LIPS: A Light Intensity–Based Positioning System for Indoor Environments

BO XIE and KONGYANG CHEN, University of Chinese Academy of Sciences, and SIAT, Chinese Academy of Sciences
GUANG TAN, SIAT, Chinese Academy of Sciences
MINGMING LU, Central South University, China
YUNHUAI LIU, Third Research Institute of Ministry of Public Security, China
JIE WU, Temple University
TIAN HE, University of Minnesota

This article presents a Light Intensity–based Positioning System (LIPS) for indoor environments. The system uses off-the-shelf light-emitting diode lamps as signal sources and light sensors as signal receivers. The design is inspired by the observation that a light sensor has *deterministic sensitivity* to both the distance and incident angle of a light signal, an under-utilized feature of photodiodes now widely found on mobile devices. We develop a stable and accurate light intensity model to capture the phenomenon, based on which a new positioning principle, *Multi-Face Light Positioning*, is established that uses three collocated sensors to uniquely determine the receiver's position, assuming merely a *single* source of light. We have implemented a prototype on both dedicated embedded systems and smartphones. Experimental results show average positioning accuracy within 0.4m across different environments, with high stability against interferences from obstacles, ambient lights, temperature variation, and so on.

Categories and Subject Descriptors: C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design

General Terms: Design, Algorithms, Performance

Additional Key Words and Phrases: Light intensity, indoor positioning, LED sensor, smartphones

ACM Reference Format:
Bo Xie, Kongyang Chen, Guang Tan, Mingming Lu, Yunhuai Liu, Jie Wu, and Tian He. 2016. LIPS: A light intensity–based positioning system for indoor environments. ACM Trans. Sen. Netw. 12, 4, Article 28 (September 2016), 27 pages.
DOI: http://dx.doi.org/10.1145/2953880

1. INTRODUCTION

The prospect of indoor location-based services (LBS) and the poor performance of indoor global positioning services (GPS) have recently fueled much interest in indoor

Guang Tan's work was supported in part by NSFC Grant No. 61379135, National Key Research Program Grant No. 2016YFB0502202, and Shenzhen Scientific R & D Funds Grant No. CXZZ20151117161747567, No. KQCX20140520154115026, and No. JCYJ20140610151856733.
Authors' addresses: B. Xie and K. Chen, University of Chinese Academy of Sciences, and SIAT, Chinese Academy of Sciences; emails: {bo.xie, ky.chen}@siat.ac.cn; G. Tan (Corresponding author), SIAT, Chinese Academy of Sciences; email: guang.tan@siat.ac.cn; M. Lu, Central South University, China; email: ming.lu@gmail.com; Y. Liu, Third Research Institute of Ministry of Public Security, China; email: yunhuai.liu@gmail.com; J. Wu, Temple University; email: jiewu@temple.edu; T. He, Department of Computer Science and Engineering, University of Minnesota; email: tianhe@umn.edu.

DOI: http://dx.doi.org/10.1145/2953880

Fig. 1. LED lamps and light sensors. (a) LED lighting in an airport terminal. (b) LED lighting in a large warehouse. (c) An infrared LED lamp, 8W. (d) A high-power visible light LED lamp, 100W. (e) A smartphone with a light sensor (enclosed by red square).

positioning techniques. Among the numerous approaches to achieving such a service, the radio-frequency- (RF) based positioning technique has attracted perhaps the most attention, due to the wide deployment of WiFi access points. A major challenge faced by this approach is that RF is subject to serious multipath effects and is vulnerable to environmental interferences. This makes it difficult to establish an accurate propagation model that allows accurate distance estimation. Recent work has focused on the RF fingerprint approach, which obviates the need of a propagation model but then requires manual effort to establish a fingerprint database in support of mapping from signal strengths to positions. The fingerprint collection process is often laborious, leading to various research efforts to reduce the cost (e.g., Chintalapudi et al. [2010] and Yang et al. [2012]). Despite significant advances made in this direction, a fully automatic solution for general indoor environments has remained open.

In this article we explore an alternative approach, using visible or infrared (IR) light signal rather than RF, to achieving indoor positioning. The main advantage of light signal over RF is that the propagation of light is more predictable than that of RF. This is because on ordinary objects, such as walls and furniture, the light signal experiences only insignificant reflection, thus the signal at a receiver is subject to negligible multipath effect. This largely eliminates the uncertainty in characterization of received signal strength (RSS), laying a sound basis for further modeling and derivation of position.

Our system, called the Light Intensity–based Positioning System (LIPS), uses commodity light-emitting diode (LED) lamps (see Figure 1) as signal sources and light sensors available on mainstream mobile devices as signal receivers. A low-end microcontroller is used to make the LED lamp switch on and off at specified frequencies, so the light signals from sources of interest can be separated from ambient ones in the frequency domain. The recovered light signal strength on the light sensor reliably reflects the sensor's distance and orientation with respect to the lamp. This allows one to establish an accurate RSS model that paves the way to *fingerprint-free* positioning. From a practical point of view, the LIPS design could re-use the existing lighting infrastructure for indoor positioning in many public environments, such as airport terminals (Figure 1(a)), warehouses (Figure 1(b)), shopping malls, and hospitals, where

lamps are extensively deployed. These lamps can conveniently serve as positioning references, provided they are distinguishable in the frequency domain with different flashing rates and have known positions.

Compared to the RF-based approach, LIPS essentially trades off obstacle penetration ability for improved predictability of signal propagation. The design is centered around two questions: (1) *how accurate and stable is a light sensor in producing position-related information*, and (2) *how do we exploit that information for positioning while minimizing the line-of-sight limitation of light signal?*

For the first question, we conduct extensive experiments to study a light sensor's characteristics. The results show that a light sensor can be used to infer not only distance but also *angular* information from light signal. The angular information turns out to be very useful for obtaining the position of a light sensor. Based on this, we propose a *Multi-Face Light Positioning* (MFLP) method for positioning, which uses three collocated sensors to uniquely determine the receiver's position, assuming merely a *single* source of light. This single-source positioning method alleviates the concern of possible high deployment density of light sources, especially in a complex environment, where the light signal's line-of-sight restriction makes it expensive to circumvent obstacles. We show that in such an environment, MFLP requires far less than 1/3 (the theoretical ratio) of the lights as required by trilateration for the same degree of coverage, thus slashing the deployment cost to more accessible levels.

We present two designs for indoor positioning, one for a dedicated receiver and the other for smartphones. We evaluate the systems in various realistic environments and show that LIPS can produce positioning accuracy to 0.4m on the average, and that is stable across significant environmental variations.

The proposed positioning scheme is particularly suitable for indoor vehicles (e.g., Automated Guided Vehicles [WIKI-AGV 2016]) or robotics (e.g., automatic cleaning robot). It can be used alone to support LBS, or it can be used in combination with other positioning techniques, depending on specific environmental features. For example, when WiFi signal is available, a WiFi fingerprint database could be established while the LIPS receiver is being used. This database can in turn help the LIPS receiver position itself in places where line of sight is unavailable.

The next section establishes the light intensity model, Section 3 describes the positioning principle, Section 4 and Section 5 describe the implementation of two prototype systems, Section 6 presents the experimental results, Section 7 discusses possible extensions of the design, Section 8 documents related work, and Section 9 concludes the article.

2. LIGHT SENSOR CHARACTERISTICS

In order to examine the characteristics of a light sensor under different light frequencies and power supplies, we considered two typical types of LED lamps,

—An IR LED lamp, with an 8W power rating and a 12V direct current (DC) power supply (see Figure 1(c)), as well as an illuminating angle of 120°. It has a sensing range of 7.5m for the sensor, covering an open area of 130m^2 on the ground when hung on a 3m-high ceiling.[1]
—A visible light LED lamp, with a 100W power rating (see Figure 1(d)), and a sensing range of 30m, covering 2000+ m^2 in open space when hung at a proper height.

Each lamp has a small illuminating chip around the size of 1.5cm × 1.5cm and thus can be approximately viewed as a point source of light from a distance of a few meters. We focus on the IR lamp and only briefly report on the visible light one. In practice, IR

[1]At the time of this writing, a lamp costs about $30 US.

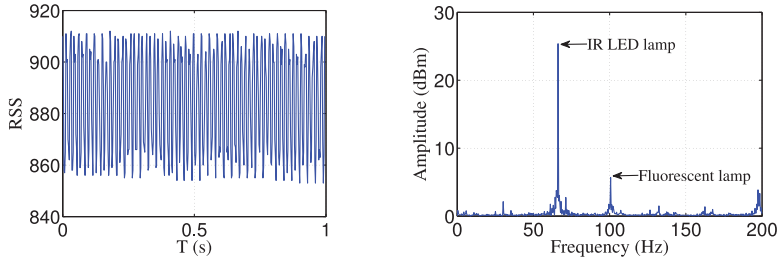

(a) Time domain (b) Frequency domain

Fig. 2. RSS in the time domain and frequency domain (excluding the DC component). The LED lamp flashes at a frequency of 65Hz, and the nearby fluorescent lamps flash at a frequency of 100Hz.

lamps can be deployed in environments where lighting devices already exist and extra visible lights are undesirable.

The light sensor used is an Intersil ISL29023,[2] which is used by the Samsung Omnia II GT-I8000 smartphone. We used a stand-alone sensor connected to a microcontroller for the experiment.

In reality, what the sensor receives is a mixture of the light signals from the sources of interest and the background, including daylight and artificial lights. (Note that visible lights also contain IR signals.) The ambient light could be so strong that the useful signals are completely overwhelmed. In order to isolate the useful light signals, we make the LED lamp switch on and off with a specified frequency using a low-end microcontroller. We then use FFT and inverse FFT to extract the signal strength at a particular frequency. Figure 2 shows a sequence of raw measurements of IR intensity and the FFT result. In this test, the IR LED lamp flashed at a frequency of 65Hz, and the nearby fluorescent lamps at the standard 100Hz, with daylight imposing a strong intensity on the sensor with reading about 850, an order of magnitude higher than the lamp's effect on the sensor. Two spikes corresponding to the IR lamp and the fluorescent lamps can be clearly identified from the figure, which allows us to recover the intensity of the IR signal faithfully. In the following, when we say a light signal sensed by a light sensor, we mean the signal after the FFT processing.

The model for light intensity, or receive signal strength (RSS), is similar to a conventional light propagation model (e.g., the model in Yang et al. [2013b] or the Lambertian model [Li et al. 2014]), with a focus on its stability in a realistic environment. The RSS, denoted by s, on a light sensor is mainly determined by three factors: the distance of the light sensor to the light bulb d, the incident angle μ of light, and the emitting angle ω. The incident angle μ plays an important role here, due to the working principle of the photodiode, which generates current under the striking of photons. When the flat contact layer is not perpendicular to the light, the energy with which the photons strike the contact layer decreases, and thus the received light energy drops. Normally, the larger the deviation to the perpendicular orientation, the more loss to the light intensity [Intersil 2012]. Finally, s decreases with increasing ω, following the characteristics of light emitting diodes that behave in the same way. We call the direction with $\omega = 0$ the *central ray*. Furthermore, we say that the lamp is *vertically oriented* when the central ray is vertical and *horizontally oriented* when the central ray is horizontal. Figure 3 shows the RSS model of a vertically oriented lamp, in which the tilted rectangle represents the sensor's surface, called its *sensing face*. Figure 4 shows the placement of the lamp of sensor in our real experiments.

[2]The ISL29023 is an integrated ambient and infrared light to digital converter [Intersil 2012]. The same family of sensors are also found on other phones such as the Motorola XT882.

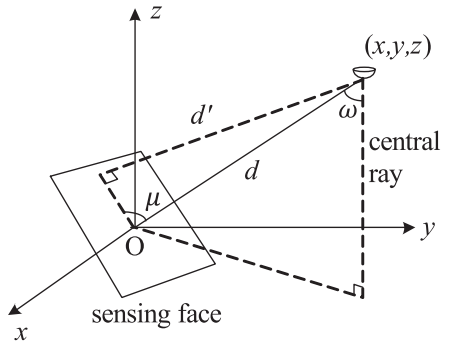

Fig. 3. RSS model of the light sensor. The light sensor's surface (sensing face) is centered at (0, 0, 0), while the lamp is located at (x, y, z). The RSS of the sensor, s, is a function of distance d, incident angle μ, and emitting angle ω.

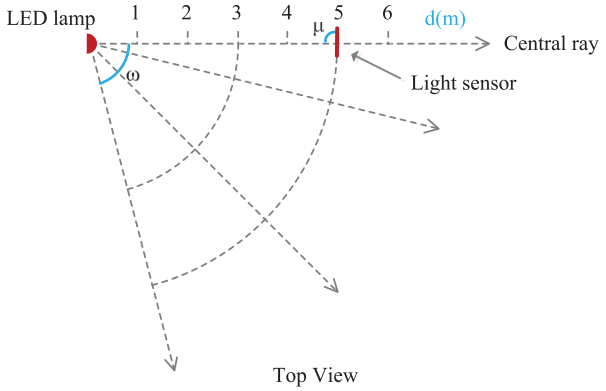

Fig. 4. Placement of the LED IR lamp and sensor in the experiments.

We need to determine three functions that respectively represent the influences of the three factors on RSS.

—$f_d(d)$, representing the impact of d on s, is obtained by varying d while fixing $\mu = \pi/2$ and $\omega = 0$, that is, making the sensing face perpendicular with the central ray.
—$f_\mu(\mu)$, representing the impact of μ on s, is obtained by varying the angle of the sensing face at a fixed d, with $\omega = 0$.
—$f_\omega(\omega)$, representing the impact of ω on s, is obtained by moving the sensor along a circle while keeping $\mu = 0$, that is, making the sensing face perpendicular with the emitting ray.

The light intensity function is then modeled as:

$$s = f_d(d) \cdot f_\mu(\mu) \cdot f_\omega(\omega). \tag{1}$$

Interference-free scenario. First, we examined the property of the light sensor in a dark room at nighttime, where no ambient light is present. The first column of Figure 5 shows the impacts of the three factors, d, μ, and ω, on the RSS. It can be seen that $f_d(d)$ quite closely follows the well-known inverse square law for light intensity. These could be accounted for by calibration but are ignored in our current design.) We can also see that $f_\mu(\mu)$ decreases with μ with fairly high predictability. The trends can

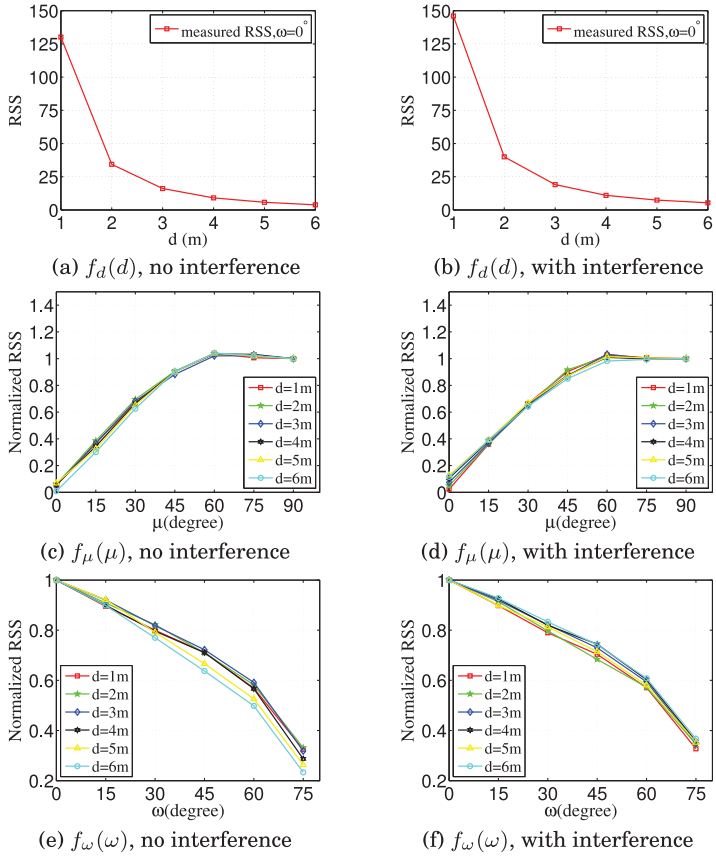

(a) $f_d(d)$, no interference (b) $f_d(d)$, with interference

(c) $f_\mu(\mu)$, no interference (d) $f_\mu(\mu)$, with interference

(e) $f_\omega(\omega)$, no interference (f) $f_\omega(\omega)$, with interference

Fig. 5. The light sensor's reading as a function of d, μ, and ω in two scenarios for the IR lamp.

be roughly captured by a sin function. For $f_\omega(\omega)$, the trend is also very deterministic, which can be modeled with a polynomial function.

Impact of ambient lights. This set of experiments was conducted in the morning, when the daylight imposed an IR intensity reading of 800 on the sensor. In addition, three fluorescent lamps were turned on, emitting periodic IR signals with frequency 100Hz. The intensity of ambient IR light was strong enough to overwhelm the RSS from the lamp at a distance of a few meters; however, the frequency domain treatment can successfully extract the component of light intensity that we are interested in. The second column of Figure 5 shows the impacts of the three factors, d, μ, and ω, on the RSS. It can be seen that the results in both interference and interference-free environments are quite consistent, with differences normally within 10% of each other.

Figure 6(a) shows the impact of ambient light intensity as experienced by the sensor at different times (e.g., nighttime and daytime) and at various places in a room. The maximum value 4,000 corresponds to the RSS near an open window around noon on a sunny day. It can be seen that the extracted IR intensity of the LED lamp remains relatively stable, with a standard deviation as small as 0.58.

Impact of light reflection. This experiment examines how light reflection from surrounding objects affects the RSS. In a dark room, we kept the lamp horizontally oriented, with its central ray parallel to a wall and the floor, additionally varying

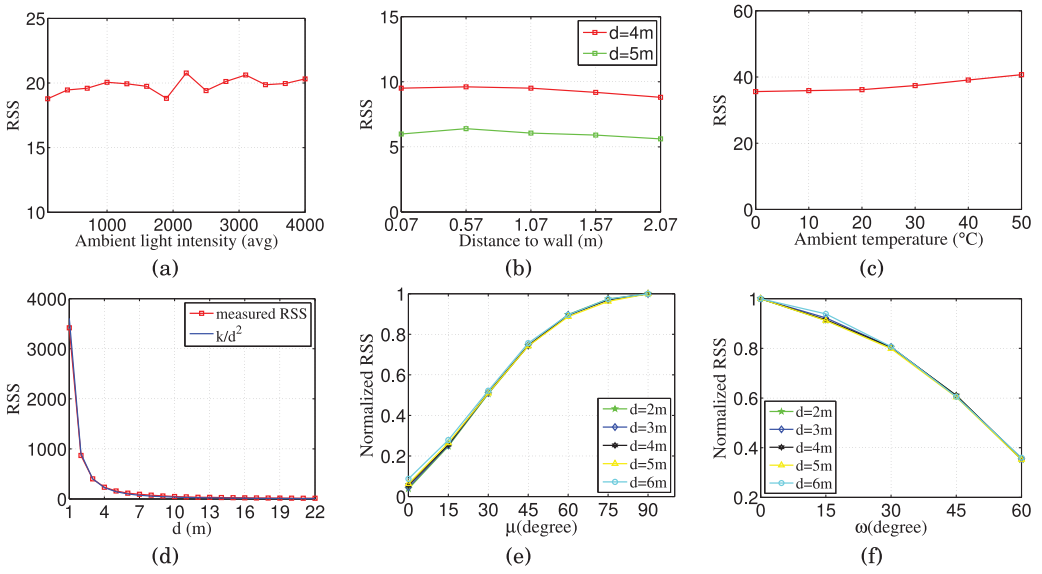

Fig. 6. Properties of a light sensor under various conditions. (a) Sensitivity to ambient light intensity. (b) Sensitivity to light reflection. (c) Sensitivity to ambient temperature. (d) Sensitivity to d under the visible light lamp. (e) Sensitivity to μ under the visible light lamp. (f) Sensitivity to ω under the visible light lamp.

the ray's distance to the wall. The light sensor was placed at a certain distance from the lamp with its sensing surface being perpendicular to the central ray. Figure 6(b) shows how the RSS changes with the distance to wall. It can be seen that the RSS experiences only insignificant changes. In our daily life, most materials (except glasses, polished metals, etc.) give no more than a few percentages of specular (i.e., mirrorlike) reflection [WIKI-DifRef 2016]; that is, most of the light, on hitting the surface of an object, is scattered in all directions, leaving only a small portion of reflected energy on the sensor. This explains the robustness of the RSS against wall reflection.

Impact of ambient temperature. Figure 6(c) shows how the RSS changes with ambient temperature. The low temperatures were produced by placing ice cubes around the sensor, and the high temperatures were generated by blowing at the sensor using a hair drier. We can see that the RSS increases with temperature, which agrees with the property of the photodiode reported in Intersil [2012]. The trend is very mild, suggesting that only small errors are introduced by the temperature factor. In addition, the trend is monotonic, so the error could be compensated for with simple calibration.

High-power LED lamp. We repeated the above experiments with the visible light lamp. Figure 6(d) shows the relationship between d and RSS, which closely matches the baseline curve of function k/d^2. Due to the much increased power, the lighting range extends to nearly 30m, and for the same d, the RSS is much higher than with the IR lamp. In this test, the lamp was horizontally oriented and placed within a narrow corridor (about 2m wide) surrounded by walls, the floor, and wooden boards, which presented complex conditions for light reflection. Compared to the low power IR lamp, the increased power causes a more noticeable reflection effect. However, the variability is still below approximately 10% of the baseline. In a practical system, the lamps will be hung on ceilings and sensor oriented primarily upward, so the sensor is unlikely to be exposed to as strong of a reflection. It is thus reasonable to assume that the reflection effect does not fundamentally invalidate our RSS model in typical environments.

Figures 6(e) and (f) show that the two functions $f_\mu(\mu)$ and $f_\omega(\omega)$ remain highly deterministic and consistent at different positions, though their particular forms differ from those of the IR lamp, due to the different photoelectric effects of the light sensor under IR and visible lights and also because of the different scattering effects of the lamp covers.

3. LIGHT POSITIONING PRINCIPLE

In this section, we describe two principles of light sensor positioning, one called MFLP, which is our emphasis, and the other following the classic trilateration method.

The basic idea of MFLP is to have three (or more) properly oriented sensors to collect signal strengths as well as the sensors' orientation measures. Along with the predefined RSS model, the measured data provide sufficient spatial constraints to locate the receiver. We call the top contact layer of the photodiode the *sensing face* of a light sensor and the plane containing it the sensor's *sensing plane*.

3.1. Multi-Face Light Positioning

Assume the considered sensing face is centered at the origin $O = (0, 0, 0)$ and the point source of light is located at $X = (x, y, z)$, where $x > 0, y > 0, z > 0, \mu \in (0, \pi/2), \omega \in (0, \pi/2)$. The corresponding sensing plane has the form $Ax + By + Cz = 0$, where A, B, C are determined by the sensor's tilt and heading. Then, the distance between X and O is

$$d = \sqrt{x^2 + y^2 + z^2},$$

and the distance between X and the sensing plane is

$$d' = \frac{|Ax + By + Cz|}{\sqrt{A^2 + B^2 + C^2}}.$$

Following Equation (1), let $f_d(d) = k/d^2$, $f_\mu(\mu) = \sin(\mu) = d'/d$, and $f_\omega(\omega) = f_\omega(\arccos(z/d))$, where $f_\omega(\cdot)$ is a monotonically decreasing function of ω (hence z). Therefore,

$$s = \frac{k}{d^3} \frac{|Ax + By + Cz|}{\sqrt{A^2 + B^2 + C^2}} f_\omega\Big(\arccos\frac{z}{d}\Big), \tag{2}$$

where $d = \sqrt{x^2 + y^2 + z^2}, x > 0, y > 0, z > 0, k > 0$, and s is a real number between 0 and some maximum reading value $s_m > 0$.

THEOREM 3.1. *When no measurement errors occur, three linearly independent sensing planes that pass through the origin and that satisfy the RSS model as specified by Equation (2) determine a unique solution of $X = (x, y, z)$.*

The system of equations generated by the mentioned three sensing planes is a high-order and nonlinear one, whose properties are in general not easy to obtain. Fortunately, the structure of A, B, C is simple enough to enable reduction among the equations, which makes it possible to establish an exact relationship between the solvability and the linear independence property of (A_i, B_i, C_i). The proof is provided in Appendix A.

Theorem 3.1 suggests that if we can create three linearly independent sensing planes on a receiver, and that these sensors can simultaneously "see" the light source (i.e., in line of sight), then one can determine a position of the light source. With a bit of coordinate transformation, we can determine the position of the receiver provided the position of the light source.

When measurement errors exist, three linearly independent sensing faces may not lead to a solution. In this case, what we look for is a least-squares solution that

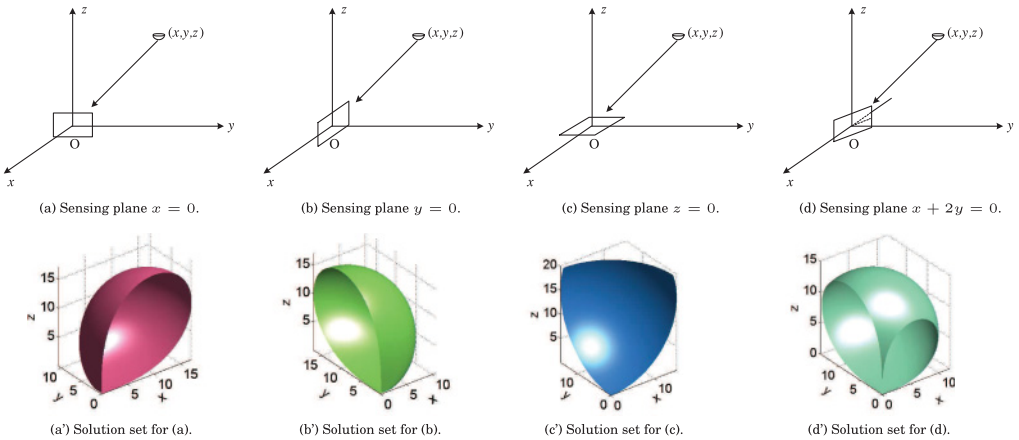

Fig. 7. The location solution sets of the sensor for different sensing faces.

minimizes the sum of the squares of the errors between each measured s and the calculated s from the corresponding equation.

3.2. Why Linear Independence of Faces

Since the geometric structure of the problem is not immediately intuitive, we now give a hypothetical example of MFLP to illustrate the necessity of the faces being linearly independent, which provides a key guideline for our system design. For the purposes of demonstration, we assume a simplified RSS model, in which $f_d(d) = 1/d^2$, $f_\mu(\mu) = \sin(\mu)$, and $f_\omega(\omega) = \cos(\omega)$. Then Equation (2) can be rewritten as

$$s = \frac{z|Ax + By + Cz|}{\sqrt{A^2 + B^2 + C^2}(x^2 + y^2 + z^2)^2},$$

where $x > 0, y > 0, z > 0$.

Consider four sensing faces centered at the origin, whose sensing planes are $x = 0$, $y = 0$, $z = 0$, and $x + 2y = 0$, as shown in Figures 7(a)–(d), respectively. The lamp is located at $X = (10, 10, 10)$ and imposes light intensity s_1, s_2, s_3, and s_4 on the sensors, respectively. Given the known position of the lamp, the theoretical light intensities should be $s_1' = s_2' = s_3' = 1/900$, and $s_4' = 1/300\sqrt{5}$. Assume no errors occurring from the light propagation and measurement, the RSS seen on the sensors should be $s_i = s_i'$.

Now we choose the first three sensing faces, $x = 0, y = 0, z = 0$, which are *linearly independent*. These faces lead to a system of equations $h \cdot z/(x^2 + y^2 + z^2)^2 = 1/900$, where $h = x$, y or z. Each of these equations will generate a solution set for the sensor, as depicted by the curved surfaces in Figures 7(a′)–(c′), respectively. Figure 8(a) shows the intersection of these three solution sets. It turns out that these surfaces intersect at a single point (10, 10, 10), which matches the true location.

Next, consider an alternative set of three sensing faces, $x = 0, y = 0, x + 2y = 0$, which are *linearly dependent*. Figure 8(b) shows the intersection of their corresponding solution sets. Differing from the first case, these sets do not intersect at a single point but instead produce a curve segment (only partly shown due to the blocking of surfaces), meaning infinitely many valid location solutions for the sensor. This comparison shows why linear independence of the sensing faces is necessary for a unique solution of the sensor's location.

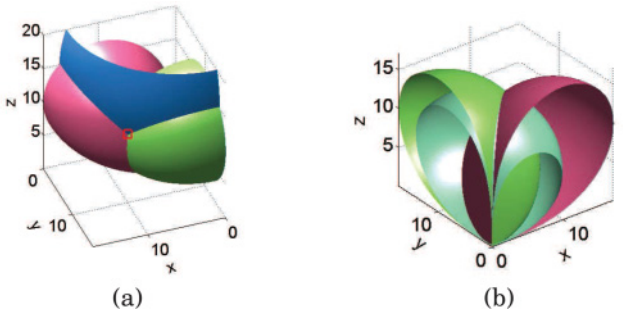

Fig. 8. The uniqueness of position solution. (a) The solution sets for the three sensing faces $x = 0, y = 0, z = 0$ intersect at a single point (red box). (b) The solution sets for the three sensing faces $x = 0, y = 0, x + 2y = 0$ intersect at a curve (partly shown).

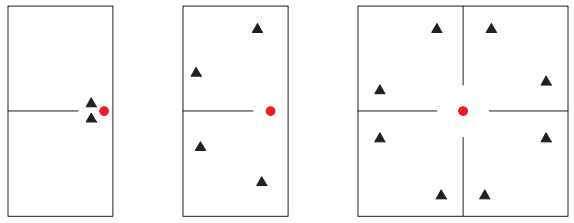

Fig. 9. Trilateration requires many more lamps than MFLP does for a full coverage of the environment. Small triangles and red dots represent the lamps. In all cases, MFLP needs only a single lamp (red dot).

3.3. MFLP vs. Trilateration

The trilateration method assumes a single sensor, with three (or more) lamps as position references [Li et al. 2014]. To reduce unknowns, the sensing face is assumed to be placed horizontally. Assume the sensor is located at (x, y, z), and there are three non-collinear lamps at distinct positions $(x_i, y_i, z_i), i = 0, 1, 2$; then, the following system of equations can be established with which one can solve for the solution:

$$\frac{k}{d_i^3} \cdot |z - z_i| \cdot f_\omega \left(\arccos \frac{z - z_i}{d_i} \right) = s, \tag{3}$$

where $d_i = \sqrt{(x - x_i)^2 + (y - y_i)^2 + (z - z_i)^2}$, $k > 0$, and s is the RSS.

While simplifying the design at the receiver side, trilateration shifts cost to the transmitter side. In theory, trilateration requires 3 times as many lamps as MFLP does. In practice, however, the difference can be much higher. Figure 9(a) shows a scenario where two adjacent rooms share a wall and a relatively narrow gate. With trilateration, every point needs to see at least three lamps. For a minimum deployment cost, the three lamps could be deployed near the gate. However, trilateration further requires that the lamps not be close by or collinear, or huge errors or position ambiguities may arise. This makes it necessary to place at least two additional lamps in each room (Figure 9(b)). Thus, trilateration ends up using 5 times as many lamps as used by MFLP. Figure 9(c) further shows a case where trilateration needs 9 times as many lamps. In a real-world environment, obstructions may appear in different forms, but a similar comparison can be drawn between the two approaches in terms of light coverage and deployment density.

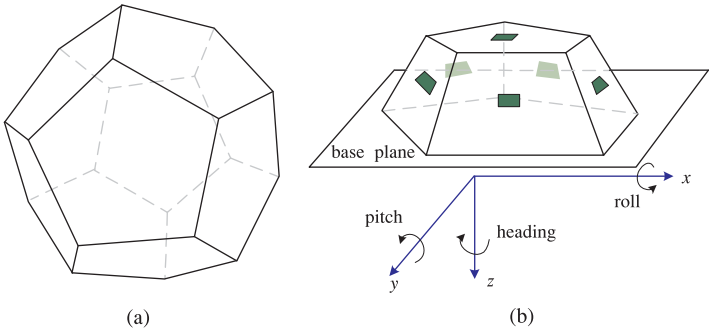

Fig. 10. (a) Dodecahedron (conceptual model). (b) Half dodecahedron (implementation model) on which six light sensors, represented by green blocks, are installed.

Therefore, when the positioning system is deployed from scratch and lamp deployment cost is a primary concern, the MFLP approach appears to be a more economical solution. On the other hand, when there already exist dense lamps (as in some public places such as shopping malls, airport terminals, etc.) that can be re-used for positioning, and when a small-sized receiver is preferred, the trilateration approach might be the choice. Since the trilateration approach is well studied and understood, we shall concentrate on the MFLP approach in the following sections.

4. LIPS RECEIVER DESIGN

In this section we describe a design for a dedicated LIPS receiver based on the MFLP principle.

4.1. Number and Placement of Light Sensors

We say that a sensor face can "see" a point p if there is a line of sight between p and all points on that sensing face. Although a receiver can be positioned with only three sensors at particular places, visibility of those sensors to a lamp may be lost when the receiver is moving around. Thus, we need more than three sensors to support positioning everywhere (assuming full coverage of light on the ground). Toward that goal, we need to answer two questions: How many sensors do we need, and how do we place them?

For convenience of implementation, we focus on a regular polyhedron framework on which the sensors are to be placed. Our choice is a dodecahedron model, as shown in Figure 10(a), in which 12 regular pentagonal faces each host a sensor. This model possesses three desirable properties:

(1) *Tri-face visibility*: Any point p in the space beyond a short distance from the dodecahedron can see at least three faces of that dodecahedron;
(2) *Tri-face linear independence*: Any three of the faces seen by the above-mentioned point p are linearly independent;
(3) *Minimal faces*: Among all regular polyhedra, a regular dodecahedron has the smallest number of faces that satisfy the above two properties.

Theorem 4.1 gives a more formal description of the tri-face visibility property. The theorem can be proved with basic trigonometric operations and is thus omitted here.

THEOREM 4.1. *Assume the edge length of a regular dodecahedron centered at the origin is a, then, for an arbitrary point at a distance D from the origin, at least three*

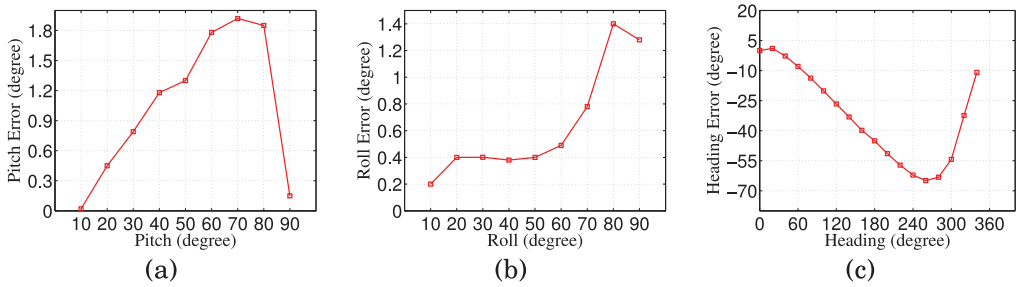

Fig. 11. Errors of pitch, roll, and heading, produced by an acceleration sensor and a magnetic sensor.

faces of the dodecahedron can see it if

$$D \geq \left(\sqrt{1+\frac{2}{5}\sqrt{5}}+\frac{1}{2}\sqrt{\frac{5}{2}+\frac{11}{10}\sqrt{5}}\right)a \approx 2.49a.$$

In LIPS, a is at the order of a few centimeters, so the dodecahedron faces can be viewed as roughly passing through the origin from the perspective of the lamp. Theorem 4.1 means that if we place 12 sensors along the faces of a regular dodecahedron, then the receiver can always be positioned, regardless of the receiver's orientation.

The tri-face linear independence property can be easily verified by examining the plane coefficients of the faces of a dodecahedron. Finally, the minimal faces property can be proved by excluding the regular polyhedra with fewer faces. For example, a cube-aligned placement of sensors may be able to position a receiver sometimes, as demonstrated in Section 3.2, but it is easy to pick a point in the space from which only a single face is visible.

To further reduce the cost, we make a simplification to the conceptual model by employing only one half of the dodecahedron, which is fixed on a *base plane* (Figure 10(b)) of the receiver. This way we need only six sensors attached to the six exposed faces. This does not affect the positioning ability as long as lamps are hung above at ceilings and the half dodecahedron is oriented upward.

4.2. Sensing Plane Coefficients

The coefficients of a sensing plane, namely A, B, and C, are obtained from acceleration and magnetic sensors. Following the aircraft convention, we use three attitude angles, *pitch*, *roll*, and *heading*, to describe a receiver centered coordinate system. Figure 10(b) shows the coordinate system, in which x, y, and z are defined as forward/right/down based on the right-hand rule, and the three attitude angles are referenced to the local horizontal plane which is perpendicular to the Earth's gravity.

Denote the pitch, roll, and heading by θ_p, θ_r, and θ_h. The first two angles can be obtained with an acceleration sensor, which produces three components of the gravity along the x, y, and z axes. Comparing these components against the acceleration of gravity can give the two angles. Figures 11(a) and (b) show that these two angles can be measured with accuracy to 2°. A standard way of obtaining θ_h is using an electronic compass. However, a compass is very susceptible to interferences. To confirm this, we rotate a compass containing a magnetic sensor and obtain the heading error with respect to a calibrated electronic compass. Figure 11(c) shows that the raw measurement of heading can vastly deviate from the true value, with errors up to 60 degrees. We discuss how to calibrate a compass in the next section.

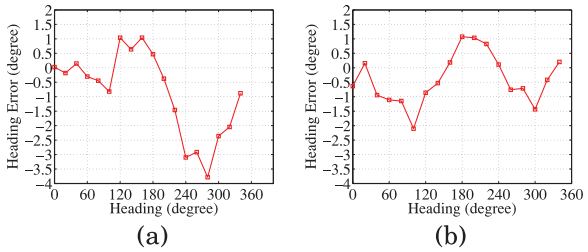

Fig. 12. Heading errors under auto-calibration. (a) Half dodecahedron calibration. (b) Circular calibration.

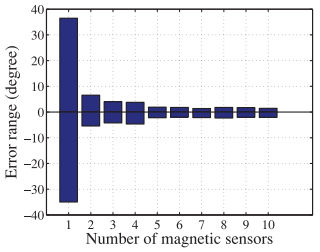

Fig. 13. Heading accuracy vs. number of magnetic sensors.

The three angles θ_p, θ_r, and θ_h entirely determine the orientation of the sensing plane and thus the coefficients A, B, and C. We use the rotation matrix [WIKI-RotMatrix 2016] to convert the angles to A, B, and C.

4.3. Accurate Heading Measurement

Magnetic sensor is known to be vulnerable to environmental interferences, because the Earth's magnetic field is a weak signal. A mobile receiver carried around may experience distinct distortion patterns at different locations. This property is exploited by Chung et al. [2011] to enable receiver positioning but causes trouble for a system that needs accurate heading information in real time. The general approach to auto-calibration of a compass requires multiple types of sensors to provide redundant measurements to correct heading errors. The redundant information can be from optical trackers [Gottschalk and Hughes 1993], inertial sensors [Hoff and Azuma 2000], or visual analysis [You et al. 1999]. In our context, these techniques are unsuitable because they either involve extra infrastructure or do not provide bounded errors in a long period of time [Hoff and Azuma 2000].

LIPS uses a new auto-calibration technique for heading measurement. The idea is inspired by the standard manual calibration method, in which sufficient 3D rotation or several full-round 2D rotations are performed to collect magnetic field strengths. Ideally, the readings along the three axes should form a sphere centered at the origin. In practice, environmental interference will distort the sphere, resulting in a tilted ellipsoid [STMicroelectronics 2010]. Given a set of manually collected magnetic field data, the least-squares fitting method can be used to discover the parameters of the ellipsoid, which are then applied to correct the raw measurements.

LIPS avoids the manual operation by using multiple magnetic sensors that are placed in a spherical or circular layout. These sensors can produce a number of magnetism readings at once, thus providing a sparse sampling of the needed magnetic data. For 3D measurement, the spherical layout is approximated with a half dodecahedron. In this experiment, we manually collected the readings along the six faces using the same magnetic sensor. Figure 12(a) gives the heading errors with the half dodecahedron calibration at the same position that produced the errors in Figure 11(c). We can see that the original error of 60+° is now reduced to around 4°.

To simplify the hardware design, we can place multiple magnetic sensors in a circular pattern when the receiver is placed horizontally. Before performing the collective calibration, the magnetic sensors are individually calibrated to achieve a consistent effect by rotating the board in an interference-free environment and collecting the readings from individual sensors. On a horizontal plane, the reference model of the magnetic field strength should be an ellipse instead of an ellipsoid, meaning that the fitting process deals with fewer parameters. As a result, the heading accuracy after calibration

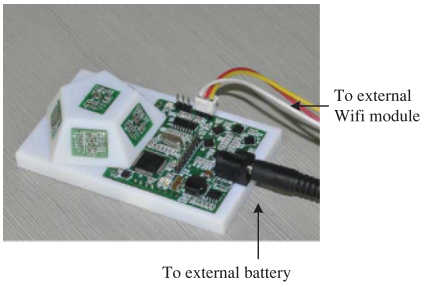

Fig. 14. The CompEye receiver, sized 9.5cm×5.7cm. The six light sensors are embedded in a half-dodecahedron model and are connected to the main board from within the model.

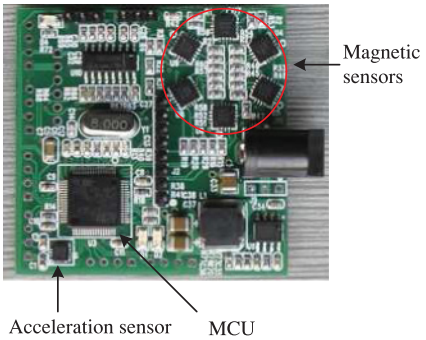

Fig. 15. Main board of LIPS receiver, sized 5cm×5cm.

will be more accurate than with the half dodecahedron calibration. Figure 12(b) shows that the heading error has now dropped to around 2° with six sensors.

The number of magnetic sensors used should strike a balance between measurement accuracy and cost. We perform a manual two-dimensional (2D) rotation of the magnetic sensor at different places, obtaining varying numbers of samples (which simulate multiple sensors) to correct the heading measurement. Figure 13 shows the result at a spot near a desktop, where the machine imposes a great interference to the magnetic field around the sensor. In the figure, the ranges of heading error for different numbers of sensors are shown. We see that the errors generally decrease with more sensors, and the trend flattens at six sensors. A similar result is observed at other places in our office. Thus, our design finally uses six sensor for heading measurement.

5. PROTOTYPE IMPLEMENTATION

We have implemented a dedicated LIPS receiver and a smartphone-based receiver. This section details the implementation.

5.1. The CompEye Receiver

Figure 14 shows the dedicated LIPS receiver, called CompEye (compound eye), for position-tracking applications. The receiver consists of two parts: the main board, shown in Figure 15, and the light-sensing component. The former is a circuit board integrating an STM32F103RC microcontroller (MCU), six AKM8975 magnetic sensors, and an ST LIS33DE accelerometer. The magnetic sensors are arranged along a circle for heading calibration. The light-sensing component comprises six ISL29023 light sensors fixed on the surfaces of a half-dodecahedron model; each sensor is linked to the main board individually. In the prototype design, the receiver is powered by an

external battery and uses a wireless serial adapter for transmitting data to a server. In the future, the power module will be replaced by a lithium battery like one used by a mobile phone, and the WiFi module will be integrated into the main board.

The MCU samples each light sensor at a rate of 640Hz and performs the FFT transformation to extract the light intensities of surrounding light sources (identified by peaks in the frequency domain). The MCU also samples each magnetic sensor at a rate of 20Hz and performs calibration, first individually and then collectively, to obtain the current heading of the receiver. The tilt and pitch are calculated from the reading of the acceleration sensor. The MCU maintains a sliding-window for each sensor and performs the above calculation every Δ seconds (e.g., $\Delta = 0.3$) and sends the light intensities, heading, roll, and pitch to the server. The server stores a digital map of the physical environment and of the deployed lamps, with which it solves for the receiver's position using a least-squares optimization algorithm.

5.2. Smartphone

We used a Samsung Omnia II GT-I8000 smartphone to evaluate the design. We installed an Android system with Linux kernel 2.6.32, which allowed us to configure the light sensor in the driver to sense infrared light instead of the default visible light.

After the configuration, the phone can give correct readings of IR intensity, as verified against a stand-alone sensor. However, under periodic sampling, the sequence of readings produced by the Android interface contains many uneven gaps in time, making the frequency analysis difficult. This is because the OS kernel contains a routine to smoothen sensed data and filter out readings with only small changes. This treatment seriously affects frequency analysis. We modified the `input_defuzz_abs_event()` function in `input.c` to disable the smoothing procedure for the light sensor and re-built the kernel, after which the sensor started working normally.

The smartphone can be positioned with the trilateration method, since there is only one sensor. We evaluate its effectiveness in the next section.

6. EXPERIMENTS

We conduct experiments in three static scenarios and one mobile scenario. We mainly consider IR LED lamps as they can be incrementally deployed without interfering with the existing lighting design of an environment.

(1) In this scenario, labeled *Empty-room* (Figure 16(a)), an IR LED lamp is hung on the ceiling of an empty room, where there are few obstructions and the floor is reflective. We selected 50 points of interest with rough uniformity across the sensing area for testing.
(2) The second scenario, *Office* (Figure 16(b)), is a crowded office environment with dense cubicles, which create complex conditions for light reflection. A single IR LED lamp was again used, and 50 points, both on the floor and on the desk, were selected for positioning.
(3) The third scenario, *Three-lamps* (Figure 16(c)), is similar to the second one, except that there were three lamps used that flashed at rates of 55Hz, 65Hz, and 75Hz. We deliberately made the lamps' sensing areas overlap more than necessary to examine the effect of increased face exposure to light signals. In this scenario, 98 points of interest were chosen, with a bias to the overlapped areas. Throughout the test, the receiver had its base plane placed horizontally on either the floor or the desk.
(4) The fourth scenario, *Three-lamps-mobile* (Figure 16(d)), consists of three lamps placed in a similar layout as in the *Three-lamps* scenario. This experiment is conducted in a different building, which creates more diversity of experimental

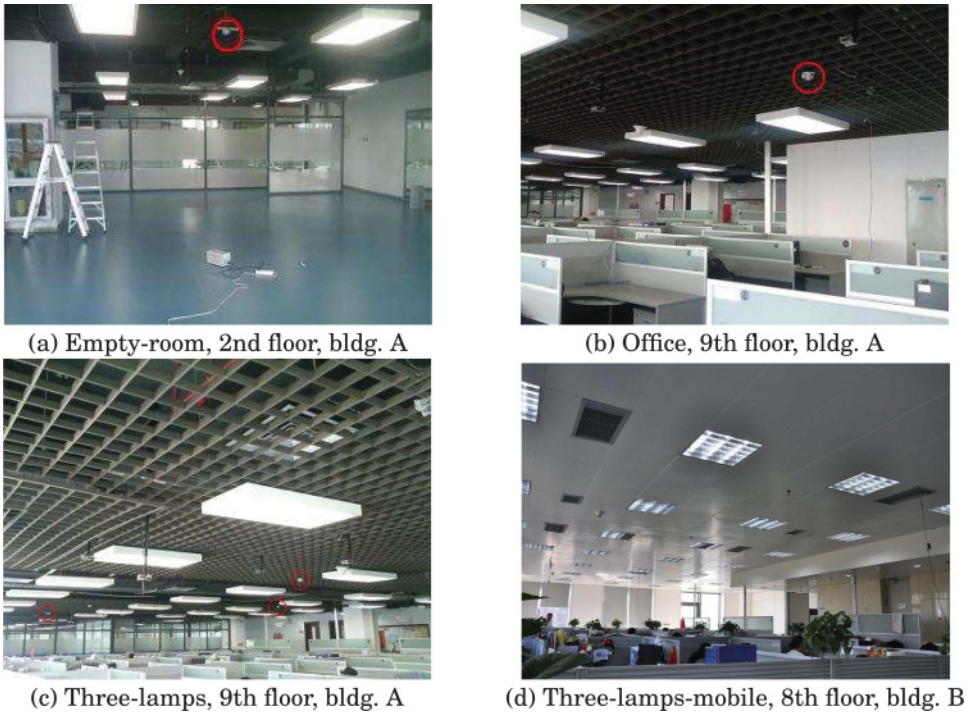

Fig. 16. Experimental scenarios for LIPS prototype. (a) Empty-room: An empty room with a single IR LED lamp. (b) Office: An office environment with a single IR LED lamp. (c) Three-lamps: An office environment with three IR LED lamps flashing at rates of 55Hz, 65Hz, and 75Hz. (d) Three-lamps-mobile: Another office environment with three IR LED lamps flashing at rates of 55Hz, 65Hz, and 75Hz; the receiver is being moved.

conditions. The receiver is moved along three straight lines while position data are collected. We compare the collected traces against the ground truth to check the system's mobile performance.

6.1. Position Accuracy

CompEye receiver. When there is no obstruction, a CompEye receiver can normally find at least three faces visible to a lamp in range; normally it can find four. We have found that the additional face generally does not improve the positioning result, because of its small incident angle μ that produces a very low RSS, potentially introducing an increased error. Therefore, for a particular lamp, we always choose the three faces with the strongest RSS for position calculation. When the receiver is covered by two or more lamps, it chooses the lamp that imposes a stronger average RSS (over three faces) for positioning.

Figures 17(a)–(c) show the maps of the various environments as well as the positioning results. In the figures, the red dots represent lamps, small dots true positions, and a line segment connects a true position to its corresponding result of positioning. The length of a line segment is thus proportional to the positioning error. In all cases, the median and average errors are below 0.4m (Table I). The consistence is also reflected in the small standard deviation. Notice that quite different environmental conditions, including floor/wall reflection, ambient temperature, and ambient light intensity (implied by time of day), and so on. are contained in these scenarios. These variations

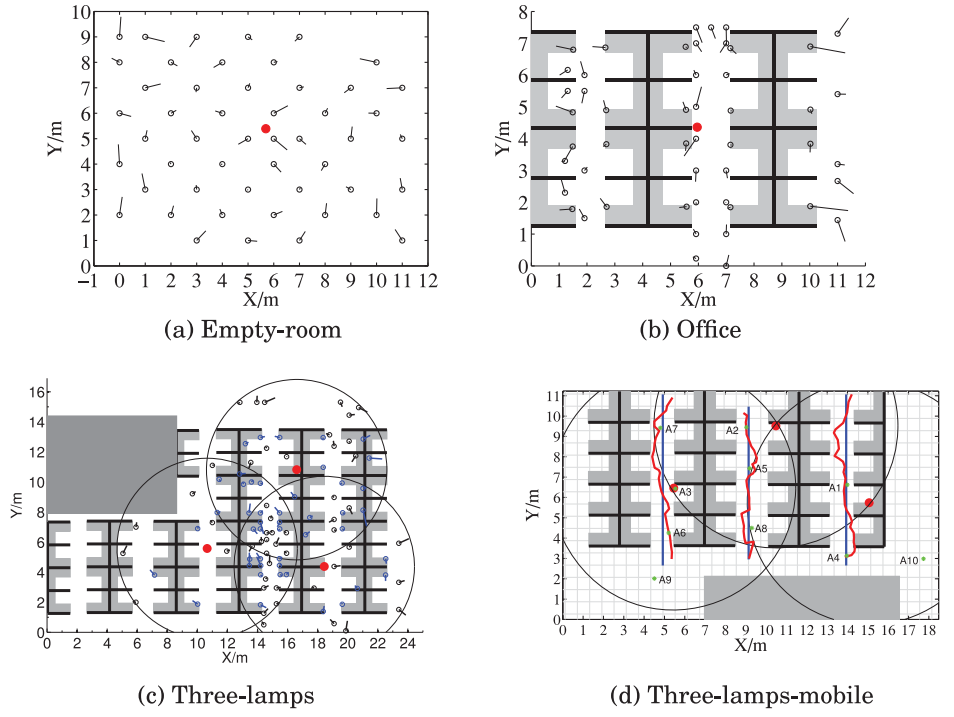

Fig. 17. Positioning results of a CompEye receiver. Gray areas represent desks or obstacles, red dots lamps, and small circles test spots. (a) Empty-room, size 12m×10m, (b) Office, size 12m×10m, (c) Three-lamps, size 24m×16m. (d) Three-lamps-mobile, 18m×12m.

Table I. Position Error Statistics in Different Scenarios. Note That the Smartphone Case Assumes a Trilateration Method, Which Requires a Denser Deployment of Lights Than the Other Cases

	Error statistics (m)			
Test cases	Median	Mean	Max	Stdev
Empty room, CompEye	0.36	0.39	0.79	0.20
Office, CompEye	0.33	0.36	0.73	0.20
Three-lamps, CompEye	0.32	0.32	1.08	0.20
Three-lamps, smartphone	0.39	0.44	1.05	0.29

cause unnoticeable differences in positioning accuracy, suggesting that LIPS is robust to environmental differences.

Figure 17(d) shows the mobile performance of the LIPS receiver. In this experiment, the receiver is placed on a small cart moving along the central lines of the three corridors. The positions of the receiver are generated at approximate intervals of 0.3s. The red lines show the position traces of the receiver; the cells are of size 0.6m × 0.6m. It can be seen that the positioned traces stay within 0.6m of the true trace lines, with an average much lower than 0.6m.

Smartphone. We also experimented with the trilateration method using a smartphone. A total of 33 points were randomly picked with rough uniformity in the intersection area of the three lamps. The phone was horizontally placed when collecting the RSS. The error statistics are given in the last row of Table I. It can be seen that the average error now goes up to 0.44m, which is worse than that of the CompEye receiver.

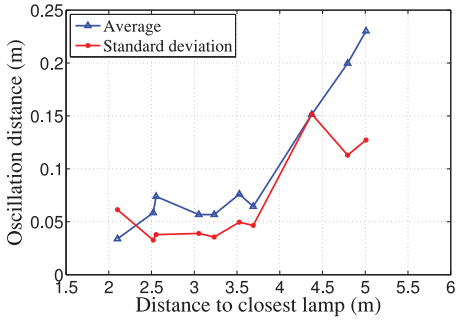

Fig. 18. Oscillation distance of ten test spots in the Three-lamps-mobile scenario.

This is mainly because of the absorbtion of light by the phone screen and the blocking of phone body, especially when the incident angle is small.

6.2. Position Stability

It is normal for a positioning system to generate position oscillation due to the instability of signal propagation, environmental interference, and various errors in the system. The degree of oscillation is an important impact factor for the user's experience. Though multiple position samples could be collected to increase stability, in delay sensitive applications oscillation is still likely to cause feelings of uncertainty and visual discomfort.

Given a spot in the field, we define the positioning system's *oscillation distance* as the average distance of the produced position samples from the centroid of those samples. In the Three-lamps-mobile scenario, we choose 10 test spots A1, A2, ..., A10 with different distances to their closest lamps, and, for each test spot, approximately 300 position samples are collected to obtain an average and a standard deviation. Figure 18 depicts the results against the distance to the closest lamp. First, it can be seen that the averages remain below 0.25m, indicating that the positions are quite stable. Second, the oscillation distance generally increases with the spot's distance to closest lamp, due to the decreased signal-to-noise ratio at a larger distance.

6.3. Impact of Lamp Density

In the three-lamps scenario, there are 98 points for testing, of which 61 are covered by a single lamp, 30 by two, and 7 by three. Thus, a receiver may have more than three faces visible to a lamp, and a sensing face may see multiple lamps. Let m denote the number of RSS readings a receiver collects. Since the receiver records only the three highest RSSs with respect to an individual lamp, and there are only three lamps, we have $3 \leq m \leq 9$. In fact, each RSS will result in an Equation (2), and we have already shown that three such equations can lead to a position. Here we want to see if using more than three RSSs will be beneficial. To that end, we pick m' highest RSSs for each point p, where the receiver has $m \geq m'$, and establish a system of m' equations, by which we solve for positions.

Figure 19 shows how position errors are affected by the number of RSS readings, m. It can be seen that in general, increasing m results in improved accuracy, reducing the average error to below 0.2m. In particular, the maximum error and standard deviation of error drops quite sharply with m, suggesting a significant benefit of multiple lamps in reducing outliers and improving positioning stability. The declining trend starts flattening after $m > 7$, implying a limiting accuracy of around 0.19m under the present RSS model.

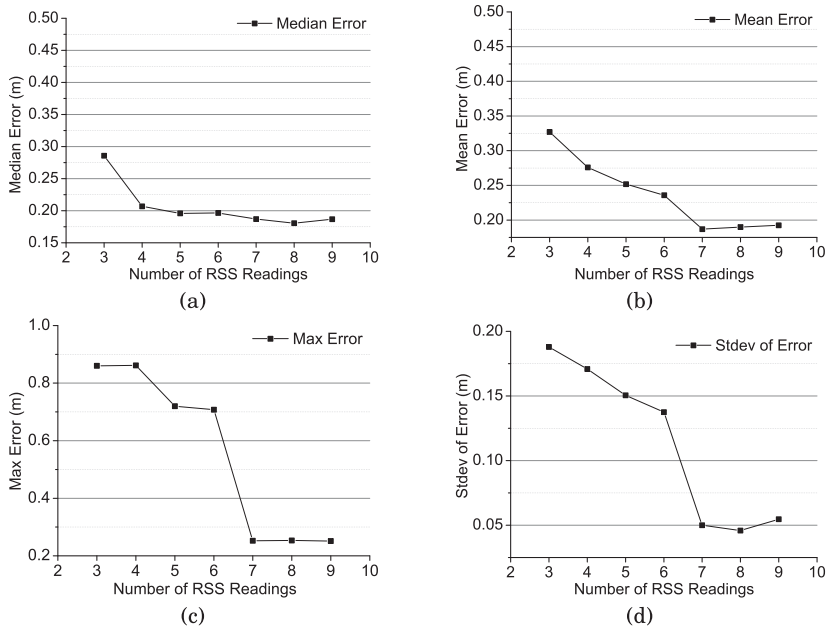

Fig. 19. Positioning errors of the CompEye receiver for various numbers of RSS readings.

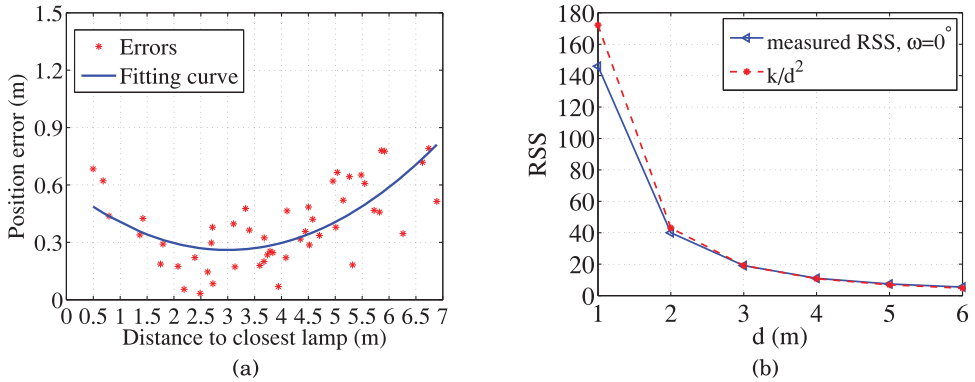

Fig. 20. Impact of distance on position errors. (a) Position error vs. distance. (b) Fitting error of $f_d(d)$.

Note that the case of more than three faces provides a generalization of MFLP, which essentially mixes the principles of MFLP and trilateration.

6.4. Impact of Distance

Intuitively, the farther the receiver is away from the lamp, the lower the signal-to-noise ratio (SNR) and thus the larger position error produced by the system. Figure 20(a) shows the position error in the empty-room scenario for different distances between the receiver and the lamp's projection on the ground, along with a fitting curve. It can be seen that for distances above 3m, the error in general increases with distance, but for a closer range, there is an opposite trend. This is due to the fitting error of the $f_d(d)$ function. Figure 20(b) shows the measured RSS and the fitting function k/d^2. Due to various errors, the fitting is not perfect, leaving some gaps for distance smaller than

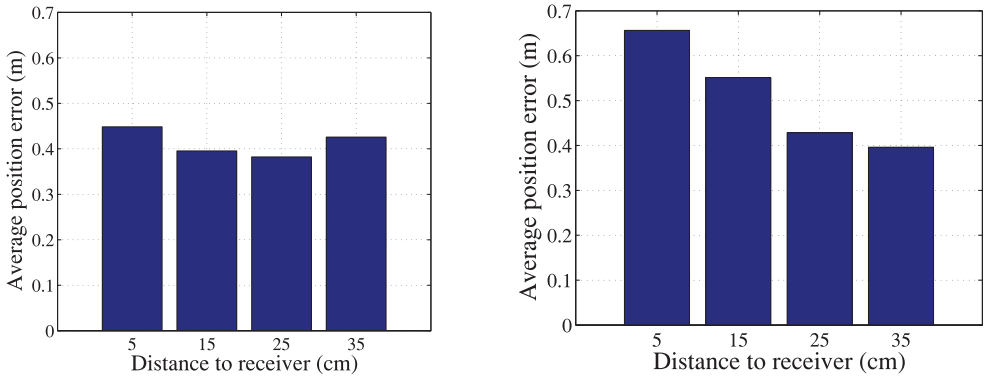

Fig. 21. Impact of light reflection on positioning accuracy.

2m. This discrepancy turns out to be influential enough to make the final position error increase noticeably. The impact of reduced SNR, as well as increased ω angle (which causes an extra error), starts to manifest itself for distances beyond 3m, yielding a growing trend of error. Note the fitting effect is determined mainly by the parameter k, and we chose such a particular k because it benefits most distances, thus minimizing the error for a large part of the field. If an alternative k were chosen in favor of small distances, then the tail of the curve will shift down, and the positioning performance beyond 2–3m will suffer.

6.5. Impact of Reflection

In this section we examine LIPS's performance under highly reflective conditions. We assume the single-lamp, Office scenario, and use a piece of glass and a piece of mirror, both of size about 40cm $\times$ 50cm, as reflectors. A reflector is placed vertically behind the receiver and has its reflective surface perpendicular to the line $p_l p_r$, where p_l is the lamp's projection on the ground and p_r the receiver's location. The distance between p_r and the reflector is denoted by d_r. The receiver is places at 20 different places of the field for each d_r.

Figure 21(a) shows that, despite the presence of a glass reflector, the accuracy remains around 0.4m, with no definite impact from d_r. The most seriously affected sensor of the receiver is the top one, which experiences an average RSS increase of 7.5%, 8.4%, 6.8%, and 5.6% for $d_r = 5, 15, 25, 35$cm, respectively. There are two other sensors whose RSS measurements are used in position calculation, but they face the lamp, and thus receive little (if any) reflected light, due to the blocking of the receiver's body.

Using a mirror as the reflector leads to quite different results. Figure 21(b) shows that a mirror in close vicinity of the receiver significantly degrades the location accuracy, because the sensors are subject to strong reflection that distorts the intensity readings. The top sensor of the receiver now has an average RSS increase of 10% to 17% for different d_r, with a maximum increase about 37%. For $d_r = 35$cm and beyond, the reflection is so weak that the average position accuracy becomes very close to the case without a reflector, that is, slightly below 0.4m.

6.6. Impact of Ambient Light

In this section, we consider how LIPS performs under different conditions of ambient light. We assume the single-lamp, Office scenario, and choose different places in the room at different times of the day, with five grades of illuminance (lux): 1, 150, 312,

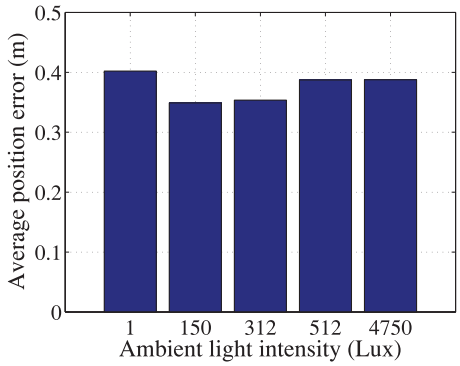

Fig. 22. Location accuracy under different intensities of ambient light.

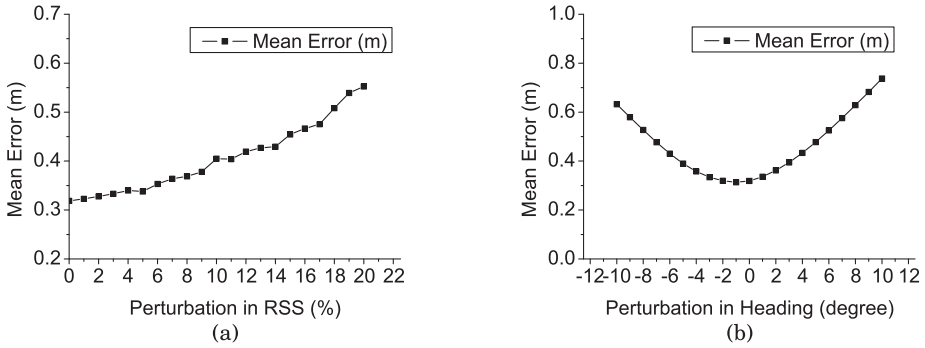

Fig. 23. Position errors changing with perturbations to RSS and heading.

512, and 4,750, where 1 lux means a dark room at night; 150 corresponds to a room at dusk, 312 in the morning, and 512 at noon, all inside the room. The maximum 4,750 luxes represents a location near the window at noon on a clear day. Figure 22 shows that for the different conditions, the average position error is quite stable and remains below 0.4m, with no noticeable impact from variation of ambient lights.

6.7. Impact of RSS/Heading Errors

In this set of experiments, we examine the sensitivity of positioning quality to two sources of error: RSS error and heading error. We introduce artificial perturbation to the measured data and then calculate the position error. For each RSS reading, a multiplicative perturbation $(1 \pm \epsilon)$ is applied, where $0 \leq \epsilon \leq 20\%$ is a parameter and the sign is chosen randomly. For each heading measurement, an additive perturbation $\epsilon_h \in [-10^\circ, 10^\circ]$ is applied. In our experiments, the fluctuation of both measures was rarely found to exceed half of the assumed ranges.

Figure 23 shows how positioning accuracy changes with the perturbations in the Office scenario. As expected, the average errors increase with larger perturbations. The increasing trend, however, is smooth and mild, having average errors well under 1m even for unrealistically large perturbations.

7. DISCUSSION

Power consumption. The LIPS receiver consumes power mainly in three tasks: sensing, computation, and transmission. Each ISL29023 light sensor has a power overhead of $70\mu A \cdot 3.6\ \text{V} = 0.25\text{mW}$ in continuous sensing mode [Intersil 2012]; each magnetic

sensor uses less than 0.2mW at a 20Hz sampling rate, and a LIS33DE accelerometer uses less than 1mW [STMicroelectronics 2009]. Overall, the sensing part accounts for less than 4mW. The MCU and the networking module normally have a much higher power consumption rate. However, their actual power consumption depends on the position sampling rate as specified by the application. The computation and transmission tasks can thus be performed with a low duty cycle to reduce power consumption.

In terms of power overhead by the sensing task itself, LIPS consumes merely 4mW. It takes less than 1s to obtain a relatively stable position from a cold start, which amounts to about 4mJ of energy. By contrast, a WiFi scan used by the WiFi-based positioning system consumes 0.55J [Lin et al. 2010], which is two orders of magnitude higher than LIPS's cost.

Limitations. The current design of the LIPS system does not support other forms of light sources. For example, our RSS model does not apply to the widely used tube-shaped lamps. In principle, the point-source-based model should be extended to a line-segment model. We have done some preliminary experiments with LED tubes and found that the RSS characteristics remain stable, though the relations between RSS and distance/angle become more complicated. This means that a geometry-based positioning method is still possible, though its mathematical foundation needs further investigation.

Currently, the full-scale range of an ISL29023 light sensor is set to 1,000 lux. In a typical indoor environment (e.g., office lighting), the daylight's illuminance is less than 500 lux [WIKI-Lux], which is well within the sensor's capacity. However, the sensor quickly gets saturated under direct sunlight, for example, in the glass-ceilinged atria of some modern buildings, making the receiver unable to position itself. This problem can be mitigated by raising the sensing range (to a maximum of 64,000) when the MCU detects a constant maximum reading. Notice that the increased range will reduce the data resolution, which implies lower sensitivity and reduced location accuracy. For higher adaptability, we will need to seek more powerful light sensors to work under more challenging circumstances.

8. RELATED WORK

In this section, we review related work from three perspectives: positioning principles, cost, and performance. For ease of reference, some of the representative schemes are listed in Table II.

8.1. Positioning Principle

There are two major working principles for indoor positioning systems: signal pattern matching and trilateration/angulation. In the first approach, a location is mapped to a particular pattern of signal characteristics. By measuring the signal pattern and looking up the mapping table, a receiver can obtain its location. RSS fingerprint, for example, is a commonly used signal pattern for WiFi positioning [Youssef and Agrawala 2005; Bahl and Padmanabhan 2000; Yang et al. 2012; Wu et al. 2012b]. In the domain of Radio Frequency Identification (RFID), Tagoram [Yang et al. 2014] localizes RFID tags by constructing an RF phase hologram with observations from several reader antennas. PinIt [Wang and Katabi 2013] leverages the multipath profiles of target tags, as well as of reference tags, to generate position. A similar technique is further used to enable robot navigation and object manipulation in RF-compass [Wang et al. 2013]. A special case of the signal pattern match approach is proximity detection, in which a receiver localizes itself simply with the positions of signal sources detected nearby [Raffaele and Franca 2003; Ni et al. 2004]. This approach can provide only coarse-grained location information.

Table II. Comparison of Different Indoor Positioning Schemes. Note the Accuracy Results Provided in the Literature Often Vary in a Wide Range Depending on Many System and Environmental Factors, So Here We Only Give a Rough Range of Errors That Cover Typical Cases

Schemes/systems	Medium type	Primary principle	Accuracy	Indoor in-frastructure	Limitations
Radar [Bahl and Padmanabhan 2000]	WiFi	RSS pattern match	3–5m	WiFi hotspots	High labor cost
Horus [Youssef and Agrawala 2005]	WiFi	RSS pattern match	<2m	WiFi hotspots	High labor cost
Tagoram [Yang et al. 2014]	RFID	Phase pattern match	<0.2m	RFID readers	Mobile objects only
PinIt [Wang and Katabi 2013]	RFID	Multipath profiles match	<0.2m	Mobile readers and ref. tags	Dense ref. tags
Ubisense [Sahinoglu et al. 2008]	UWB	TDOA+AOA	<0.3m	Dedicated transmitters	Expensive devices
ABS [Tarzia et al. 2011b]	Ambient sound	Spectrogram match	Room level	None	Low accuracy
Geo-Magnetism [Chung et al. 2011]	Magnetic field	Magnitude pattern match	4.7m	None	High labor cost
FM-fingerprint [Chen et al. 2012]	FM broadcast	RSS+PHY info match	Room level	None	Low accuracy
Lighthouse [Romer 2003]	Light	Trilateration	N/A	LED luminaires	Complicated design
VL Landmark [Rajagopal et al. 2014]	Light	Proximity	Meters	LED luminaires	Low accuracy
Luxapose [Kuo et al. 2014]	Light	AOA	<1m	LED luminaires	Dense lamps
Epsilon [Li et al. 2014]	Light	Trilateration	<1m	LED luminaires	Coverage degree ≥ 3
LIPS	Light	Multi-sensor	<1m	LED luminaires	Dedicated receiver

The trilateration/angulation approach can produce accurate position when ranging/angulation ability is available on the signal transmitter or receiver. Various methods, such as time-of-arrival, time difference of arrival (TDOA), angle of arrival (AOA) [Niculescu and Nath 2003, 2004], and two-way sensing [Liu et al. 2012], can be used to obtain distance or angle, based on which a position can be calculated. To improve accuracy, sometimes the physical layer information is exploited for more accurate distance estimation [Wu et al. 2012a; Sen et al. 2013; Mariakakis et al. 2014].

Light signal-based indoor positioning has been considered in many previous studies. Rajagopal et al. [2014] propose to use LED lamps as light landmarks and exploiting the rolling shutter effect of camera sensors. ByteLight [ByteLight] is a commercial project that appears to adopt a similar idea. At the time of this writing, however, there is no detailed information about how the ByteLight system operates. For better accuracy, light intensity (i.e., RSS) is often used with a multilateration method to locate the receiver [Barry 1994; Zhou et al. 2012; Yang et al. 2013a; Yoshino et al. 2008]. In Prince and Little [2012], a hybrid RSS/AoA algorithm is proposed for location optimization. Luxapose [Kuo et al. 2014] uses AoA and image analysis techniques to improve the location accuracy with a smartphone camera, assuming a dense deployment of LED lights.

In Yang et al. [2013b], a positioning system is designed that uses a single transmitter and multiple receivers based on visible light produced by white LEDs. The multiple receivers are independent, and thus the positioning principle behind them differs fundamentally from ours. The Epsilon [Li et al. 2014] system also uses light signal and is

primarily based on the classic trilateration method. A variation of the method is devised to reduce the requirement on number of light sources, but the proposed method requires special manual operations from the user, which makes automatic positioning and tracking impossible.

8.2. Infrastructure Dependence and Deployment Cost

Early indoor positioning systems use dedicated devices such as Bluetooth [Raffaele and Franca 2003], RFID [Ni et al. 2004], or sensor nodes to realize positioning via proximity detection but, on the other hand, require a relatively dense deployment of signal transmitters. Higher positioning accuracy is provided by more accurate ranging techniques such as a combination of radio and sound signals [Priyantha et al. 2000; Tarzia et al. 2011a], Ultra Wide Band (UWB) technology [Sahinoglu et al. 2008], and so on. These techniques require synchronization and coordination among signal transmitters or expensive devices, which greatly increases the deployment cost.

A class of positioning techniques require little or no dedicated infrastructure. They use ambient signals such as cellular radio [Varshavskya et al. 2007], FM radio [Chen et al. 2012], magnetism [Chung et al. 2011], or ultrasound [Borriello et al. 2005] to create position fingerprints. These solutions often provide only coarse-grained positioning, for example, at room-level granularity [Borriello et al. 2005]), or work for only special environments (e.g., steel-rich buildings [Chung et al. 2011]). The WiFi RSS fingerprint-based scheme exploits pre-existing Access Point (AP) hotspots to create signal fingerprints. Depending on the density and placement of APs, the positioning accuracy varies from below 1m to more than 10m [Turner et al. 2011].

Infrastructure cost depends not only on hardware investment but also on human effort in site survey. The latter factor is a serious concern for fingerprinting-based solutions, represented by the WiFi scheme. Recently, researchers have proposed various schemes that use little or zero explicit human effort [Rai et al. 2012; Yang et al. 2012; Wu et al. 2012b]. These techniques leverage inertial sensors on smartphones to automatically infer a user's real position while collecting RSS data. These solutions are not generic, since they require the environments to possess special structural characteristics. When the indoor environment contains large free spaces (e.g., in a factory or a large warehouse), or has a symmetric layout, position ambiguity may arise that prevents exact positions from being inferred from the user's trace.

LIPS uses off-the-shelf LED lamps as signal sources, with a cheap microcontroller attached to each lamp to control flashing. On the receiver side, a dedicated receiver uses a number of light and magnetic sensors, each costing no more than a few US dollars.

8.3. Accuracy and Stability

Different applications require different levels of positioning accuracy. For example, in-building pedestrian route guidance may work well with an accuracy to a few meters, while automated handling requires positioning accuracy within 1cm [Sahinoglu et al. 2008]. The WiFi scheme mostly offers an accuracy of 1m to 3m and is known to be instable in positioning quality; this is because the RSS at a fixed position varies significantly over time and is sensitive to obstacle presence, receiver orientation, and type of device. Also, the distinctiveness of fingerprints depends heavily on AP density [Chandrasekaran et al. 2009]. It is reported that the same scheme can produce drastically different performance across different environments. For example, the classic RADAR scheme is found to generate median accuracies of 1.3m and 5m in small and large buildings, respectively [Chintalapudi et al. 2010]. An empirical study [Turner et al. 2011] shows that an algorithm can yield 5 times worse accuracy than in its original test environment. The same problem is shared by some other fingerprint-based schemes, such as those using sound signal [Tarzia et al. 2011b], magnetic field

magnitude [Chung et al. 2011], and FM broadcast [Chen et al. 2012]. Some RFID systems [Yang et al. 2014; Wang and Katabi 2013] use phase pattern match for positioning and can achieve very high accuracy (e.g., below 0.2m), assuming some special equipment such as mobile objects, mobile antenna, dense reference tags, and so on, which restricts their applications. Light-based positioning systems can generally achieve accuracy below 1m [Kuo et al. 2014; Li et al. 2014] with good stability.

9. CONCLUSION

We have presented a light intensity–based positioning system, *LIPS*, for indoor environments. LIPS exploits ordinary lighting devices such as LED lamps as signal transmitters and uses light sensors as signal receivers. Several light sensors on a receiver can jointly determine the receiver's position with the measured RSS. The main contribution of LIPS is that it explores a new way of indoor positioning, with fairly high accuracy and high stability. The design is fingerprint free, requiring little human intervention other than the establishment of an RSS model. In the future, we will extend the design to include more complex lighting environments.

APPENDIX A: PROOF OF THEOREM 3.1

PROOF. Consider three linearly independent sensing planes, $A_i x + B_i y + C_i z = 0, 1 \leq i \leq 3$, that generate three nonzero RSS values s_1, s_2, and s_3. Substituting these variables into Equation (2), we can get a system of nonlinear equations:

$$\begin{cases} \dfrac{A'_1 x + B'_1 y + C'_1 z}{(x^2+y^2+z^2)^{3/2}} f_\omega \left(\arccos \dfrac{z}{\sqrt{x^2+y^2+z^2}} \right) = \dfrac{s'_1}{k}, & (4) \\ \dfrac{A'_2 x + B'_2 y + C'_2 z}{(x^2+y^2+z^2)^{3/2}} f_\omega \left(\arccos \dfrac{z}{\sqrt{x^2+y^2+z^2}} \right) = \dfrac{s'_2}{k}, & (5) \\ \dfrac{A'_3 x + B'_3 y + C'_3 z}{(x^2+y^2+z^2)^{3/2}} f_\omega \left(\arccos \dfrac{z}{\sqrt{x^2+y^2+z^2}} \right) = \dfrac{s'_3}{k}, & (6) \end{cases}$$

where $A'_i = A_i / \sqrt{A_i^2 + B_i^2 + C_i^2}$, $B'_i = B_i / \sqrt{A_i^2 + B_i^2 + C_i^2}$, $C'_i = C_i / \sqrt{A_i^2 + B_i^2 + C_i^2}$, and $s'_i = s_i \in (0, s_m/k]$ or $s'_i = -s_i \in [-s_m/k, 0)$, depending on the symbol of $A'_i x + B'_i y + C'_i z$. Performing dividing among these equations gives

$$\begin{cases} \left(A'_1 - \dfrac{s_2}{s_1} A'_2 \right) x + \left(B'_1 - \dfrac{s_2}{s_1} B'_2 \right) y + \left(C'_1 - \dfrac{s_2}{s_1} C'_2 \right) z = 0, & (7) \\ \left(A'_1 - \dfrac{s_3}{s_1} A'_3 \right) x + \left(B'_1 - \dfrac{s_3}{s_1} B'_3 \right) y + \left(C'_1 - \dfrac{s_2}{s_1} C'_3 \right) z = 0. & (8) \end{cases}$$

It can be shown that Equation (7) and Equation (8) are linearly independent; otherwise, it can be verified that the vectors (A_i, B_i, C_i) are linearly dependent, which contradicts with the assumption. With this linear independence, we can represent x and y with z as $x = c_1 z$ and $y = c_2 z$, where c_1 and c_2 are functions of s_i, A_i, B_i and C_i. Substituting them into Equation (4) and using the fact $z > 0$ can solve for z, which then gives x and y. Thus we obtain a unique solution of (x, y, z). □

REFERENCES

P. Bahl and V. N. Padmanabhan. 2000. RADAR: An inbuilding RF-based user location and tracking system. In *INFOCOM*.

J. R. Barry. 1994. *Wireless Infrared Communications*. Springer (1994).

G. Borriello, A. Liu, T. Offer, C. Palistrant, and R. Sharp. 2005. WALRUS: Wireless acoustic location with room-level resolution using ultrasound. In *MobiSys*.

ByteLight. http://www.bytelight.com/.

G. Chandrasekaran, M. A. Ergin, J. Yang, S. Liu, Y. Chen, M. Gruteser, and R. P. Martin. 2009. Empirical evaluation of the limits on localization using signal strength. In *6th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*.

Y. Chen, D. Lymberopoulos, J. Liu, and B. Priyantha. 2012. FM-based indoor localization. In *ACM MobiSys*.

K. Chintalapudi, A. P. Iyer, and V. N. Padmanabhan. 2010. Indoor localization without the pain. In *Mobicom*.

J. Chung, M. Donahoe, I. Kim, C. Schmandt, P. Razavi, and M. Wiseman. 2011. Indoor location sensing using geo-magnetism. In *MobiSys*.

S. Gottschalk and J. Hughes. 1993. Autocalibration for virtual environments tracking hardware. In *ACM SIGGRAPH*.

B. Hoff and R. Azuma. 2000. Autocalibration of an electronic compass in an outdoor augmented reality system. In *IEEE/ACM Intl. Symposium on Augmented Reality*.

Intersil. 2012. *The ISL29023 Datasheet*.

Y. Kuo, P. Pannuto, K.-J. Hsiao, and P. Dutta. 2014. Luxapose: Indoor positioning with mobile phones and visible light. In *ACM Mobicom*.

L. Li, P. Hu, C. Peng, G. Shen, and F. Zhao. 2014. Epsilon: A visible light based positioning system. In *Proc. NSDI*.

K. Lin, A. Kansal, D. Lymberopoulos, and F. Zhao. 2010. Energy-accuracy trade-off for continuous mobile device location. In *MobiSys*.

H. Liu, Y. Gan, J. Yang, S. Sidhom, Y. Wang, Y. Chen, and F. Ye. 2012. Push the limit of WiFi based localization for smartphones. In *ACM Mobicom*.

A. T. Mariakakis, S. Sen, J. Lee, and K.-H. Kim. 2014. SAIL: Single access point-based indoor localization. In *Proc. of MobiSys*.

L. Ni, Y. Liu, C. Yiu, and A. Patil. 2004. LANDMARC: Indoor location sensing using active RFID. In *WINET*.

D. Niculescu and B. Nath. 2003. Ad hoc positioning system (APS) using AOA. In *IEEE INFOCOM*.

D. Niculescu and B. Nath. 2004. Vor base stations for indoor 802.11 positioning. In *ACM Mobicom*.

G. B. Prince and T. D. Little. 2012. A two phase hybrid RSS/AoA algorithm for indoor device localization using visible light. In *IEEE Global Communications Conference (GLOBECOM)*.

N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. 2000. The cricket location-support system. In *MobiCom*.

B. Raffaele and D. Franca. 2003. Design and analysis of a bluetooth-based indoor localization system. In *Workshop on Personal Wireless Communications*.

A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen. 2012. Zee: Zero-effort croundsourcing for indoor localization. In *Mobicom*.

N. Rajagopal, P. Lazik, and A. Wowe. 2014. Visual light landmarks for mobile devices. In *IPSN*.

K. Romer. 2003. The lighthouse location system for smart dust. In *ACM MobiSys*.

Z. Sahinoglu, S. Gezici, and I. Guvenc. 2008. *Ultra-Wideband Positioning Systems: Theoretical Limits, Ranging Algorithms, and Protocols*.

S. Sen, J. Lee, K.-H. Kim, and P. Congdon. 2013. Avoiding multipath to revive inbuilding WiFi localization. In *ACM Mobisys*.

STMicroelectronics. 2009. *LIS33DE Datasheet*.

STMicroelectronics. 2010. *Using LSM303DLH for a Tilt Compensated Electronic Compass, AN3192 Application Note*.

S. Tarzia, P. Dinda, R. Dick, and G. Memik. 2011a. Indoor localization without infrastructure using the acoustic background spectrum. In *ACM MobiSys*.

Stephen P. Tarzia, Peter A. Dinda, Robert P. Dick, and Gokhan Memik. 2011b. Indoor localization without infrastructure using the acoustic background spectrum. In *ACM Mobisys*.

D. Turner, S. Savage, and A. Snoeren. 2011. On the empirical performance of self-calibrating wifi location systems. In *Local Computer Networks (LCN)*.

A. Varshavskya, E. D. Lara, J. Hightower, A. LaMarca, and V. Otsason. 2007. GSM indoor localization. In *Pervasive and Mobile Computing*.

J. Wang, F. Adib, R. Knepper, D. Katabi, and D. Rus. 2013. RF-compass: Robot object manipulation using RFIDs. In *ACM Mobicom*.

J. Wang and D. Katabi. 2013. Dude, where's my card? Rfid positioning that works with multipath and non-line of sight. In *Proc. of ACM SIGCOMM*.

WIKI-AGV. 2016. *Automated Guided Vehicle*. https://en.wikipedia.org/wiki/Automated_guided_vehicle.

WIKI-DifRef. 2016. *Diffuse Reflection*. http://en.wikipedia.org/wiki/Diffuse_reflection.

WIKI-Lux. *Lux*. http://en.wikipedia.org/wiki/Lux#cite_note-9.

WIKI-RotMatrix. 2016. *Rotation Matrix*. http://en.wikipedia.org/wiki/Rotation_matrix.

C. Wu, Z. Yang, Y. Liu, and W. Xi. 2012b. WILL: Wireless indoor localization without site survey. In *INFOCOM*.

K. Wu, J. Xiao, Y. Yi, M. Gao, and L. M. Ni. 2012a. FILA: Fine-grained indoor localization. In *IEEE INFOCOM*.

L. Yang, Y. Chen, X.-Y. Li, C. Xiao, M. Li, and Y. Liu. 2014. Tagoram: Real-time tracking of mobile RFID tags to high precision using COTS devices. In *Proc. of Mobicom*.

S.-H. Yang, E.-M. Jeong, D.-R. Kim, H.-S. Kim, Y.-H. Son, and S.-K. Han. 2013a. Indoor three-dimensional location estimation based on LED visible light communication. *Electron. Lett.* 49, 1 (Jan. 2013).

S.-H. Yang, E.-M. Jung, and S.-K. Han. 2013b. Indoor location estimation based on LED visible light communication using multiple optical receivers. *IEEE Commun. Lett.* 17, 9 (Sept. 2013).

Z. Yang, C. Wu, and Y. Liu. 2012. Locating in fingerprint space: Wireless indoor localization with little human intervention. In *ACM Mobicom*.

M. Yoshino, S. Haruyama, and M. Nakagawa. 2008. High-accuracy positioning system using visible LED lights and image sensor. In *IEEE Radio and Wireless Symposium*.

S. You, U. Neumann, and R. Azuma. 1999. Hybrid inertial and vision tracking for augmented reality registration. In *IEEE Virtual Reality*.

M. Youssef and A. Agrawala. 2005. The horus WLAN location determination system. In *MobiSys*.

Z. Zhou, M. Kavehrad, and P. Deng. 2012. Indoor positioning algorithm using light-emitting diode visible light communications. *Opt. Eng.* 51, 8 (2012).

Received January 2015; revised December 2015; accepted June 2016