SpinLight: A High Accuracy and Robust Light Positioning System for Indoor Applications

Bo Xie
University of Chinese
Academy of Sciences
SIAT, Chinese Academy of
Sciences, China
bo.xie@siat.ac.cn

Guang Tan
SIAT, Chinese Academy of
Sciences
Shenzhen, China
guang.tan@siat.ac.cn

Tian He
University of Minnesota-Twin
Cities
Minneapolis, MN 55455, U.S.
tianhe@cs.umn.edu

ABSTRACT

This paper presents *SpinLight*, an indoor positioning system that uses infrared LED lamps as signal transmitters, and light sensors as receivers. The main idea is to divide the space into *spatial beams* originating from the light source, and identify each beam with a unique timed sequence of light signals. This sequence is created by a coded shade that covers and rotates around the LED, blocking the light or allowing it to pass through according to pre-defined patterns. The receiver, equipped with a light sensor, is able to determine its spatial beam by detecting the light signals, followed by optimization schemes to refine its location within that beam. We present both 2D and 3D localization designs, demonstrated by a prototype implementation. Experiments show that SpinLight produces a median location error of 3.8 cm, with a 95th percentile of 6.8 cm. The receiver design is very low power and thus can operate for months to years from a button coin battery.

Categories and Subject Descriptors

C.5.3 [**COMPUTER SYSTEM IMPLEMENTATION**]: Microcomputers—*Portable devices (e.g., laptops, personal digital assistants)*

General Terms

Design, Experimentation

Keywords

Indoor positioning; Light; Sensing

1. INTRODUCTION

Fine-grained indoor localization can enable a multitude of applications. In the domain of manufacturing and warehouse management, robots with accurate location information can help to place, fetch or assemble objects. The same

concept can be used for healthcare applications such as aid for the visually impaired, or as part of a home service for the elderly. In supermarkets and shopping malls, exact location of products can greatly improve the customer's shopping experience, promote targeted advertisements, and enable customer analytics [12]. In a recent survey [27], the market for location-based marketing and advertising is predicted to surpass 10 billion by 2018. To date, both the industry and academic community have made tremendous efforts to realize inexpensive, reliable, and easy-to-use location services, exploring a plethora of techniques, yet the adoption of high accuracy indoor localization is still in its infancy.

Recently, there has been an increase of interest in indoor localization using visual light communication (VLC) technologies [3, 12–14, 24, 34, 44], due to the pervasiveness of lighting infrastructure. Compared with radio frequency (RF) technologies, the propagation of light is more predictable, thus allowing more reliable location estimation using geometric models. For example, light intensity or RSS-based VLC localization has been explored in many previous studies [14, 34, 44]. Recent work shows that this approach can achieve accuracy to a half meter. The Angle of Arrival (AoA) approach [18] has also been considered for localization using smartphone cameras [3, 13].

In this paper we present a new approach to light-based localization, targeting object tracking applications and smartphone users. Apart from high accuracy, we want to achieve several goals that are important to the considered applications, and yet not completely fulfilled by prior work: (1) robustness against noise and ambient interference, (2) insensitivity to receiver orientation (provided line of sight), and (3) low energy consumption of location sensing, which is important to small-sized location tags.

In our system, called *SpinLight*, the 3D space is divided into *spatial beams* originating from the transmitter, whose coordinates are known in advance. Each beam is identified with a unique timed sequence of infrared (IR) signals. Here, we adopt infrared (IR) instead of visible light for avoiding visual disturbance to human eyes. The signal sequence is created by a hemispherical shade that covers and rotates around the LED under the drive of a programmed step motor. The shade is divided into cells by its longitude and altitude lines, and some cells are removed to allow the light to pass through. The rotation of the shade generates unique light patterns in different spatial beams. The receiver, equipped with a light sensor, is able to determine the spatial beam containing it by detecting and decoding the

signal pattern. With a single transmitter, the receiver can perform 2D localization, assuming a known height. An additional transmitter can be used to realize full 3D localization by creating two spatial beams whose intersection gives the receiver's location.

SpinLight features a digital-only approach in which sensors make binary decision as to the existence of a light signal. This is in contrast to existing light positioning approaches that rely on analog signals, i.e., RSS measurements. Given its digital nature, our system is significantly less affected by noise, interference, and multipath effect than RSS based systems, enjoying an order of magnitude improvement in location accuracy. Our digital approach also makes SpinLight insensitive to receiver orientation, which is shown to be a critical factor for RSS or AoA measurements and complicates receiver design [13, 14].

Very recently, there have been a few localization systems that achieve centimeter-level accuracy using RFID technology [35, 36, 43]. While these systems have the advantage of being able to work without line of sight, they fall short in several aspects, such as requiring highly dense deployment of reference tags [36], or only working for mobile objects [43]. These techniques and our design therefore have complementary functions and characteristics.

In summary, this paper makes two major contributions:

1. We propose a light positioning approach based on the idea of spatial division. Fundamentally differing from traditional methods such as RSS or AoA, this approach allows more robust localization at the receiver, with reduced requirements on receiver design and user's cooperation, and at a very low power consumption rate.

2. We present a series of optimization schemes to refine the location results of the basic design, and deal with the problem of gray areas, which are a major source of errors. Experiments show that with a single transmitter, the median location error is 3.6 cm. An additional transmitter allows 3D localization, with median accuracy of 3.8 cm. In both cases, the 95th percentiles are within 9.2 cm. The results are robust against environmental dynamics and interference.

The next section provides an overview of our system; Section 3 describes a coarse-grained localization approach; Section 4 presents optimization schemes to refine the location and deal with errors; Section 5 describes how to perform 3D localization; Section 6 presents the experimental results; Section 7 discusses several practical issues of the system; Section 8 introduces related work, and Section 9 concludes the paper.

2. SYSTEM OVERVIEW

In this section, we outline the basic design of SpinLight. The transmitter's 3D coordinates are known in advance. For simplicity, we assume that the receiver is placed on the floor plane, that is, its height $z = 0$. This assumption will be removed when we discuss 3D localization.

The transmitter device contains an IR LED chip embedded in the center of a circular base plane, which is mounted on the ceiling. Current LED chips at a few watts are quite small in size (e.g., 0.5 cm×0.5 cm), so can be approximately considered a point light source in the field. The LED chip is controlled by a microcontroller (MCU) to flash at specified
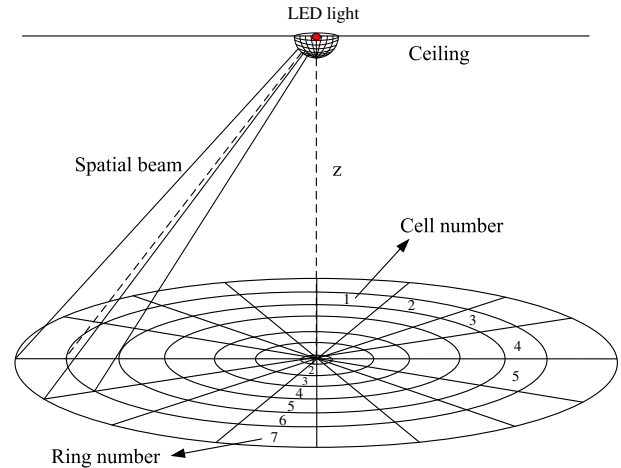
Figure 1: Spinlight overview. The system primarily comprises an IR LED lamp covered by a hemisphere shade. The shade is divided into cells that define the spatial beams. The task of the receiver is to find the beam it belongs to.
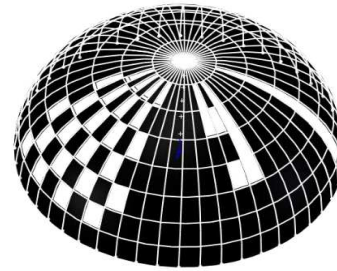
Figure 2: A shade divided into 8 s-rings, each further divided into 48 s-cells. A hollow cell represents a 1, and a closed cell 0.

rates, in order to distinguish itself from other transmitters and ambient lights in the frequency domain.

There is a hemispherical shade that covers the LED chip from below, and can rotate horizontally around the chip at a constant speed under the drive of a step motor. The shade is evenly divided to *shade rings* (or s-rings for short) and *shade cells* (s-cells) following longitude and latitude lines with pre-defined intervals, see Figure 1 for an illustration. If the shade stays still, then from the viewpoint of the light source, each s-cell is projected onto the floor, forming an area which we call a *cell*. A shade ring's projection on the floor is called a *ring*. Along with the transmitter's coordinates, the shade's radius and division parameters are also known to the receiver. The information of system configuration can be provided by a location server. Thus, assuming a known height, the receiver can calculate the layout of the projected rings and cells of a transmitter on its plane.

On the shade, some s-cells are removed to create holes, or hollow s-cells, allowing the light to pass through and form solid beams of light. These light beams generate light spots on the floor. If we label an open s-cell '1', and a closed one '0', then the arrangement of the cells creates a *coding scheme* of the shade, see Figure 2 for an example. When the shade spins around, the light beams rotate in the air, and the light spots on the floor move around in a way like the spotlight shadows on a dance floor. When a light spot sweeps across
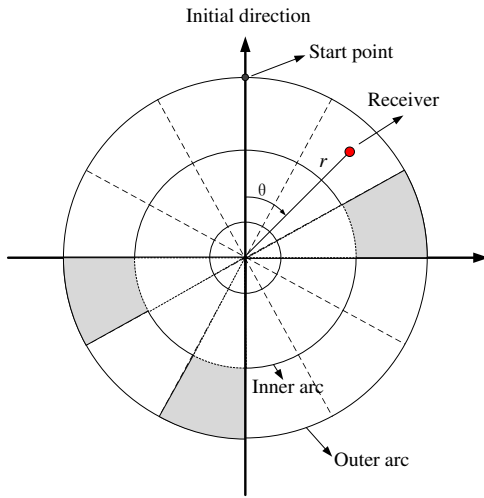
Figure 3: Projection of the coded shade.

the receiver, the light sensor will detect a window of light signal, called a *signal window*. For each s-ring, we create holes in different ways, so that receivers located in different rings can detect different flash patterns, so as to find their own rings.

Due to symmetry, receivers within different cells of the same ring will see the same sequence of signal. To differentiate the cells, we introduce a synchronization mechanism. We mark a *start point* on the upper edge of the shade, and use its angular offset clockwise from some initial direction (e.g., geographical north) to define the shade's *current direction*. Also, the LED light uses two different flash rates, and switches between the two when the shade's start point hits the initial direction. By detecting the switch of the flashing rate, the receiver knows that a new round of shade rotation begins. Thereafter, the receiver counts the number of signal windows passed to estimate its angular offset to the initial direction, called its *angle*, and determine its cell number. Section 3 describes how the receiver determines its cell number.

Identifying the receiver's cell only gives a rough location result. We propose a method to fine-tune the receiver's location. The idea is to use a specially shaped cell in each s-ring. This cell is shaped in such a way that the receiver will detect a signal window whose length is a function of the receiver's offsets to the cell's edges. With this approach, the receiver can estimate its relative position within its cell. Section 4 covers the details.

When z is unknown, position ambiguity will arise – at each height the receiver can find a spot that satisfies the timed flash pattern. In fact, there is a ray from the light source that contains infinitely many candidate positions for the receiver. In this case, we use two transmitters to create two candidate rays, whose intersection is the receiver's real position. Section 5 describes how the position is obtained.

3. 2D LOCALIZATION

In this section we first introduce the rotation control of the shade, whose main purpose is to let the receiver know when a new round of shade rotation begins, in order to calculate its angle with respect to the initial direction. We then describe
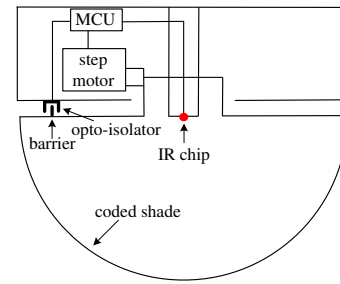

Figure 4: Structure of transmitter.

the coding design and how a receiver processes the received signal to determine its cell and final location.

We focus on a 2D plane, onto which the shade is projected. Figure 3 shows an example of the projected area in the form of a polar grid, where the shade's current direction is zero. In this polar coordinate system, the receiver's location is represented by radius r and angle θ. With known 3D coordinates of the transmitter, (r, θ) can be translated to a 2D location. There are M rings, each consisting of N cells. There is a one-to-one correspondence between a cell on the floor and an s-cell on the shade. A ring consists of two circles, an *inner circle* and an *outer circle*. A cell has two arcs, one on its inner circle, called the *inner arc*, and the other on its outer circle, called the *outer arc*.

3.1 Rotation and flash control

The shade rotates horizontally around the LED chip at a constant speed in revolutions per minute (rpm). Thus, it takes the shade a constant time, called a *cell period*, to rotate the width of an s-cell. Apart from controlling the step motor, the MCU also switches on and off the light at two pre-defined frequencies. Figure 4 gives an illustration of the lamp structure. To enable the receiver to obtain its angle, an opto-isolator [21] is used. The opto-isolator is a U-shaped component, in which the IR signal is transmitted from one side to the other. We also install a small piece of plastic as a barrier on the start point of the shade. When the barrier passes through the opto-isolator, the IR light is blocked and a state change is triggered and transferred to the MCU, which then switches to another flash rate. Upon detecting a switch of flash rate, the receiver starts to count the number of cells passed, each corresponding to a signal window, to estimate its angle.

3.2 Shade coding and cell recognition

A hollow s-cell on the shade allows the light to pass through, so is labeled '1'; a closed s-cell blocks the light, so is labeled '0'. The code of each s-ring primarily comprises three parts: (1) leading bits, which help the receiver to identify the start point of the shade; (2) ring ID bits, used to identify rings; and (3) extension bits. There are also some spacing bits between the different parts.

The leading bits are simply a sequence of 1 cells, which generate a sequence of signal windows that are easy to identify. The design of the ring IDs is more complicated. Apart from being unique, it needs to tolerate errors. This requires it to encode more information. We adopt a layered structure, in which a ring ID is represented by several *macro bits*, each further represented by a number of regular bits. The
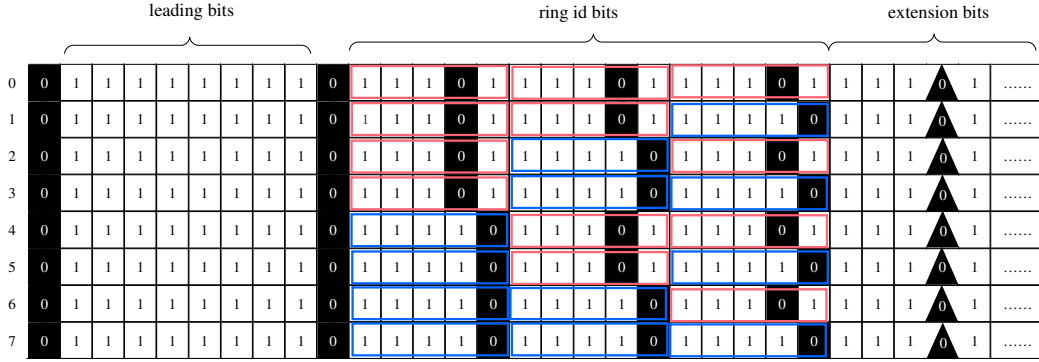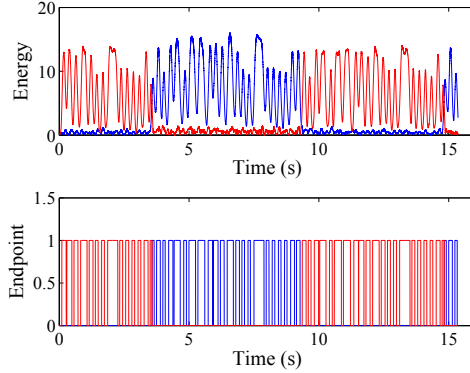
Figure 5: An example of shade code. A blue rectangle represents a macro bit **1**, and a red rectangle a macro bit **0**.
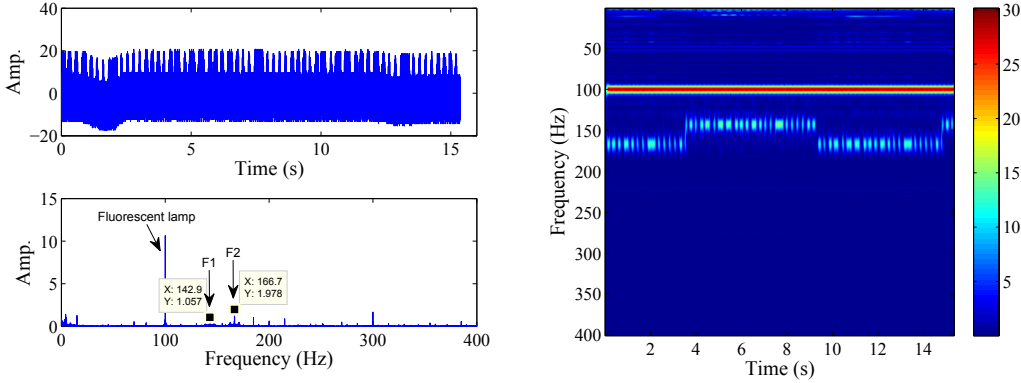
(a) Signals in time and frequency

(b) Spectrogram of signals

(c) Endpoint Detection

Figure 6: Signal processing for cell recognition.

extension bits are for optimizing location accuracy, and may be used for other purposes such as advanced error detection, which we leave for future study. The leading bits must be longer than a macro bit, and in the same ring, no two 0 bits are allowed to appear together. Figure 5 gives an example of shade coding, in which cells are displayed in squares for clarity. In this example, there are eight s-rings, numbered from **000** to **111**. A macro bit **0** is represented by regular bits 11101, and a macro bit **1** by 11110.

On the receiver, the received signal is a mixture of light signals from multiple sources, including possibly more than

one transmitters, daylight, as well as ordinary lamps. For example, Figure 6(a) shows a sequence of signals obtained by a receiver and its frequency components. Note that the average amplitude is subtracted from the raw signal, and the fluctuation of the signal is due to interference from a person walking by while collecting the data. The transmitter uses two flash rates: 143 Hz and 167 Hz, mixed with IR from fluorescent lamps at 100 Hz. We use the wavelet transform as our signal processing tool. Figure 6(b) shows the spectrogram of the signal that contains three frequency components. A bright spot corresponds to a time interval during
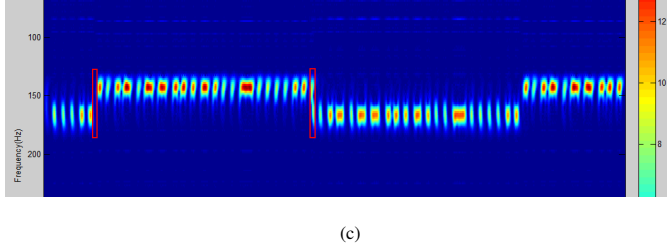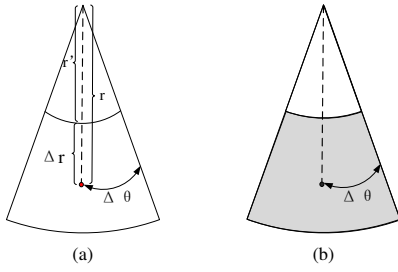
Figure 7: The receiver's location offsets, Δr and $\Delta \theta$, to its cell edges. (a) Receiver is located in a 1 cell. (b) Receiver is located in a 0 cell. (c) Wavelet spectrogram of signals used to estimate $\Delta \theta$ for a receiver initially located in a 0 cell.

which a light beam scans across the light sensor, while a gap between bright spots means a time interval during which the light is blocked. The brighter the spots, the stronger the signal. The top level is from the surrounding fluorescent lamps, which together emit IR about seven times stronger than the light from the transmitter. From the receiver's viewpoint, the sequence of bright and dark spots reflects the ring code. For example, the longest segment of consecutive bright spots represent the leading bits.

From the results of wavelet transform, we can further get the energy function of the signal, as shown in the upper panel of Figure 6(c). By filtering out irrelevant frequency components (in this case the 100 Hz one), using an endpoint detection method like those adopted in the field of speech recognition, and normalizing the energy, we can obtain the ring code shown in the lower panel of Figure 6(c). In accordance with Figure 6(b), the width of a bar is proportional to the number of 1 bits. The ring ID part directly gives the receiver's ring number, while the time interval between the first flash rate switch point and the beginning of the leading bits implies the receiver's cell number. The center of the determined cell (i.e., intersection of the cell's diagonals) is taken as the receiver's location.

4. 2D LOCALIZATION: OPTIMIZATIONS

The receiver's ring number and cell number only provide a rough estimate of its real location. To improve accuracy, one could manage to create highly dense s-cells on the shade. Suppose we want to limit the size of a (possibly distant) cell on the floor to a few centimeters, an s-cell may need a sub-millimeter level size. This will generate very blurry cell edges, because the holes are smaller than the LED chip. As a result, the receiver faces more uncertainty in detecting cells, and likely generate more errors. Moreover, a delicate construction of the shade complicates the design of the transmitter, which increases the cost. In this section we describe

how to refine the location result in a software approach, and how to deal with some reliability issues arising in this process.

4.1 Fine-grained localization

To locate the receiver at a finer grain, we need to obtain its relative location to the edges of its cell. As illustrated in Figure 7(a), define $\Delta r = |r - r'|$ as the receiver's radial offset, where r is the receiver's radius, and r' the radius of its ring's inner circle. Also, define the angular offset, $\Delta \theta = |\theta - \theta'|$ as the angular distance to the side of cell that is closer to the initial direction clockwise. We label a cell '1' if its corresponding s-cell is 1, and '0' otherwise.

4.1.1 Determining $\Delta \theta$

There are two cases to deal with when determining $\Delta \theta$. First, if the receiver's location is in a 1 cell (Figure 7(a)), then the receiver is likely to see a non-integral multiple of cell periods since the shade starts a new round of rotation. Figure 6(c) gives an example in which the first bar of the blue wave is only a fraction of a regular signal window. This fraction multiplied by a cell's width in degree, that is, $360/N$, is then equal to $\Delta \theta$.

If, however, the receiver is located in a 0 cell (Figure 7(b)), then $\Delta \theta$ cannot be obtained using the above method, because the receiver will miss the light signal for an interval of $\Delta \theta$ at the beginning of a new round of rotation, and cannot know the duration of this blackout (note we do not assume clock synchronization between the transmitter and receiver). We propose a *shifted round synchronization* strategy to solve the problem. Recall that when the shade's current direction becomes zero, the MCU will instruct the LED to switch to another flash rate. In our strategy, the rate switch does not happen immediately. Instead, before switching to a certain flash rate, for example the higher one, the MCU stays one more cell period at the lower rate. Since there exist no two 0 bits together in a ring, the one-cell shift will ensure that a receiver initially located in a 0 cell will be exposed to light at the beginning of the next round of rotation, as if it were moved to the next cell clockwise in the ring. Following the method in Figure 7(a), the receiver can obtain $\Delta \theta$. Also, by detecting the pre-determined flash rate (in this case the higher one), the receiver knows it is in a shifted round of rotation, and will compensate for a cell period when calculating θ.

Figure 7(c) provides a concrete example of estimating $\Delta \theta$ for a receiver initially located in a 0 cell. The spectrogram of signal shows the two frequencies of the light, the higher frequency f_h represented by the lower level of light spots, and the lower frequency f_l by the upper level of light spots. As the shade's current direction becomes 0, an $f_h \rightarrow f_l$ switch takes place immediately. At the same time the receiver is covered by a dark (i.e., 0) cell, so will experience a blackout for a partial cell period, see the first blank rectangle in red line in Figure 7(c). From this blackout period, the receiver cannot know its relative position in the cell. Now look at the second rectangle, which represents the time of frequency change $f_l \rightarrow f_h$. With our one-cell shift strategy, the frequency change will be delayed by a cell period, allowing the receiver to be exposed to light for a partial cell period. This can be seen from the upper part of the second rectangle. The remaining cell period is left to f_h. From the split of the cell period, the receiver can obtain $\Delta \theta$.

(a)

(b)
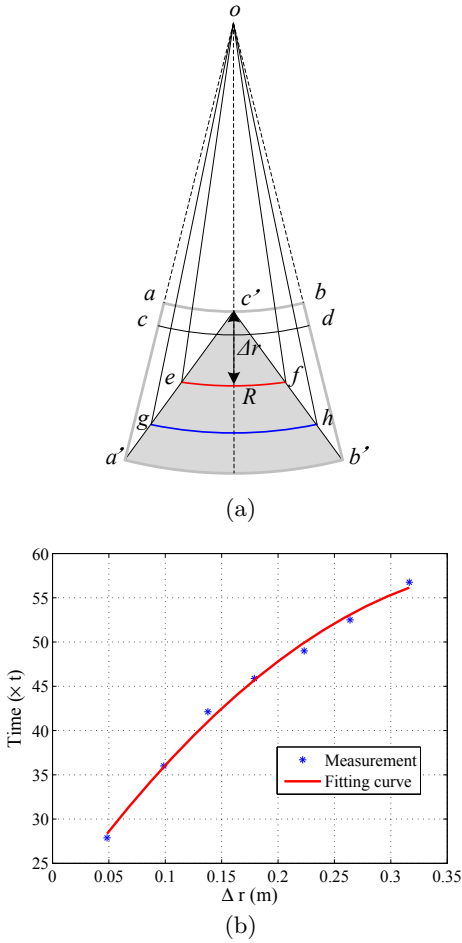
Figure 8: Illustration of the method to obtain Δr. (a) A triangle-like area in the cell $abb'a'$ that generates variable length of scan-arcs (e.g., red and blue arcs) for different Δr. (b) Measured scan-arc length in time for various Δr in a special cell, where t is the data sampling interval.

4.1.2 Determining Δr

For Δr, we propose a solution based on the idea of *variable window length*. Because a cell has a constant angular width, a receiver of different radius within the same cell will receive signal windows of equal length. It is this uniformity that causes the ambiguity of a receiver's radial coordinate. Figure 8 (a) shows a 1 cell, $abb'a'$, in which the two arcs ab and cd, generate signal windows of the same length, and thus are indistinguishable to a receiver on these two arcs. To solve this problem, we choose a special 1 cell on the shade, select a subarea in the shape of a spherical triangle, and close it. The spherical triangle takes the outer arc of the cell as its base side, and the middle point of the inner arc as one of its vertices. This subarea casts a shadow on the floor, see $a'b'c'$ in Figure 8 (a). Given a point with radial offset Δr in the cell, define the *scan-arc* as the intersection of this shadow area and the (0,0) centered circle passing through this point, see for example the red arc in Figure 8. The sides of $a'b'c'$ have the property that the scan-arc has monotonically increasing central angle with Δr. In Figure 8 (a), this means that $\angle eof < \angle goh$ for increasing Δr. This way, we establish a monotonic function that maps a Δr to
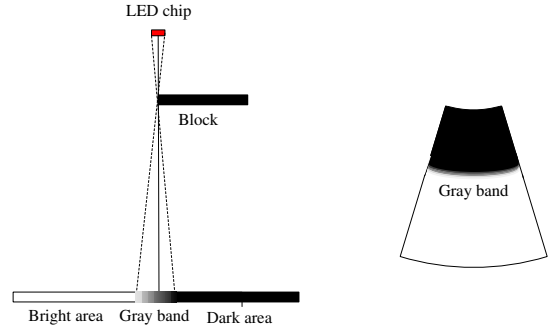
Figure 9: IR beads create gray bands between dark and bright areas.
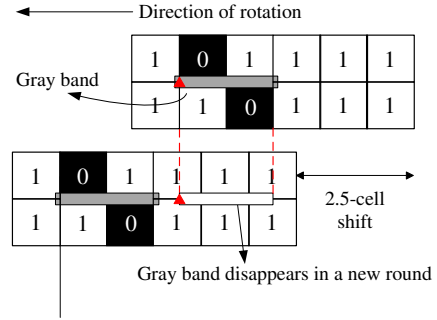
**Figure 10:** Avoiding gray areas by shifting 2.5 cell periods.

the angular distance of a scan-arc. When the receiver detects a signal window from this special cell, it measures the length of the scan-arc, and obtains Δr with this function.

In practice, the narrow tip of a triangular area may cause error of Δr due to the gray area problem (to be discussed in the next section). Thus, we choose to use a trapezoid-like shape, with the only difference that the tip is now an upper base. Figure 8 (b) shows the lengths of scan-arcs for various Δr measured for a special cell. We can see that the relationship between Δr and scan-arc length is highly predictable.

4.2 Dealing with gray areas

So far we have assumed that the LED offers a point source of light. In practice, however, this is not the case. The illuminating LED beads are actually embedded in a small plane with certain spacing, which will create gray bands between cells. Figure 9 shows how gray bands are formed by a non-point of light source. The gray band is an area with continuously decreasing gray degree from full darkness to full brightness. Thus, depending on the threshold value chosen for distinguishing 0 and 1, a cell's effective area may extend to a neighboring cell. In the example of Figure 9, a narrow area of the 0 cell above the geometrical border of the two cells may be seen as part of the 1 cell below, or vice versa. In these cases, we say that a cell *encroaches on* the other cell.

Gray bands can occur between cells within the same ring, or across two neighboring rings, with width much smaller than a cell's side length. The former can be coped with the shifted round synchronization method described earlier, here we consider the *cross-ring gray bands*. Such areas only

216
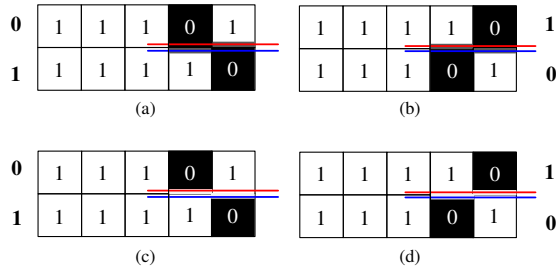
Figure 11: Erroneous cell readings due to gray bands. Red and blue lines are two examples of near-border locations where errors may occur. (a,b) Dark cells encroach on bright ones, causing erroneous reading 00. (c,d) Bright cells encroach on dark ones, causing erroneous reading 11.

Figure 12: 3D localization with two transmitters.

occur on the border between a 1 cell and a 0 cell from two neighboring rings. In general, a receiver located in such an area may mistakenly determine the value of its cell. This leads to two potential problems for the receiver's localization: inaccurate estimation of $\Delta\theta$ and wrong determination of ring number (hence r).

Figure 10 shows a possible result of the first problem. In the figure, the receiver (red triangle) is located in the gray area between the two rings. Suppose the light signal is too weak for it to detect, then it will miss the light signal for at least two cell periods before two vertically (along radial direction) neighboring 1 cells come across. Thus the receiver cannot properly estimate $\Delta\theta$, even by shifting one cell period. Our solution for this problem is to have the shade shift two and a half cell periods instead of one, where the extra half cell period is to avoid the small gray areas around a cell's vertices. By the coding rule, no two 0 cells appear together within a ring, and any 10 and 01 cell combination must be surrounded by three consecutive 1 cells, so the 2.5 cell shift must allow an initially shadowed receiver to be exposed to the light in a new round. In our final design, the 2.5 cell shift strategy is implemented with the lower flash rate. Upon detecting the lower flash rate for a half period cell, the receiver can obtain $\Delta\theta'$, and then $\Delta\theta$ is obtained by adding the degree of 2.5 cell periods, that is $2.5 \times 360/N$, to $\Delta\theta'$.

For the second problem, we show that the potential error can be corrected with the coding rules and the triangle cells. Again, the errors can only happen between two vertically neighboring and different cells. This further means that an error can only occur between two vertically neighboring and different macro bits. Figure 11 shows the possible situations of two neighboring rings where such a situation may happen.

Suppose the 0-1 determination threshold value is set such that a 0 cell encroaches on a vertically neighboring 1 cell. Consider the example in Figure 11(a). If the receiver is located on the upper red line, then it will determine a 00 for the two cells crossed by the red line. The same goes for the lower blue line in Figure 11(a). In other words, if a 0 cell encroaches on a 1 cell, then 00 will always be determined for the last two (regular) bits of each corresponding macro bit. By the coding rules, bits 00 are illegal, so the receiver will recognize such an error. At the same time, the receiver will detect a relative wide dark band from the special cell, and thus realize it is near the border of two neighboring rings. Now look at the ring ID parts coded in
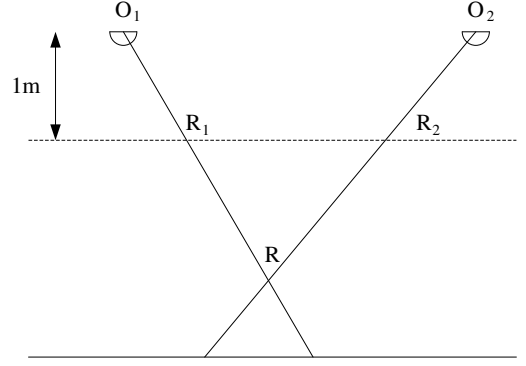
macro bits, and examine the possible situations of two rings, k and $k+1$, whose ID parts may induce such an error. Let $k = \mathbf{b}_1\mathbf{b}_2 \ldots \mathbf{b}_s$, where macro bit $\mathbf{b}_i \in \{\mathbf{0}, \mathbf{1}\}$, then there are three cases where an increase from number k to $k+1$ will cause changes of macro bits:

1. $\mathbf{b}_1\mathbf{b}_2 \ldots \mathbf{b}_{s-1}\mathbf{0}$ to $\mathbf{b}_1\mathbf{b}_2 \ldots \mathbf{b}_{s-1}\mathbf{1}$, for example **0000** to **0001**,

2. $\mathbf{b}_1\mathbf{b}_2 \ldots \mathbf{b}_{s-2}\mathbf{01}$ to $\mathbf{b}_1\mathbf{b}_2 \ldots \mathbf{b}_{s-2}\mathbf{10}$, for example **0001** to **0010**,

3. $\mathbf{b}_1\mathbf{b}_2 \ldots \mathbf{b}_{s-t-1}\mathbf{01}_1 \ldots \mathbf{1}_t$ to $\mathbf{b}_1\mathbf{b}_2 \ldots \mathbf{b}_{s-t-1}\mathbf{10}_1 \ldots \mathbf{0}_t$, for example **0011** to **0100**

In all three cases, the two ring IDs will differ in their last few macro bits, which are aligned vertically in the code. Any pair of vertically adjacent and different macro bits will generate erroneous readings 00. Thus, the two ring IDs may generate one or more such erroneous readings. For example, if the receiver detects 00 for two consecutive macro bits, then it learns that its adjacent two rings must have **01** and **10** as their ending macro bits. More generally, from the number of consecutive macro bits that go wrong, the receiver can determine the last few macro bits of its surrounding two rings. Combining this with the other correctly recognized macro bits, the receiver is able to recover k and $k+1$.

If instead the 0-1 determination threshold value is set such that a 1 cell encroaches on a vertically neighboring 0 cell, as illustrated in Figures 11(c) and (d), then the above argument still holds except that the erroneous reading 00 becomes 11, and the receiver can find k in a similar way.

With k determined, the receiver simply takes the radius of the border circle between ring k and ring $k+1$ as its r. Combining $\Delta\theta$ obtained earlier and r it can obtain its location.

5. 3D LOCALIZATION

So far we have assumed a known height of the receiver, and obtained a 2D location. This is useful for some applications such as robot navigation [6], or object tracking on fixed height carriers. Other applications may require a full 3D location. Notice that within a 3D spatial beam generated by the cells on the lamp shade, the receiver will detect exactly the same flash pattern. Given a 2D location spot p with an assumed height, we draw a ray from the LED to
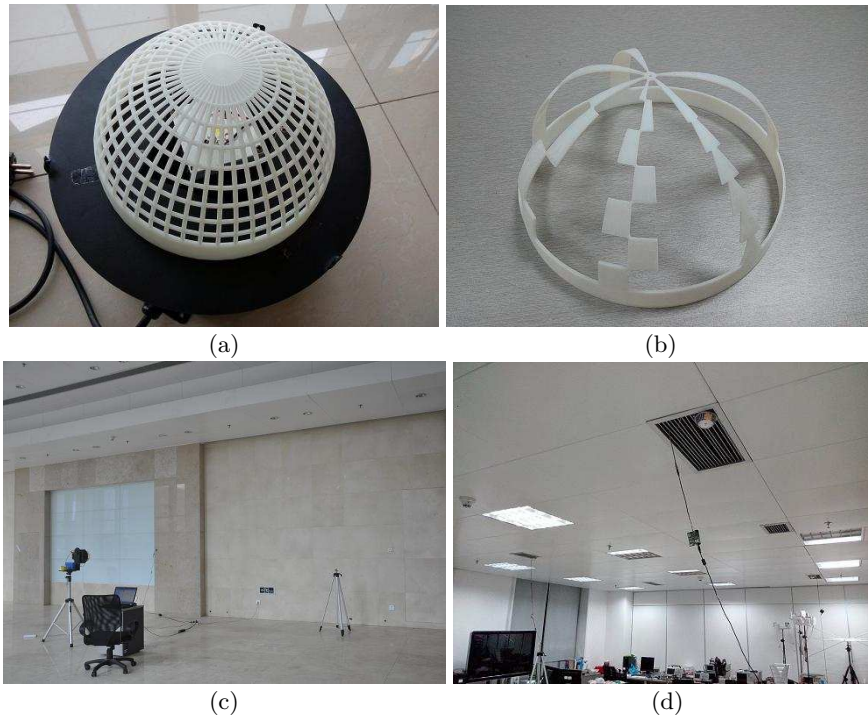
Figure 13: Experimental environments. (a) A SpinLight transmitter with a shade in its general form, (b) An example of the final shade with fixed coding, (c) Test environment for SpinLight, (d) Test environment for Epsilon, Multi-Sensor, and Weighted-Avg, where three LED lamps are mounted on the ceiling.

p, then the ray contains infinitely many points that satisfy the flash pattern constraints. Therefore, without a known height, we need at least an additional transmitter that generates another ray to solve for a 3D location. In our design, for every transmitter detected, the receiver first assumes it has a fixed height (e.g., 1 m) from the ceiling, and obtain a location, generate a ray. For multiple transmitters, the receiver finds a point that has minimum accumulative distance to all the rays using a least square method. Figure 12 provides an example in which two transmitters help the receiver at R determine its location by first finding two candidate points R_1 and R_2, creating two rays O_1R_1 and O_1R_2, whose intersection gives the receiver's location.

6. EXPERIMENTAL EVALUATION

We have implemented a prototype system of SpinLight. Figure 13(a) shows the transmitter, containing an LED chip of 3 W, covered by a 3D printed shade whose hollow cells can be covered by opaque paper to create closed cells. The used code of the shade follows Figure 5. By default, we use a single transmitter flashing at the two frequencies depicted in Figure 6. During the experiments, ordinary fluorescent lamps are turned on, which generate both visible and IR light signals at 100 Hz. For convenience of experimentation, we let the transmitter face a wall at a distance of 2.8 m to simulate a ceiling mounted lamp, see Figure 13(b). The light covers a circular area of radius 5.5 m. For 2D localization, the distance to the wall is assumed to be known to the receiver. For 3D cases, we need signal from two transmitters. This is realized by moving the transmitter to two locations 3 m apart, and collecting data separately. The raw data from the two sources are then combined before processing. Note

that the environmental conditions remain the same when the transmitter works at two locations, implying that the mixed data are actually subject to stronger noise and interference than in the case of two transmitters emitting signal simultaneously.

The receiver comprises an MCU and an Intersil ISL29023 sensor [11] for light sensing, which can be configured to sense IR or visible light, and has been used in smartphones such as the Motorola XT882. For the 2D case, 131 points in the covered area are picked with uniform randomness, at which the receiver is tested. For the 3D case, we pick 59 points in the intersection area of the two lamps.

We compare five methods of light localization:

1. SpinLight-2D: the 2D version of SpinLight using a single transmitter;

2. SpinLight-3D: the 3D version of SpinLight using two transmitters;

3. Epsilon [14]: it uses the classic trilateration method based on a Lambertian RSS model, which relates light intensity with distance and angles of the signal [5].

4. Multi-sensor: in contrast to trilateration, this approach uses a single light but multiple light sensors to locate the receiver, based also on the Lambertian model. The approach is discussed in [44], simulated with manual operation in [14], and implemented in [40].

5. RSS Weighted Average (Weighted-Avg): it is an intuitive solution in which the receiver locates itself based on a weighted average of the transmitters' locations in range, where the weight is proportional to the square root of RSS;
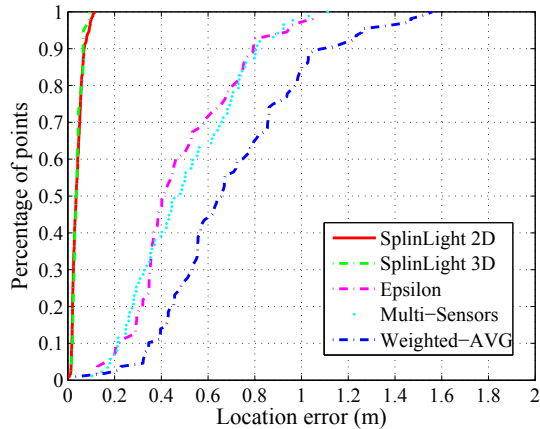
218

Figure 14: CDF of location error of different methods.

Figure 15: Location accuracy for various distance to the transmitters.

Figure 16: Effect of reflection.

For some experiments, we will focus on 2D localization. In this case we simply use 'SpinLight' to represent 'SpinLight-2D'. For Weighted-Avg and Epsilon, we choose three IR LED lights of 8 W, and mount them onto the ceiling (Figure 13(c)). During the experiments, the receiver board is horizontally placed and the light sensor always faces upward. For Multi-sensor, we manually create multiple orientations to simulate multiple sensors, following [14].

6.1 Localization accuracy

Figure 14 presents the cumulative distribution functions (CDFs) of the different methods' location errors, and Table 1 gives some statistics of the errors. The Weighted-Avg method performs the worst, with a median error of 0.76 m, due to the coarse grained proximity detection. Epsilon and Multi-sensor produce median errors of 0.413 m and 0.481 m, respectively, which are comparable with the results in the literature. The improvements come from a more accurate RSS/AoA model, that is, the Lambertian radiation pattern. Compared with these analog signal based methods, SpinLight provides a dramatic improvement in location accuracy with a pure digital design, reducing the median error to within 0.04 m, which is an order of magnitude smaller than the errors of Epsilon and Multi-sensor. Also, it can be seen that both Epsilon and Multi-sensor have a standard deviation an order of magnitude higher than that of SpinLight, because RSS is more sensitive to noise and environmental interference.

6.2 Effect of distance

The receiver's distance to the transmitter largely determines light intensity, hence signal-noise ratio (SNR). For RSS or AoA based localization methods, a lower SNR often introduces greater error in distance or angle estimation. In this experiment, we consider the 2D case of SpinLight, and measure the receiver's distance to the transmitter's projection on the floor (or center of projections when there are multiple transmitters). The distances are divided into intervals of 1 m; for each interval, we calculate the average location error of each method. Figure 15 compares the performance of these methods. We can see that SpinLight performs stably over different distances, while the other methods generally have increasing location error for larger distances to
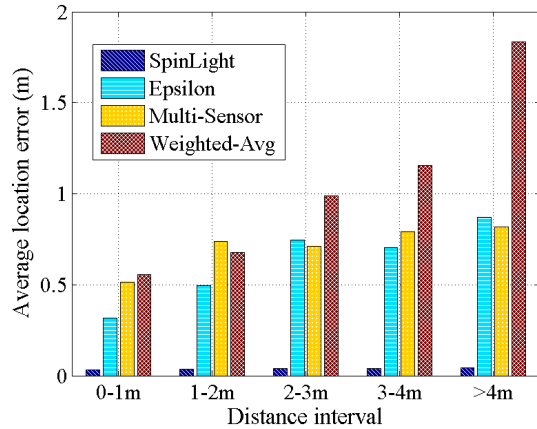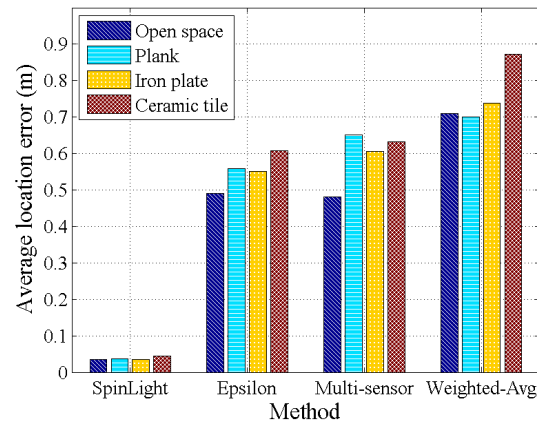
the transmitter. For Epsilon, the location error increases by about three times from within 1 m to beyond 4 m. This is because RSS attenuates according to a square law, which implies that at a relatively large distance, the RSS becomes very small, and the same level of noise and environment interference will cause more uncertainty in distance estimate than at a small distance. The same problem exists for incident angle estimation. For Weighted-Avg, a similar trend can be found, due to the same reason as for Epsilon.

6.3 Effect of reflection

We create various reflection conditions to examine the different methods' performance in more realistic settings. Objects of various types, including no object (open space), plank, non-polished iron plate, and ceramic tile, are considered. The objects are placed near the receiver and about 5 cm farther away from the transmitter, with their plane facing the light, or one of the lights in cases of multiple transmitters (as with Epsilon and Weighted-Avg). Figure 16 shows the average location error under the different conditions. It can be seen that for SpinLight, the location error is relatively stable under the different conditions, mainly be-

	Error statistics (m)					
Method	Min	Median	Mean	Max	95th Prtl	Stdev
SpinLight 2D	**0.004**	**0.036**	**0.041**	0.118	0.092	0.023
SpinLight 3D	0.016	0.038	**0.041**	**0.092**	**0.068**	**0.020**
Epsilon	0.124	0.413	0.490	1.112	0.923	0.230
Multi-sensor	0.103	0.481	0.505	1.050	0.938	0.236
Weighted-Avg	0.184	0.760	0.780	1.05	1.344	0.313

Table 1: Location error statistics of different methods.
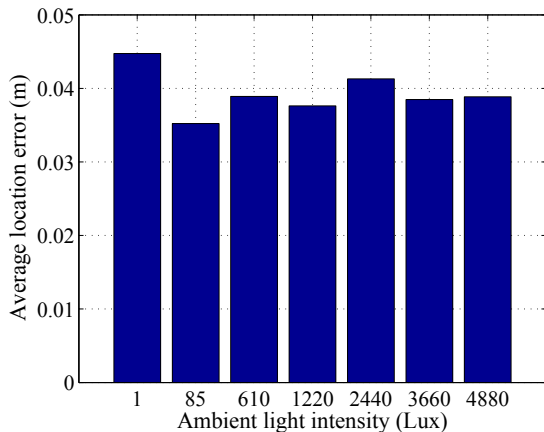
Figure 17: Effect of ambient lights.
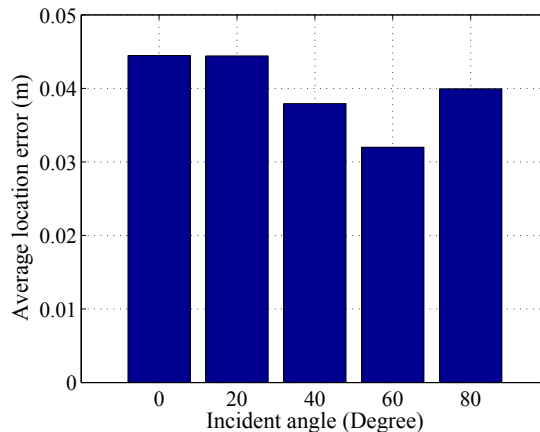
Figure 18: Effect of receiver orientation.

cause the impact of light intensity (RSS) is subdued by the binary decision in SpinLight's localization algorithm. For other methods, the location error increases quite noticeably in the presence of reflection. In particular, the ceramic tile is among the materials that have the greatest impact on location accuracy, possibly because it has a smooth surface and thus a high reflection ratio.

6.4 Effect of ambient lights

In this section we consider how a SpinLight receiver behaves under different conditions of ambient lights. We choose different places in the room at different times of the day, and select seven grades of illuminance (lux): 1, 85, 610, 1220, 2440, 3660, 4880, where 1 lux means a dark room, and 85 corresponds to a room with a few fluorescent lamps on but without daylight. The other grades represent different conditions of illuminance; for example, 4880 lux represents a location near the window on a clear day, but without direct sunlight. Figure 17 shows that for the different conditions, the average location error remains around 0.04 m, with no noticeable impact from variation of ambient lights.

6.5 Effect of orientation

A light sensor is known to be very sensitive to the incident angle of light. For RSS or AoA based localization methods, the incident angle is a key factor in the location model, and has to be carefully accounted for. Past systems use accelerometers and magnetic sensors to estimate the angle of the sensor's surface, which introduce extra errors. For example, due to magnetic interferences from metals and electronic devices, a magnetic sensor inside a building can generate errors up to tens of degrees [14, 40], which makes the

location result very unreliable. We examine SpinLight's sensitivity to the receiver's orientation, by placing the receiver at a fixed location, and adjusting the incident angle from $0°$ (perpendicular to light) to $80°$, with a $20°$ step size. We consider the single-transmitter, 2D case, and select 13 test points across the covered area with various distances to the center. Figure 18 shows the location error for the five incident angles at 13 test locations (colorful bars). Overall, there is no clear correlation between incident angle and location accuracy, which fluctuates between 0.03 m and 0.045 m in most cases. This indicates that SpinLight is robust against receiver orientation, thus provides more flexibility in application than RSS/AoA based solutions.

6.6 Effect of location optimizations

In Section 4, we have introduced how to use specially shaped cells and partial cell recognition to refine the location results obtained by the basic design, in which a receiver takes the center of its cell as its location. In this section we examine how these optimizations improve location accuracy. Figure 19 compares the accuracies of SpinLight with and without the optimizations. We can see the accuracy is improved by 3.9 times to 12.0 times with optimizations, depending on distance. Moreover, the larger the distance to the light source, the larger the error without optimizations. This is because the size of cells grows as the ring number increases, which makes the location obtained more coarse grained.

7. DISCUSSION

In this section we discuss a number of practical issues of the system.

Figure 19: Effect of location optimizations.

7.1 Deployment issues and usability

In our current implementation, an LED light with only 3 W power rate covers an area of around 95 m². Given a higher power level (e.g., 10 W as in [14]), the coverage can potentially extend to hundreds of square meters. This obviates the need for dense deployment of transmitters. The system targets both smartphone users and passive objects localization in LOS scenarios. Today, almost all smartphones have a front-facing light sensor, which can be used by an active user for localization. Sometimes the operating system does not allow the application to sample light signals fast enough; in this case the device driver can be modified to accomplish the goal [40], or a light sensor board can be used, with connection to the phone via the audio jack [14]. For object tracking applications, we envision that a tag can be fixed to the object with exposure to light. Possible application scenarios include robot navigation in unobstructed rooms and product tracking on open assembly lines.

7.2 Localization delay

In our current prototype, the lamp shade does not run at a fast speed due to the restriction of the step motor used, so a receiver takes one to two seconds to perform localization. Here we provide an estimate of the delay in a more advanced implementation. In the worst case, the receiver needs to wait nearly two rotations to obtain its location, thus the localization delay depends on the speed of the step motor. A low cost motor for home appliances such as household electric fans can rotate at a few thousand revolutions per minute [33]. Assuming a moderate rotational speed of 600 rpm, our system would yield a 200 ms delay of localization. There is another restriction from the light sensor's sampling rate. The ISL29023 light sensor has a sampling rate between around 50 Hz to 40 KHz [11], depending on the data granularity (to which our system is not sensitive). Our experiments show that 4096 data samples per revolution produce sufficient accuracy, thus assuming a moderate sampling rate of 20 KHz, the localization delay caused by the two-round rotation is 0.4 s. It is expected that a smaller delay can be attained by optimizing the signal processing algorithm, and by choosing light sensors with a higher sampling rate. This will improve the real-time performance of the system.

7.3 Energy consumption of receiver

The light sensor used for collecting location information consumes very low power. The ISL29023 sensor, for example, consumes 0.25 mW power in normal operation [11]. This is three orders of magnitude lower than a CMOS image sensor [15], and two orders of magnitude lower than a Bluetooth Low Energy (BLE) module in active state [20]. It is even lower than an accelerometer's consumption in continuous operation [1]. For a receiver that comprises a light sensor, a processor, and a BLE module, we expect that the receiver has comparable operation time with an ordinary iBeacon tag, which lasts month to years with a coin cell battery.

8. RELATED WORK

SpinLight is motivated and inspired by research in a few lines. In this section we discuss the related work.

8.1 RF based schemes

A large class of localization techniques are based on RF, including WiFi, cellular radio [32], FM radio [9], although some employ magnetism [10] or ultrasound [7], etc. A commonly used property of RF is RSS, which is easy to obtain on standard wireless devices. However, RSS is not a reliable estimator for distance or location, as it is susceptible to indoor multipath reflections and shadowing. As a result, these solutions, be it propagation model or fingerprint based, provide only low to medium accuracy, typically to a few meters [4, 46]. Various efforts have been made to improve the accuracy. For example, physical layer information is exploited for a better distance estimation [16, 26, 39]. Another commonly used technique of RF localization is AoA [18, 19]. In recent work [41], antenna arrays are used to track RF phase to achieve sub-meter accuracy.

In the domain of RFID, both the AoA and synthetic aperture radar (SAR) approaches have been used for enhancing position accuracy. PinIt [36] leverages the multipath profiles of target tags, as well as of reference tags to generate a small position error, with a highlight on its capability to deal with non-light of sight conditions. However, PinIt requires a highly dense deployment of reference tags (e.g., with a 15 cm spacing). The techniques of PinIt is further used to enable robot navigation and object manipulation in RF-compass [35]. Tagoram [43] provides very fine-grained location for RFID tags by constructing an RF phase hologram with observations from several reader antennas. For general mobility condition, (i.e., with no a prior knowledge of the trajectory), the median position error is 12.3 cm. Tagoram assumes a moving object that simulates an antenna array, thus it does not work for stationary objects.

8.2 Light signal based schemes

Recent research on light signal based localization has been primarily carried out in the context of visual light communications (VLC). Traditional localization principles such as proximity detection, RSS, or AoA have been well explored in various respects. Rajagopal et al. [24] propose to use LED lamps as visual light landmarks by encoding information in the flash patterns of light and exploiting the rolling shutter effect of camera sensors, addressing many practical concerns. ByteLight [2] is a commercial effort that appears to follow a similar approach. For improved accuracy, RSS

information, commonly derived from the Lambertian pattern [5], is used with a multilateration method to locate the receiver [42, 45, 47]. In [22], a hybrid RSS/AoA algorithm is explored for improved location accuracy. In these solutions, the receiver's RSS is sensitive to its orientation, so the receiver needs a fixed orientation (e.g., perpendicularly upward), or needs extra sensors to help rectify the RSS, in order to solve for location. Recently, Luxapose [13] uses AoA and careful image analysis to improve the accuracy to 10 cm, assuming dense deployment of LED lights and high image quality of the camera.

Some other work proposes to use optical sensors instead of cameras as receivers. As optical sensors consume much less power, this approach makes it more practical for small battery powered receivers. The Epsilon [14] system uses an optical sensor to derive RSS from at least three lamps and use the classic trilateration method to solve for location. Epsilon reports position accuracy to about a half meter. In contrast to the multi-transmitter, single receiver approach, the authors in [44] propose to use a single transmitter and multiple receivers for localization. A similar principle has been realized in [40], which offers an accuracy of 40 cm.

Using mobile light sources for positioning has been explored in some early work as well. The Lighthouse location system [25] proposes methods to generate rotating parallel light beams that scan around and help determine a receiver's distance. With three transmitters, a receiver can obtain three distance measurements, and then calculate its location. In the Spotlight system [30], the instrument for localization is an unmanned aerial vehicle, which flies over the receivers and generates encoded light signals (or events). A receiver detects the light event, records the event time sequence, and then sends it to the vehicle for location analysis. SpinLight is similar to these systems in its digital nature, but uses a different design, and works on a distinct principle of localization.

8.3 Other related work

SpinLight can also be compared with the existing work from alternative perspectives. There are some previous systems that also achieve centimeter-level accuracy, such as Active Badge [37], Bat [38], and Cricket [23]. They use a combination of ultrasound and RF sensors for localization, and are expensive and time consuming to deploy. Similar to SpinLight, the work in [19, 36, 43] also considers antenna motions, with differences in application scenarios, signal characteristics, motion style, as well as in performance objectives.

9. CONCLUSION

We have presented SpinLight, a light based localization system for indoor environments. The main innovation is its use of a rotating lamp shade to realize spatial division and coding for the environment. By detecting the flash pattern of the light signal, the receiver is able to identify the spatial beam which it belongs to, followed by optimization techniques to refine its location to a few centimeters. Compared with other light positioning systems, our design has the advantages of being less sensitive to environmental conditions and receiver orientation, and has the benefit of low power consumption. The system could be used in for example robot navigation and product assembly lines. Dead reckoning could be used to fill in the gaps of line of sight.

The transmitter in SpinLight involves a revolving shade, which may introduce instability or increase error after long periods of operation. We will investigate the impact of mechanical vibration on location accuracy, and improve the positioning algorithm to better tolerate the instability. Our future work also includes more careful signal processing for a smaller localization delay, in order to achieve real time positioning in mobile scenarios.

10. ACKNOWLEDGMENTS

References

[1] Analog Devices, Inc. ADXL330: Small, Low Power, 3-Axis ±3g iMEMS Accelerometer (Data Sheet), 2007.

[2] ByteLight. http://www.bytelight.com/

[3] J. Armstrong, Y. A. Sekercioglu, and A. Neild. Visible light positioning: A roadmap for international standardization. *IEEE Communications Magazine*, 51(12), 2013.

[4] P. Bahl and V. N. Padmanabhan. RADAR: An Inbuilding RF-based User Location and Tracking System. In *INFOCOM*, 2000.

[5] J. R. Barry. *Wireless infrared communications*, Springer, 1994.

[6] J. Biswas and M. Veloso. Wifi localization and navigation for autonomous indoor mobile robots. In *IEEE International Conference Robotics and Automation (ICRA)*, 2010.

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

[8] A. Carroll and G. Heiser. An Analysis of Power Consumption in a Smartphone. *USENIX Annual Technical Conference 2010*.

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

[10] J. Chung, M. Donahoe, I. Kim, C. Schmandt, P. Razavi, M. Wiseman. Indoor location sensing using geomagnetism. In *MobiSys*, 2011.

[11] Intersil. The ISL29023 Datasheet. July 17, 2012.

[12] A. Jovicic, J. Li, and T. Richardson. Visible light communication: Opportunities, challenges and the path to market. *IEEE Communications Magazine*, 51(12), 2013.

[13] Y.-S. Kuo, P. Pannuto, K.-J. Hsiao, and P. Dutta. Luxapose: indoor positioning with mobile phones and visible light. In *Mobicom* 2014.

[14] L. Li, P. Hu, C. Peng, G. Shen, and F. Zhao. Epsilon: A Visible Light Based Positioning System. *Proc. NSDI 2014.*

[15] R. LiKamWa, B. Priyantha, M. Philipose, L. Zhong and P. Bahl. Energy Characterization and Optimization of Image Sensing Toward Continuous Mobile Vision. *Proc. of MobiSys 2013.*

[16] A. T. Mariakakis, S. Sen, J. Lee, K.-H. Kim. SAIL: single access point-based indoor localization. In *MobiSys*, 2014.

[17] L. Ni, Y. Liu, C. Yiu, and A. Patil. LANDMARC: Indoor Location Sensing Using Active RFID. In *WINET*, 2004.

[18] D. Niculescu and B. Nath. Ad hoc positioning system (APS) using AOA. *IEEE INFOCOM*, 2003.

[19] D. Niculescu and B. Nath. Vor base stations for indoor 802.11 positioning. *ACM Mobicom*, 2004.

[20] Nordic semiconductor. *nRF8001 Product Specification 1.2.* 2013.

[21] Opto-isolator. http://en.wikipedia.org/wiki/Opto-isolator.

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

[23] N. Priyantha, A. Chakraborty, and H. Balakrishnan. The Cricket location-support system. In *ACM Mobi-Com*, 2000.

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

[25] K. Romer. The Lighthouse Location System for Smart Dust. *ACM MobiSys 2003.*

[26] S. Sen, J. Lee, K.-H. Kim, and P. Congdon. Avoiding Multipath to Revive Inbuilding WiFi Localization. In *ACM Mobisys*, 2013.

[27] G. Sterling and D. Top. Mapping the Indoor Marketing Opportunity. *Opus Research Report.*

[28] STMicroelectronics. STM32 MCU family – 32-bit Flash microcontrollers powered by ARM Cortex^{TM} – M3, May 2008.

[29] STMicroelectronics. LIS33DE datasheet, April 2009.

[30] R. Stoleru, T. He, J. A. Stankovic, D. Luebke. A High-Accuracy, Low-Cost Localization System for Wireless Sensor Networks. *ACM SenSys 2005.*

[31] S. Tarzia, P. Dinda, R. Dick, G. Memik. Indoor Localization without Infrastructure using the Acoustic Background Spectrum. In *ACM MobiSys*, 2011.

[32] A. Varshavskya, E. d. Lara, J. Hightower, A. LaMarca, V. Otsason. GSM indoor localization. *Pervasive and Mobile Computing*, 2007.

[33] Vexta PK245 technical data sheet. Velmex, Inc.

[34] VLC for location, positioning and navigation. http://visiblelightcomm.com/vlc-for-location-positioning-and-navigation/

[35] J. Wang, F. Adib, R. Knepper, D. Katabi and D. Rus. RF-Compass: Robot Object Manipulation Using RFIDs. In *ACM Mobicom*, 2013.

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

[37] R.Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Trans. on Information Systems*, 10(1):91–102, Jan. 1992.

[38] A. Ward, A. Jones, and A. Hopper. A new location technique for the active office. *IEEE Personal Communications*, 4(5):42–47, Oct. 1997.

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

[40] B. Xie, G. Tan, Y. Liu, M. Lu, K. Chen, T. He. LIPS: A Light Intensity Based Positioning System For Indoor Environments. http://arxiv.org/abs/1403.2331

[41] J. Xiong and K. Jamieson. Arraytrack: A fine-grained indoor location system. In *Usenix NSDI*, 2013.

[42] S.-H. Yang, E.-M. Jeong, D.-R. Kim, H.-S. Kim, Y.-H. Son, and S.-K. Han. Indoor three-dimensional location estimation based on LED visible light communication. *Electronics Letters*, 49(1), January 2013.

[43] L. Yang, Y. Chen, X.-Y. Li, C. Xiao, M. Li, and Y. Liu. Tagoram: Real-Time Tracking of Mobile RFID Tags to High Precision Using COTS Devices. In *Proc. of Mobicom*, 2014.

[44] S-H. Yang, E-M. Jung, S-K. Han. Indoor location estimation based on LED visible light communication using multiple optical receivers. *IEEE Communications Letters*, 17(9), Sept. 2013.

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

[46] M. Youssef and A. Agrawala. The Horus WLAN Location Determination System. In *MobiSys*, 2005.

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