square wave V_{SYNCH} of the same frequency and feeds it into the microcontroller. The rising edge of this square wave is then used as a trigger to start a new cycle, and the microcontroller uses its internal timer to divide this cycle into N time slots, where N is the number of light nodes in our system. Consequently, all the nodes agree on the timing of each time slot.

The transmitters in our system are the LED lamps along with their flickering circuit. We would like each lamp to use its light flickering pattern to notify its identity besides providing illumination. To do so, each lamp turns ON almost the whole cycle, except its own time slot. The OFF time of each bulb can be seen as its own ***light beacon*** sent to the photosensors: for each time slot inside each cycle, the photosensors read the effect of only *one* lamp. To avoid noticeable flickers to the human eyes, the flickering rate should be above 1 kHz.

Each microcontroller connected to a lamp will pick a different time slot for its lamp to flicker. Each lamp has its own ID, which is a number between 1 and the number of lamps inside the room. Each node uses its ID as a way to pick a unique time slot to flicker the lamp and avoid collisions with other lamps.

2.3 Server-based blockage detection

The processing server is a central machine used for processing the readings coming from all the nodes. We assume that from the node IDs the server knows where the nodes are by mapping the IDs to a map of all the nodes throughout the building.

Our sensing algorithm is based on detecting shadows on the floor resulting from blocking light beams. Suppose there are N nodes A_1, A_2, ..., A_N in our system. When there is blockage (human, doors, etc.) between node A_i and A_1, it would cast shadow from A_i direction to the point on the floor where A_1 points to. Therefore, node A_1 would not see the light beacon in the i^{th} time slot. By looking at the readings reported from node A_1, the processing server would detect the blockage between node A_1 and node A_i.

Blockage detection is the basic building block for further localization and activity detection. Blockages can be translated to humans present/moving, doors opening/closing etc. depending on the known relationship between two light nodes (same room vs. different rooms). For example, after detecting a person inside a room, the system could track this person, by correlating the readings coming from the nodes and their neighbors. Using this technique, we could detect the direction of movement to be between the nodes observing shadow changing.

We focus here on another example: detecting when a door is opened by using two light nodes on opposite sides of the door. Assuming the door is normally closed, when the system detects there is no blockage between these two nodes, it can tell that the door is opened. Depending on the door type, we use different algorithms to detect when the door is opened.

Solid door. When a person enters the room, the receiver sees the light beacon from the transmitter. When the light turns on, the readings at the receiver decrease. Since the transmitter and the receiver are synchronized, the decrease in the readings happens at the same time slot over consecutive cycles. Fig. 9 demonstrates an example, where the light bulb turns on during the first 100 ms

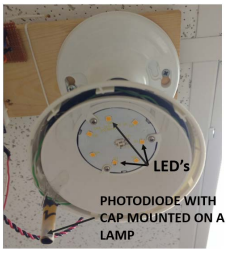

Figure 5: Photosensor co-located with LED lamp.

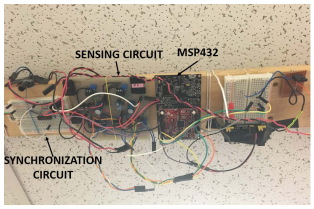

Figure 6: Receiver circuit

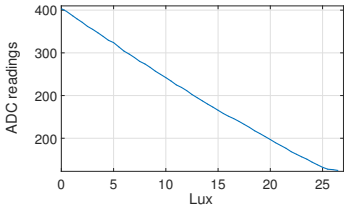

Figure 7: Sensitivity curve of TEPT5600. MSP432 reads voltage (ranging from 0 to 3.3V) and maps it to digital readings from 0 to 1023. Note that the higher the ADC reading, the lower the light intensity.

of each 800 ms cycle, and the sampling rate is 10 ms per reading. When the door is open, we observe a "dip" at the receiver in the first ten readings in almost all 5 consecutive cycles. We do not see this observation when the door is closed since the light from the transmitter is blocked by the door. We use this observation in our detection algorithm.

Glass door. In this case, the door lets light pass through, so the receiver sees V_o decreasing during the first time slot for all cycles when the door is closed. While the door is open and somebody walks through it, because of the occlusion caused by the person and the non-glass part of the door, there would be consecutive cycles where the decrease cannot be observed. We use this observation in our detection algorithm.

Detection algorithm. We compare $m1$, the mean of the first 10 readings, with $m2$, the mean of the last 70 readings of each cycle. If $m1 < m2 - threshold$, we set the light beacon count of this cycle to be 1. Over 5 consecutive cycles, if there are more than 3 light beacons in the solid door case, or less than 2 light beacons in the glass door case, the algorithm reports a door opening event. The threshold value is used in the evaluation of our algorithm.

Future work can focus on detecting more complex activities. This could be simplified by observing the output of motion detection module and correlating it with the knowledge of the purpose of the room. For example, being still in the same place for more than couple of minutes in the bedroom could be translated to a sleeping person. Also, from the motion detection output, we can directly translate the motion to certain activities like standing, sitting, walking, etc.

3. IMPLEMENTATION

3.1 Prototype

LED lamp. We use off-the-shelf Ecosmart 65W BR30 LED bulbs that operate at 42V, which we supply using a constant DC voltage source. The only modification done is the removal of the internal AC to DC transformer and the voltage regulator circuit to allow blinking control of the lamp. The LED lamp is placed on the ceiling roughly 10 feet above the ground. A prototype of the photodiode co-located with the LED lamp is shown in Fig. 5.

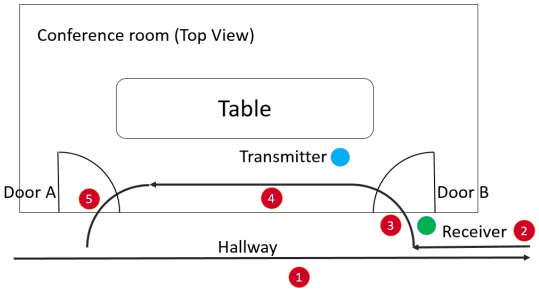

Figure 8: Experimental Scenario.

Receiver circuit. The receiver circuit follows the description in Section 2. We use an INA126 as an instrumentation amplifier. The microcontroller is an MSP432 board, which has an onboard ADC with 10-bit resolution and uses a CC3100 BoosterPack WiFi module to send readings to the server. An image of the prototype is shown in Fig. 6.

Photosensor. The TEPT5600 is an NPN phototransistor whose sensitivity curve is shown in Fig. 7. We produced that figure by dimming the light source while logging the TEPT5600 readings along with off-the-shelf calibrated light sensor (ground truth) readings.The same receiver circuit that was explained earlier is used here and in the rest of paper. As can be seen, this photosensor has a linear current response to the light intensity. The digital ADC readings (and correspondingly voltage at the ADC) decrease with increasing light intensity. For the sensing circuit in Fig. 3 we see that (V_o) follows Eq.(1). When light of higher intensity is present, more current flows through the phototransistor, hence more voltage appears over the series resistor ($R2$) in the Wheatstone bridge which results in a decrease in the voltage drop over the phototransistor. This decrease in the phototransistor voltage (V_{PD}) makes it less than the potentiometer voltage (V_{POT}) (assuming the bridge was balanced before the light change), resulting in lower ADC readings (V_o) following Eq.(1). For noise reduction in the readings and to increase unidirectional sensitivity, we cover the photodiode with a dark cylinder about 3cm long.

Component values. The component values in the prototype are as follows: $R_1 = 1.8K\Omega$, $R_2 = 1M\Omega$, $R = 10K\Omega$, $R_3 = 4.7K\Omega$, $R_4 = 1.8K\Omega$, $R_{pullup} = 4.7K\Omega$, $C_L = 10nF$.

3.2 Experimental setup

The TEPT5600 photosensor has high gain (provides sufficiently large current even when light level is low), but slow switching time. When switching at 60Hz frequency, the photosensor outputs indistinguishable responses to ON and OFF light levels from the light sources. Therefore, we test our idea with a prototype that runs slower than described above. Instead of using a 60Hz AC signal, we use a function generator to provide a 1.2Hz sine wave signal to each module. Thus, there are 800ms in each period, and we divide this into 8 time slots of 100ms each. We measured the accuracy of our synchronization scheme, and found that the time slot is shifted by at most 50ms (about half a time slot). In the current prototype, the TEPT5600 is used here to demonstrate our overall system idea. For future work, we will replace the TEPT5600 photosensor with a photodiode with faster switching time, together with a low-noise front-end to be able to detect small changes in light intensity. We will also look at a more accurate synchronization method to prevent disambiguity between time slots.

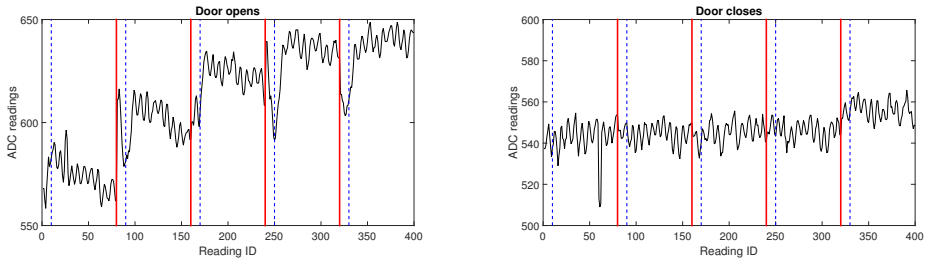

(a) (b)

Figure 9: Sensor readings over consecutive cycles. The red solid lines mark the beginning of each cycle, while the blue dash lines mark the end of the first 10 readings of each cycle. (a) A person enters the room, (b) Door remains closed. When the door is open, its readings are lower during the first 10 readings (the first time slot) consistently over a series of continuous cycles. These decreases are not observed when the door is closed. Note that the higher the ADC reading, the lower the light intensity.

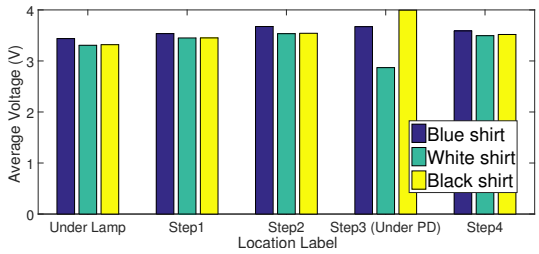

Figure 10: Tracking different users between two nodes. Note that the lesser the light intensity, the higher the voltage.

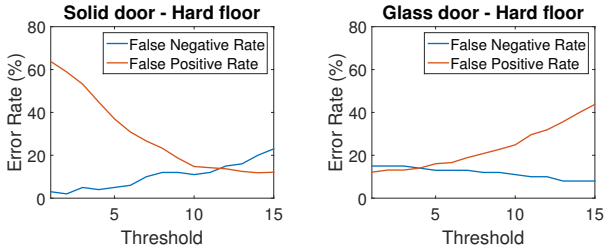

(a) (b)

Figure 11: Performance of our system for different types of doors.

In the ideal scenario, all LED bulbs would flicker following the TDMA flickering scheme described in Section 2.1: each LED bulb remains on during each cycle except its own time slot. The LED bulbs would illuminate the room sufficiently while each bulb can still broadcast its own identity. However, we were not able to modify all existing high power lamps in our room, so we demonstrate our system through a simple experiment with the same effect: while keeping the existing high power lamps as ambient light, we use one LED bulb to send one light pulse during its own time slot. The only difference is that during its time slot, the lamp *adds light to* instead of *subtracts light from* the overall light level.

In particular, we evaluate our system in detecting a first activity: opening a door. The transmitter node is placed inside a room (approx. 4 feet from the door) and the receiver node outside (approx. 2 feet from the door) in the hallway. The LED on the transmitter node turns on during the first time slot (100ms) and turns off for the remaining 700ms of each cycle. The experiment path is shown in Fig. 8. A person walks along 1 in the hallway, returns along 2 and finally enters the room along 3. The person then walks through the room and exits from the door on the other side. As the receiver is placed in the hallway, its readings might also be affected when people pass under the photosensor in the hallway. Hence it is necessary to be able to differentiate between a person walking in the hallway and entering the room. This experiment was repeated with 10 participants, and each person was asked to walk along the defined path 10 times.

4. PRELIMINARY RESULTS

4.1 Location differentiation

We demonstrate the feasibility of using shadow-based visible light sensing for detecting the presence of a human being indoors. We set up two nodes on the ceiling and made a person walk from one node to the other in increments of a single step (60 cm) to see if we can obtain an estimate of the position of the person in motion.

As the person moves towards the second node, there is a change in length of shadow that is cast which changes the sensor readings on the second node. In addition to the shadow that is cast, we find that when the person is standing right under the photodiode, the color of the persons' hair and clothes directly affects the amount of light being reflected to the sensor; when the person wears a white shirt the light intensity increases while it decreases when the person wears a black shirt. This can be seen in Fig. 10. Recall that voltage over the phototransistor is inversely proportional to the photocurrent (light intensity). This result shows promise in tracking humans indoors.

4.2 Detecting open doors

In this section, we evaluate the door opening detection experiment described in Section 3.2. We evaluated our prototype for two different types of doors: solid door, where no light can pass through, and glass door, where a portion of the door lets light through.

Equal Error Rate. Fig. 11 shows False Negative Rate (FNR) and False Positive Rate (FPR) for the two scenarios above with different thresholds. False negatives are the miss events, i.e., when the door is opened but not detected, while false positives are the false alarm events, i.e., when the door is not opened but the algorithm raises a detection event. In the case of solid door and hard floor, the equal error rate (FPR equals FNR) is 13% when the threshold is 12. In the case of glass door and hard floor, we achieve the equal error rate of 12% when the threshold is 4.

Note that the evaluation above is for when our system is synchronized with a low frequency signal (1.2Hz). When the system operates at high speed by synchronizing by 60Hz signal, there would be more cycles per each door opening events, thus more "dips" can

be observed. Therefore, the detection confidence of the system can be improved. We put this evaluation in the future work.

Other reflective material. We also tested our system with a dark carpeted floor. However, results show far less distinguishable "dips" during the time slot when the LED lamp sends its light beacon, causing a large error rate. The reason may be that the dark carpet absorbs and diffuses most of the light, and only reflects little light to the photosensor. We are working on methods to increase the light falling on the detection module, such as using lens on the photosensors or using an array of photosensors to increase resolution.

5. CONCLUSION & FUTURE WORK

In this paper, we proposed an activity and motion sensing system using the existing ceiling lighting, without requiring sensors on the floor or any device on the users. The system relies on a highly sensitive photosensing circuit to detect shadows on the floor and a time division flickering scheme to differentiate the light nodes causing shadows on the floor. We constructed a prototype and our first experiments showed promise in differentiating several positions of an occupant and in detection open doors.

Future work could expanded on this towards more general activity detection and tracking across rooms. This will likely require further improving the sensitivity of the system by increasing sensing resolution (i.e., adding more directed photodiodes per node). Second, future work should target enhancing the sensing circuit to be more sensitive and fast enough to capture beacons from fast flickering light sources. Third, it is desirable to improve the accuracy of the synchronization scheme to support more concurrent nodes.

6. REFERENCES

[1] C.-J. Huang, Y.-L. Wei, C. Fu, W.-H. Shen, H.-M. Tsai, and C.-J. K. Lin. Libeamscanner: Accurate indoor positioning with sweeping light beam. In *Proc. of the 2nd Int. Workshop on Visible Light Commun. Syst.*, pages 33–38. ACM, 2015.

[2] S.-Y. Jung, S. Hann, and C.-S. Park. Tdoa-based optical wireless indoor localization using led ceiling lamps. *Consumer Electronics, IEEE Trans. on*, 57(4):1592–1597, 2011.

[3] M. Keally, G. Zhou, G. Xing, J. Wu, and A. Pyles. Pbn: Towards practical activity recognition using smartphone-based body sensor networks. In *Proc. of the 9th ACM Conf. on Embedded Networked Sensor Systems*, pages 246–259. ACM, 2011.

[4] H.-S. Kim, D.-R. Kim, S.-H. Yang, Y.-H. Son, and S.-K. Han. An indoor visible light communication positioning system using a rf carrier allocation technique. *Lightwave Technology, Journal of*, 31(1):134–144, 2013.

[5] E. D. Lascio, A. Varshney, T. Voigt, and C. Perez-Penichet. Poster abstract: Localight - a battery-free passive localization system using visible light. In *2016 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–2, April 2016.

[6] J. Lei, X. Ren, and D. Fox. Fine-grained kitchen activity recognition using rgb-d. UbiComp '12, pages 208–211. ACM, 2012.

[7] L. Li, P. Hu, C. Peng, G. Shen, and F. Zhao. Epsilon: A visible light based positioning system. In *USENIX NSDI 14*, pages 331–343, 2014.

[8] T. Li, C. An, Z. Tian, A. T. Campbell, and X. Zhou. Human sensing using visible light communication. In *Proc. of the 21st Annual Int. Conf. on Mobile Computing and Networking*, pages 331–344. ACM, 2015.

[9] Z. Li, W. Chen, C. Li, M. Li, X.-Y. Li, and Y. Liu. Flight: Clock calibration using fluorescent lighting. In *Proceedings of the 18th annual international conference on Mobile computing and networking*, pages 329–340. ACM, 2012.

[10] K. Panta and J. Armstrong. Indoor localisation using white leds. *Electronics letters*, 48(4):228–230, 2012.

[11] Y. Wang, J. Liu, Y. Chen, M. Gruteser, J. Yang, and H. Liu. E-eyes: device-free location-oriented activity identification using fine-grained wifi signatures. In *Proc. of the 20th Annual Int. Conf. on Mobile Computing and Networking*, pages 617–628. ACM, 2014.

[12] J. Yun and S.-S. Lee. Human movement detection and identification using pyroelectric infrared sensors. *Sensors*, 14(5):8057–8081, 2014.

[13] C. Zhang, J. Tabor, J. Zhang, and X. Zhang. Extending mobile interaction through near-field visible light sensing. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 345–357. ACM, 2015.

[14] Z. Zhou, M. Kavehrad, and P. Deng. Indoor positioning algorithm using light-emitting diode visible light communications. *Optical Engineering*, 51(8):085009–1, 2012.