CELLI: Indoor Positioning Using Polarized Sweeping Light Beams

Yu-Lin Wei†, Chang-Jung Huang†, Hsin-Mu Tsai†, Kate Ching-Ju Lin‡
†Department of Computer Science and Information Engineering, National Taiwan University, Taiwan
‡Department of Computer Science, National Chiao Tung University, Taiwan
{r03922027,r02922061,hsinmu}@csie.ntu.edu.tw, katelin@cs.nctu.edu.tw

ABSTRACT

Existing visible light positioning (VLP) systems leverage the high image resolution of a receiving camera to support high positioning accuracy. However, power-hungry cameras might not always be applicable for many scenarios, such as smart factories, in which small objects require to be accurately localized and tracked. In this paper, we introduce CELLI, an indoor VLP system that only uses a single luminary as the transmitter and requires only a simple light sensor to achieve an extremely high accuracy with centimeter-level error. The key idea is to provide the spatial resolution capability from the transmitter instead of the receiver, so that the complexity of the receiver can be minimized. In particular, a small LCD is installed at the transmitter to project a large number of narrow and interference-free polarized light beams to different spatial *cells*. A receiving light sensor identifies its located cell by detecting the unique polarization-modulated signals projected to that cell. CELLI further incorporates a number of novel designs to overcome the technical challenges such as reducing the positioning latency, which is typically limited by the long optical response time of an LCD, and transforming a cell coordinate to the global 3D position using only a single light. We have prototyped our design using off-the-shelf optical and electrical components, and experimentally shown that CELLI achieves a median 3D positioning error less than 11.8 cm and a median 2D positioning error to less than 2.7 cm.

1. INTRODUCTION

Indoor positioning enables location-based services for a wide range of commercial applications [8]. Seeking to maximize their efficiency, smart factories in the next decade requires assets and personnel to be constantly tracked, automated vehicles navigating themselves to transport goods, and robots in the assembly lines manipulating objects with precision. To turn these into reality, it is essential to have a positioning technology with centimeter-level accuracy, while being applicable even for small objects or able to be directly attached to existing equipment. The use cases set stringent energy consumption, computation, and cost constraints that need to be overcome by the positioning technology.

MobiSys'17, June 19-23, 2017, Niagara Falls, NY, USA

ISBN 978-1-4503-4928-4/17/06…$15.00

DOI: http://dx.doi.org/10.1145/3081333.3081352

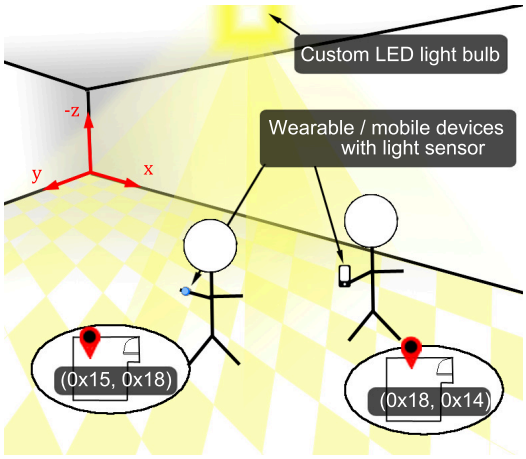

Figure 1: Positioning with CELLI. CELLI provides accurate positioning using one light bulb and a simple light sensor.

While RFID-based positioning technologies [18–22, 26] are especially suitable for object localization or tracking, their accuracy, however, could degrade significantly in a multipath-rich environment, such as a factory with metallic objects everywhere. Recent works [12, 14–16, 24] have demonstrated that visible light signals can be utilized to get rid of the critical multipath problem and deliver accurate positioning estimation with sub-meter accuracy. In these systems, taking advantage of the high image resolution, a commodity camera is able to capture at least three transmitting luminaries in the same image and leverages some geometric properties to estimate its location. However, the need for multiple captured luminaries prevents those designs from being applied on small objects, which might only be able to be equipped with a simple light sensor (i.e., single-pixel photodiode). Also, supporting positioning services with dense modified lighting equipments creates additional difficulty for widely deployment.

In this paper, we propose CELLI, a novel visible light positioning system that can perform centimeter-level 3D positioning with only a single transmitting luminary and a simple receiving light sensor. The key innovative idea of CELLI is to transfer the spatial resolution capability from the receiver to the transmitter. Making use of a small liquid crystal display (LCD) panel, CELLI's transmitter projects a large number of fine-grained, parallel, and yet interference-free light beams toward the service area. As a result, CELLI only needs a small light sensor, instead of a commodity camera, as the receiver, which can be easily integrated with small objects or devices. At a certain location, the light sensor receives the light beam from a particular pixel of the LCD, which carries

the information about *the coordinate of a fine-grained cell*. The positioning signals are modulated using the polarization direction of a light source. Hence, human eyes perceive non-flickering white light from the transmitting luminary, preserving its main function – illumination.

Realizing the idea, however, is not trivial, due to the following challenges: First, CELLI's transmitter changes the states of LCD pixels to modulate the light beams with positioning information. However, the long optical response time of an LCD, i.e., the time for a pixel to completely change between states, could take hundreds of milliseconds. If we naïvely transmit the bits of the location identifier, it would take several seconds to complete positioning once, due to the large number of state changes (i.e., bit flipping). To address this issue, we modulate the positioning signals using a modified *pulse position modulation* (PPM) scheme, which incorporates only 4 state changes in a complete positioning cycle. In particular, taking advantage of PPM's single pulse property, we propose a *pre-switching* mechanism to reduce the positioning cycle from 1.8 seconds to 0.6 seconds, but still keep each LCD pixel staying at a state for a sufficient time to settle, maintaining a high signal strength for accurate positioning.

In addition, in CELLI, the light sensor only receives the 2D coordinate from a transmitting light. This renders conventional triangulation and trilateration techniques inapplicable and makes 3D positioning a challenging task. To further convert the received 2D coordinate to the actual 3D position, the system should be able to estimate the receiver height (the projected distance along the Z axis in Fig. 1) as well as the corresponding size of the projected cell. To achieve this goal, CELLI employs a novel *dual-lens optical design* to project two independent sets of light beams to the service area. A receiver would capture two light beams respectively from both sets of projection, enabling it to accurately estimate the Z-axis distance and, thereby, its 3D position, using an angle-of-arrival based technique.

Finally, since CELLI uses the polarization direction of a light beam to represent the positioning information, the received signal strength of an angle-of-polarization-modulated signal could drop significantly [24] when the receiver rotates to certain orientations with respect to the transmitter. To ensure that a receiver gets a signal strength high enough for accurate positioning, CELLI incorporates an electronic polarizer in front of the receiving light sensor. Combined with a simple search algorithm, the electronic polarizer adaptively rotates its polarization direction that best matches the incoming light polarization and maintains sufficient received signal strength. CELLI's receiver can, hence, perform accurate positioning in any of the 360 degree orientations.

We have implemented a prototype of CELLI using commodity optical compoments and development boards, and tested the impact of different types of LCDs. The key findings from our testbed experiments are as follows:

- CELLI combines PPM with a simple but effective pre-switching scheme to shorten the positioning cycle from 1.8 seconds (required by the traditional bit-pattern approach) to 0.6 seconds, while keeping mostly the same positioning accuracy.
- CELLI resolves the problem of intensity contrast drop due to arbitrary rotation by applying an additional electronic polarizer at the receiver. By dynamically adapting the polarization direction, the receiver can always observe a similar intensity contrast, independent of its orientation.
- The interference from a neighboring cell 2 cm further away is negligible. Taking advantage of the dense narrow yet interference-free light beams generated by a high-resolution LCD, CELLI produces a 2D positioning error no larger than 3.3cm.

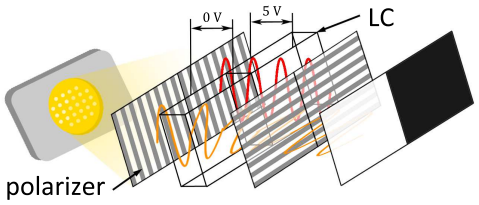

Figure 2: LCD operation principle. Each LC pixel twists the polarization direction of the light according to the applied voltage.

- With the novel dual-lens design, CELLI can further support 3D positioning with a median error of only 11.8 cm.

2. LCD PRIMER

2.1 Principle of Operation

LCD is a widely-used display technology on all types of screens, e.g., computers, mobile devices, TV, projectors, etc. An LCD is composed of a liquid crystal (LC) array between two polarization filters that are *perpendicular to each other*, as shown in Fig. 2. A polarizing filter (or polarizer) is an optical filter that allows only one particular polarization of the incoming light to pass through while filtering out all the remaining polarizations. Therefore, the first polarizer of an LCD can convert an unpolarized light source to a polarized light beam. Each LC cell (or pixel) in the array then twists the polarization direction of the incoming light according to the voltage applied. The intensity of the polarized light passing through the second polarizer, I_θ, varies with the angle between the polarization direction and the second filter, θ, which can be expressed as

$$I_\theta = I_0 \cos^2 \theta, \tag{1}$$

where I_0 is the initial intensity of the polarized light before passing the second filter [10]. Hence, by controlling the voltage applied to each LC pixel, we can tune the angle θ and change the intensity a narrow light beam accordingly. For example, when no voltage is applied, the pixel twists the polarized light by 90 degrees, making the polarization direction of the light become parallel to the second polarizer. The polarized light can, hence, completely pass through the LCD. On the contrary, when full voltage is applied, the LC pixel does not twist the polarization direction of light, and, hence, the polarized light stays perpendicular to the second polarizer. In this case, the polarized light will be fully blocked and a black pixel is displayed. CELLI leverages the spatial resolution of an LCD to simultaneously emit a large number of well-modulated narrow light beams and achieve find-grained positioning.

2.2 Types and Characteristics

Most LCDs on the market can be categorized into *passive matrix* and *active matrix* according to the control circuit design [17]. Passive matrix LCDs, such as STN LCDs, cannot adjust the voltage of each pixel independently, but can only control a column of pixels or a row of pixels at a time. As a result, a passive LCD suffers from a lower contrast ratio and long response time, but features low price, high optical transparency and energy efficiency. For example, a common low-cost STN LCD (~$10 each) can support a resolution up to 128 by 64 pixels.

On the contrary, active matrix LCDs are able to control each pixel of the LC array independently using a thin transistor film (TFT) array over the LC layer. Active LCDs usually support a greater contrast ratio with the cost of lower optical transparency and

LCD Type	Passive Matrix	Active Matrix
Addressing	row/column control circuit	thin film transistor
Energy Consuming	Low	High
Cost	Low	High
Transparency	High	Low
Contrast Ratio	Low	High
Response Time	Long	Short

Table 1: Comparison between passive and active matrix LCDs.

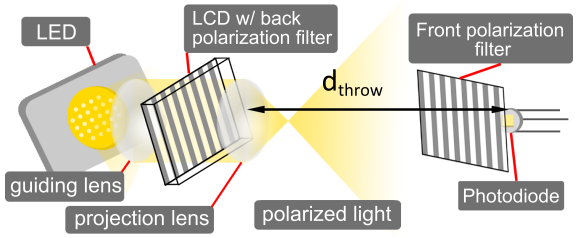

Figure 3: CELLI system design. CELLI utilizes a LCD to generate polarized light beams, each of which is then projected to a small cell in the service area.

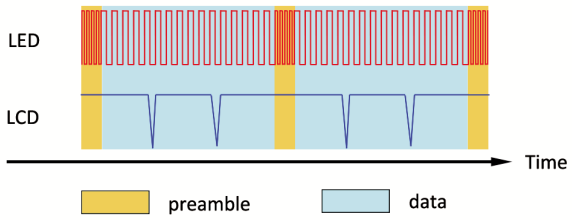

Figure 4: Positioning cycles. CELLI modulates the LED back-light using the preamble frequency to indicate the beginning of a positioning cycle.

higher energy consumption required by the additional TFT layer. Also, the response time of active LCDs is much shorter than passive LCDs since the delay of changing the state of one LCD pixel is smaller than that of updating a group of pixels or the entire LCD frame. The detailed comparison between passive and active LCDs is summarized in Table 1, and we will check in Sec. 5.1 CELLI's positioning performance when different types of LCDs are applied.

3. SYSTEM DESIGN

3.1 Overview

CELLI aims to provide accurate indoor positioning using only a single luminary. On the transmitter side, we place an LCD between two lenses to modulate the light from an LED light source, which emits unpolarized light, as shown in Fig. 3. The *guiding lens*, which sits in front of the LED light source, refracts the light to enter the LCD in parallel. Then, each pixel of the LCD modulates the angle of polarization of the incoming light to generate a narrow and polarized light beam. The *projection lens* refracts the parallel light beams out of the LCD, and projects them toward non-overlapping locations, denoted as *cells* hereafter, in a positioning service area. By carefully selecting this lens, CELLI can flexibly control the tradeoff between accuracy and the dimensions of the service area. We detach the second filter from the LCD, but attach it in front of the receiving light sensor, such that, between the transmitter and the receiver, only the angle of polarization varies. As human eyes are insensitive to this variation, flickers are not observed. The second polarizer makes sure that the light sensor can observe the intensity changes induced by the change of each pixel's polarization angle. In particular, we explicitly adjust the states of the LCD's pixels to generate positioning signals such that the receiving light sensor at different cells observes different and unique bright-dark sequences over time (see Sec. 3.2 for the details). Taking advantage of the high spatial resolution of the LCD, CELLI is able to achieve fine-grained positioning.

Realizing such cell positioning, however, requires us to overcome two practical challenges. First, with CELLI, the receiver can only identify its coordinate, instead of the absolute position. To convert coordinate to position, the receiver further needs the information about its *height*, which is defined as the *throw distance* between the transmitter and the cell plane of the receiver, as shown in Fig. 3. To allow a receiver to estimate its own height, we introduce a dual-lens design at the transmitter such that each receiver can capture two different positioning signals from two projection lenses (see Sec. 3.3). Once detecting two different bright-dark sequences, the receiver can leverage some fundamental geometric properties to estimate its height and thereby its absolute location. Second, the receiver may be rotated to an arbitrary orientation during positioning. Since rotation changes the angle between the LCD at the transmitter and the polarizer installed at the receiver, the bright-dark pattern might become indistinguishable at the receiver at certain angles [24], resulting in received signal strength insufficient for positioning. We resolve this issue by replacing the receiver's polarizer with a liquid crystal cell (LCC), i.e., a "single pixel LCD", which acts as an electronic polarizer whose orientation can be dynamically adjusted. With the LCC, the receiver can search for the polarizer direction of its LC that best matches the polarization of the incoming signal, maximizing the contrast of the bright-dark signals, and thereby improving detection accuracy (see Sec. 3.4).

3.2 Coordinate Detection

CELLI partitions time into disjoint positioning cycles, as shown in Fig. 4. Each positioning cycle can be further divided into T time slots. CELLI modulates the LED back-light using two frequencies: the preamble frequency and data frequency. The preamble frequency is transmitted for a small period of time, e.g., 10 ms, at the beginning of each positioning cycle, as the reference starting time before every round of positioning data signals, which are then sent with the back-light modulating at the data frequency. A receiver starts the positioning procedure by correlating the received signal with the preamble carrier in order to detect the preamble and tag the reference starting point.

Modulating positioning signals. To transmit a unique bright-dark sequence toward each cell, CELLI utilizes the characteristics of the LCD mentioned in Sec. 2.2 to realize *Pulse Position Modulation* (PPM) [9], which transmits only one or a few signal pulses in the T possible time slots. In particular, each pixel in the x-th column and the y-th row is assigned a coordinate (x, y) and broadcasts the positioning messages to the corresponding cell. To represent (x, y) using PPM, CELLI partitions the duration of a position cycle into the *column scanning* phase and the *row scanning* phase. Let T_x and T_y denote the durations of the column and row scanning phases, respectively, and, hence, $T = T_x + T_y$. Each pixel then switches

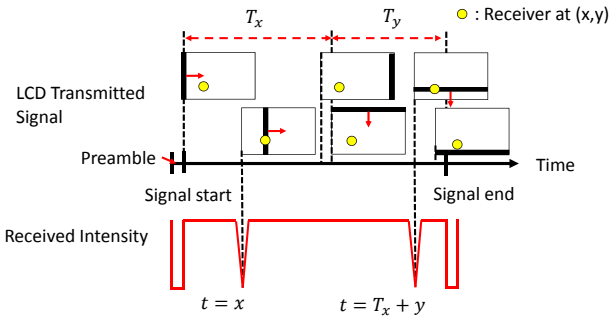

Figure 5: CELLI's PPM modulation. The transmitter delivers PPM signals by sweeping a strip horizontally and then vertically.

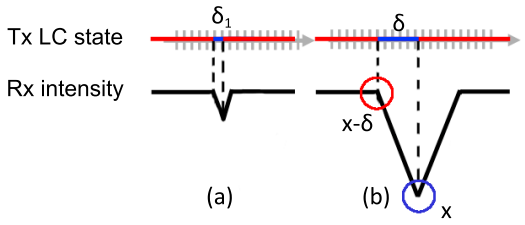

Figure 6: Pre-switching for increasing the intensity contrast. CELLI lets each pixel pre-switch its state so as to ensure complete state transition.

its polarization angle $\theta_{x,y}(t)$ at time slot t as follows:

$$\theta_{x,y}(t) = \begin{cases} 90^\circ & \text{, if } t = x \text{ or } t = T_x + y \\ 0^\circ & \text{, otherwise.} \end{cases} \quad (2)$$

A nice property of this PPM scanning scheme is that each pixel rarely changes its state. Specifically, the traditional bit-pattern scheme uses black and white states as bit 0 and bit 1 to represent the binary identifier of each cell, and, hence, needs to frequently switch the state of each pixel between "0" (dark) and "1" (bright). However, CELLI's PPM allows each pixel (x, y) to switch its state only *four times*, i.e., at time slots x and $T_x + y$ (from 0° to 90°) and at time slots $x + 1$ and $T_x + y + 1$ (from 90° to 0°).

Assume that the receiver's polarizer is placed with an angle of 0°. Then, it receives the dark intensity at only time slots x and $T_x + y$, but receives the bright intensity at all the remaining time slots, as shown in Fig. 5. Hence, with PPM, the positioning signals sent by the transmitter is visually similar to the effect of sweeping a vertical strip column by column, within the column scanning phase T_x, and then sweeping a horizontal strip row by row, within the row scanning phase T_y, as shown in Fig. 5. However, for a receiver to be able to detect the dark signal, the duration of each time slot, denoted by T_{switch}, should at least be larger than the optical response time. The main reason is that each transmitting LCD pixel cannot *immediately* change it polarization state from one angle to another. Hence, if a time slot is too short, the contrast between the bright and dark intensities could be insufficient to distinguish, as shown in Fig. 6(a), since the transmitter might not completely change its angle from 0° to 90°.

To further shorten the positioning cycle without sacrificing the positioning resolution, we need to design an approach to reduce the duration of each time slot, while maintaining the contrast of the bright-dark signals. To achieve this goal, our solution is to shorten the duration of a state transition time-slot T_{switch} but, alternatively,

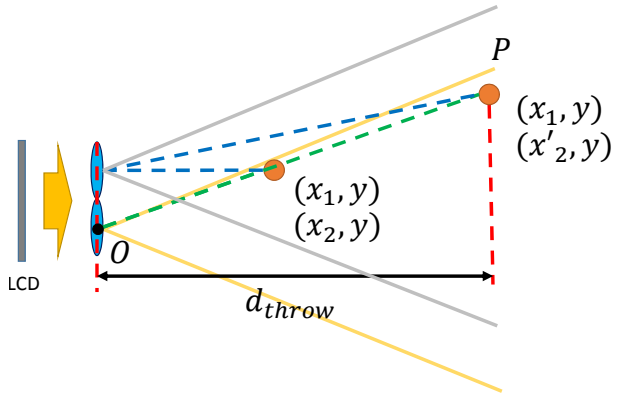

Figure 7: Height estimation by dual-lens. A receiver could detect the same 2D coordinate at different heights. Our dual-lens design allows a receiver to estimate its height even if it is served by only a single transmitting light.

modulate the positioning signals as follows:

$$\theta_{x,y}(t) = \begin{cases} 90^\circ & \text{, if } x - \delta < t \le x \text{ or} \\ & (T_x + y - \delta) < t \le T_x + y \\ 0^\circ & \text{, otherwise.} \end{cases} \quad (3)$$

We can see from the above equation that each pixel (x, y) *pre-switches* to the 90° state at time slot $x - \delta + 1$ and $T_x + y - \delta + 1$, respectively, and stays in this state for δ time slots. To maximize the contrast of the bright-dark signals, we should carefully pick the value of δ such that the total time a pixel stays in the 90° state is larger than its response time. Fig. 6(b) illustrates that this pre-switching scheme allows each transmitting pixel (x, y) to gradually switch to the *actual* 90° state, i.e., reaching the darkest intensity, at time slots x and $T_x + y$, respectively. Therefore, the receiver can still detect the deep of the received signals and estimate its location accurately. With this pre-switching, CELLI now reduces the duration of each time slot, thereby shortening the whole positioning cycle. Therefore, CELLI can achieve fine-grained positioning resolution even with some cheap and low-end LCDs with larger response time (i.e., longer latency for a state transition).

Combining PPM with bit-pattern. While TFT LCDs are capable of updating the state of each pixel independently, some STN LCDs can only update *a group of pixels* at a time. Hence, if we incorporate an STN LCD in the transmitter, CELLI's PPM should treat a group of pixels as a coordinate. Each pixel, again, sends the PPM signals according to its assigned coordinate. However, such group-based scanning can only provide block-level positioning accuracy. To overcome this limitation, we further adopt a hybrid approach that combines CELLI's PPM with the traditional bit-pattern scheme. That is, to achieve pixel-level accuracy, after sending our PPM signals, each pixel within a group further delivers a unique assigned bit-pattern to the corresponding cell. It is worth noting that, since the bit-pattern scheme cannot take advantage of the previously mentioned pre-switching scheme due to its frequency state changes, the duration of a time slot needs to be increased to allow sufficient time for pixels to change their states.

3.3 From 2D Coordinate to 3D Positioning

CELLI's PPM signals only allow receivers to identify their coordinates, instead of absolute locations. Consider the example shown in Fig. 7. If a receiver identifies its coordinate (x_1, y), it still does not know what its exact location is since any points along the line

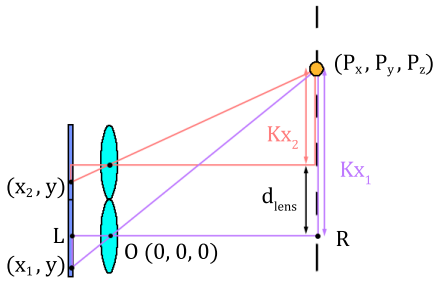

Figure 8: Converting the coordinate to the actual position. The receiver leverages some geometric properties to estimate its height, and thereby 3D position.

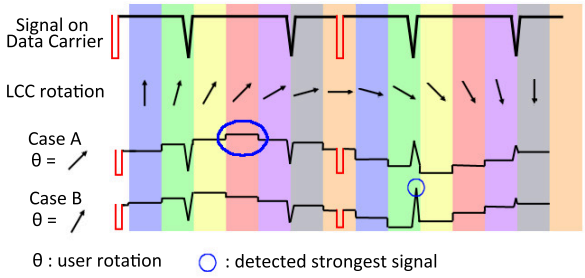

Figure 9: Polarization adaptation using LCC. CELLI uses a LCC to search for a proper polarization angle that gives a high intensity contrast.

segment $\overline{OP}$ receive the same signals from the LCD pixel (x_1, y). To enable a receiver to convert its coordinate to the location, we introduce a dual-lens design at the transmitter. In particular, we replace the original single projection lens with two side-by-side lenses, each of which is covered by half of the pixels of the LCD. With dual lenses, the transmitting light projects to two areas, mostly overlapped with each other, on a surface, as shown in Fig. 7. The LCD pixels, thus, are also partitioned into two sets. Each set has its own coordinate system, with the origin located at the center of each lens. Then, each pixel still broadcasts the positioning signal according to its assigned coordinate using Eq. 3. With these two coordinate systems, the receiver now captures two dark signals in the *column scanning* phase, each of which is from one of the two projection lenses[1]. We can observe from Fig. 8 that, though any point along $\overline{OP}$ receives the same coordinate from the bottom lens, it, however, receives a different coordinate from the top lens.

With such a dual-lens design, a receiver can use the two detected coordinates, (x_1, y) and (x_2, y), from the two lenses, to calculate its global 3D location (P_x, P_y, P_z). Here, we assume the origin of the global coordinate system locates at the center of the bottom lens. To find P_x and P_y, we define K as an unknown scaling factor representing the conversion ratio of the actual physical size of a cell to the unit of our coordinate system, which is determined by the receiver height, P_z. That is, the scaler K is a function of the height, i.e., $K = f(P_z)$. Then, the actual position can be represented by

$$P_x = Kx_1 = Kx_2 + d_{\text{lens}} \text{ and } P_y = Ky, \tag{4}$$

which can be rewritten as

$$K = d_{\text{lens}} \frac{1}{x_1 - x_2}, \tag{5}$$

where d_{lens} is the distance between the origins of the two lenses. Thus, K can be calculated from the detected values of x_1 and x_2, if d_{lens} is known.

The next question is how do we know the value of the constant d_{lens}? Directly measuring the distance between the two lenses might generate measuring error that introduces additional inaccuracy. However, we can see from Eq. 5 that the unknown scaling factor K is a linear function of $\frac{1}{x_1 - x_2}$. This also implies that, since all the cells on a service plane at distance P_z have the same scaling factor K, they must observe the same value of $x_1 - x_2$. Therefore, to more precisely determine the coefficient d_{lens}, we instead measure the true values of K and $\frac{1}{x_1 - x_2}$ at a few known receiver positions, and adopt linear regression to estimate d_{lens}.

[1]Note that, with our dual-lens design, a receiver still receives only one y-axis coordinate since the two lenses are placed side-by-side horizontally.

Given that the value of K is now known, the receiver can calculate its global position along the first two dimensions, P_x and P_y, based on Eq. 4. The receiver can further make use of some geometric properties to learn P_z, which is equal to the throw distance between the LCD origin and the plane where the receiver is located, denoted by $\overline{OR}$ as shown in Fig. 8. Using the property of similar triangles, we can estimate the distance $\overline{OR}$, i.e., the position along the z-axis P_z, by

$$P_z = \overline{OR} = \left(\frac{\overline{OL}}{d_x}\right) K, \tag{6}$$

where $\overline{OL}$ is the distance between the LCD plane and the lens center and d_x is the size of a pixel of the LCD along the x-axis. Again, the distance $\overline{OR}$ is a linear function of K, and we adopt linear regression to learn the constant coefficient $\frac{\overline{OL}}{d_x}$, which is independent of the receiver height, in a similar way[2]. In summary, by estimating the scaling factor K, the receiver can easily learn its three-dimensional global position (P_x, P_y, P_z) based on Eq. 4 and Eq. 6.

3.4 Dealing with Rotation

So far, we assume that the orientation of the polarizer in front of the receiving light sensor perfectly matches the orientation of the LCD in the transmitter (i.e., the relative angle $= 0°$). Thus, when the transmitter alternates the polarization angle of the emitted light beam between $0°$ and $90°$, the receiver can observe the maximum difference between received intensities (i.e., contrast) of these two states, according to Eq. 1. However, in reality, a receiver might have an arbitrary rotation with respect to the transmitter. Since the received intensity is determined by the inter-polarization angle between the polarization angle of the emitted light beam and the receiver's polarizer, with this additional rotation, the angle no longer alternates between $0°$ and $90°$. In the most critical case when the receiver is rotated to a relative angle of $45°$, the inter-polarization angle stays the same at $45°$ when a pixel is set to either a polarization angle of $0°$ or $90°$. As a result, the receiver observes the same intensity in both the bright and dark states and, therefore, can hardly recover the positioning signal.

To address this problem, we replace the receiver's polarizer with a *liquid crystal cell* (LCC) [2], i.e., a one-pixel LCD, which serves as an electronic polarizer. By applying different voltages to the LCC, we can rotate the polarizer orientation to any desirable angle

[2]This linear relationship will be destroyed if the conjunction point of the two lenses is not exactly aligned with the center of the LCD. This offset can be easily calibrated by geometric transformations. We skip the details due to space limitation.

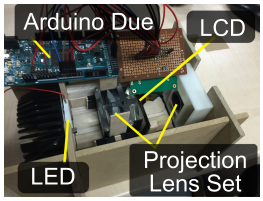

(a) Transmitter

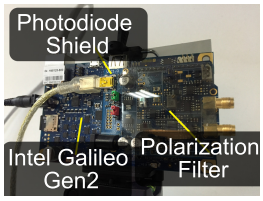

(c) Receiver

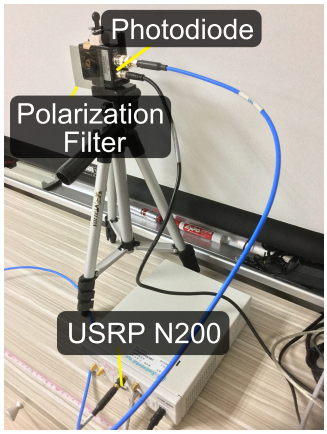

(b) Channel measurement receiver

Figure 10: STN prototype

between 0° and 180°.[3] Hence, by tuning the voltage of the LCC, we can search for a proper polarization angle that maximizes the contrast of the received intensities. To reduce the search complexity, in our implementation, we use only 11 LCC rotation angles. In addition, since it is time-consuming to receive the signal of a complete positioning cycle for every tested LCC rotation angle, to further reduce the search time, we apply the following optimization. For each tested rotation angle, the receiver only measures the maximum received intensity during a small fixed period of time, e.g., 50ms, which is much shorter than a positioning cycle. Then, after testing all candidate angles, the receiver rotates the LCC by the tested angle that reaches the maximum intensity during the whole testing interval, as shown in Fig. 9. The rationale behind this design is that the identified angle should be the one that switches the receiver's polarization such that the relative angle between its polarization direction and the LCD's polarization is closest to either 0° or 90°, thereby producing the maximum contrast.

Note that, due to our PPM signal design, for most of the time, CELLI's transmitter emits light beams with a polarization angle of 0°. Hence, the identified angle producing the maximum intensity is usually the one that best aligns the LCD's orientation and the direction of the receiver's electronic polarizer (i.e., relative angle $= 0^\circ$), e.g., case A in Fig. 9. Then, the receiver can apply our decoding algorithm to detect the dark signals and estimate its location without any change. However, for a very special case where the receiver happens to observe the maximum intensity when the transmitting light beam is set to the 90° state, i.e., case B in Fig. 9, the relative angle between the LCD and the orientation of the receiver's electronic polarizer is approximately 90°. In this case, the receiver now observes the dark signals for most of time, while seeing the bright signals only when the transmitting light beam is switched to the 90° state. This, however, can be easily solved by taking the inverse of the bright-dark pattern, i.e., $y' = y_{\max} - y$, where y denotes the original received signals and $y_{\max}$ is the maximum intensity of the received signals. After such inversion, the receiver can again apply the normal decoding algorithm to find its location.

[3]A commodity LCC usually only supports rotation up to 90°. To further extend the supported range, we place two LCCs back-to-back to rotate the polarizer to any angle between 0° and 180°.

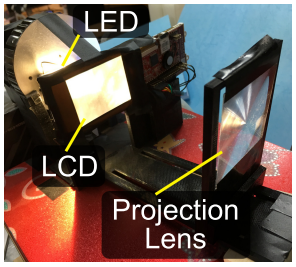

(a) Transmitter

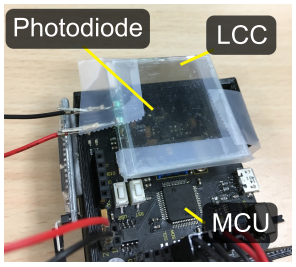

(b) Receiver

Figure 11: CELLI prototype

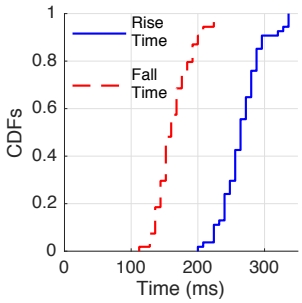

(a) STN

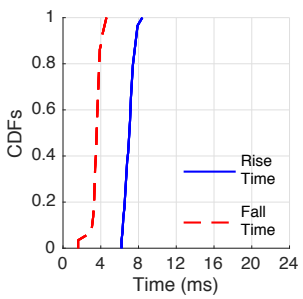

(b) TFT

Figure 12: Response time comparison. The response time of some LCDs is quite large, meaning that reducing the positioning latency is a critical challenge.

4. IMPLEMENTATION

We implement two prototypes of CELLI's transmitter (Tx), one using a passive LCD and the other using an active LCD to check how their characteristics (discussed in Sec. 2) affect CELLI's positioning accuracy. For the passive LCD, we use the STN (supertwisted nematic) LCD with a resolution of 128×64 pixels from NewHeaven Display [3], and control it by an Arduino development board, as shown in Fig. 10(a). For the active LCD, we use the TN (twisted nematic) TFT LCD with a resolution of 320×240 pixels and a controller built in from 4D Systems [1]. Both prototypes use XLamp® CXA2540 as the backlight LED, as shown in Fig. 11(a).

At the receiver side, we implement three different prototypes according to the characteristics of the optical channel. All the three receiver prototypes consist of a photodiode with human eye sensitivity, a polarization filter and a microprocessor. The first prototype is designed for benchmark channel measurement, where we use an experimental photodiode [7] as the receiver to increase the precision of channel characterizations, as shown in Fig. 10(b). For the STN Tx, we use SFH206K [4] as the receiving photodiode. Since the positioning cycle is longer, we choose Intel Galileo Gen2, which provides a larger memory size and computational power, as the motherboard, as shown in Fig. 10(c). Finally, since the TFT Tx features a shorter positioning cycle, we use a microprocessor [6] combined with a photodiode [5] to implement real-time positioning, as shown in Fig. 11(b).

5. RESULTS

In this section, we first conduct benchmark measurements to characterize two common types of LCDs, and check the feasibility of CELLI's dual lens design. Experiments are carried out in a typical conference room with normal daylight and all 12 light fixtures turned on. Each light fixture has four 14W T5 fluorescent tubes.

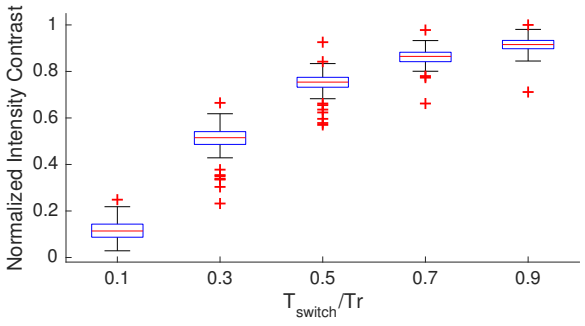

Figure 13: Impact of the state transition duration on the received intensity contrast. The intensity contrast drops significantly if an LCD pixel is not given sufficient time to change its state.

This represents a worst-case indoor scenario where both daylight and light fixtures add noises to the received signal. For ease of experiments, we place the transmitter horizontally, similar to how a projector is usually placed, and setup the receiver on the projected surface, set at distances 175 cm and 225 cm, respectively, away from the transmitting light. The two tested distances correspond to the typical distances between the ceiling and human shoulder or chest, respectively. The performance metrics used in our evaluation are defined as follows:

- *Normalized intensity:* It equals the ratio of the received intensity of a light signal to the maximum light intensity.
- *Signal strength:* The received raw signals are converted to the frequency domain signals via FFT. The signal strength is defined as the amplitude of the peak frequency.
- *Intensity contrast:* It is defined as the difference between the maximum and minimum received intensities.

5.1 Benchmark

As mentioned in Sec. 2.2, different types of LCDs could have very different characteristics. We first empirically compare both the passive STN LCD and the active TFT LCD in terms of the response time, cross-cell interference and the transmittance. We then check the performance of different sampling strategies as finding the regression models used in our dual-lens height estimation.

5.1.1 LCD Response Time

We first measure the duration required by each LCD pixel to complete a state transition.

Experiment setup: In this experiment, a region of 10 by 10 pixels at the center of the LCD flash at a frequency of 0.5 Hz (i.e., switching once every two seconds) for a backlight operating on the frequency of 1000 Hz. This flashing rate (0.5 Hz) is configured to guarantee that each LCD pixel can completely change its state so that the photodiode can measure the maximum intensity contrast between the dark and bright states of the received signals. To find the latency required to ensure a high enough intensity contrast, the transmitter performs 30 rounds of state transitions, and the receiver measures the time gap that the *normalized intensity contrast* raises from 10% to 90% or falls from 90% to 10%. Here, the normalized contrast is the expected contrast divided by the maximum contrast.

Results: Fig. 12 plots the CDFs of the response time for both types of LCDs. The results show that the response time of the STN LCD ranges from 100 ms to 250 ms, while the response time of the TFT LCD is no more than 8 ms. Due to this, the positioning cycle may be very long if we use an STN LCD. If we force the STN LCD to change the state faster than its response time, the achievable intensity contrast of the received signals would be greatly reduced, leading to a weak signal strength. To verify this point, we further measure the intensity contrast when the duration of state switching T_{switch} is smaller than the LCD's response time, denoted by T_r, and compare it to the maximum intensity contrast. Fig. 13 plots the *normalized* intensity contrasts, i.e., the measured contrasts divided by the maximum contrast, for various settings of T_{switch}. The results show that the intensity contrast decreases significantly even if the switching duration is only reduced to half of the response time. CELLI's PPM addresses this practical challenge by reducing the time interval of each switching time-slot but leveraging pre-switching to ensure that each state can still last for a long enough duration by setting $T_{\text{switch}}\delta$ equal to the response time. Hence, our design significantly shortens the positioning cycle, while still keeping the signal strength high.

5.1.2 Interference from Neighboring Cells

We next check whether the light signal from one LCD pixel to a cell will leak to the neighboring cells and degrade the signal recovery probability in those cells.

Experiment setup: In this experiment, a receiver is placed in the origin at distance 175 cm and 225 cm, respectively, apart from the transmitter. That is, the global location of the receiver is set to $(0, 0, 175)$ and $(0, 0, 225)$, respectively. We then let each of the 10 by 10 pixels at the center of the LCD *sequentially* transmit the modulated light signals, which switch between the states 0° and 90° at 0.5 Hz. Namely, at any point of time, only a single LCD pixel is transmitting. The receiver, locating in the origin $(0, 0)$, measures the strength of an interfering signal that steers its beam to a neighboring cell (x, y) rather than $(0, 0)$. Then, the measured signal strength is refereed to as the interference from cell (x, y) at the same height, i.e., on the same projected plane.

Results: As shown in Fig. 14, the leakage interference from an STN LCD is slightly less than that from a TFT LCD. Fig. 14(b) shows that, in the case of the TFT LCD at distance 175 cm, the interfering signal strength from the cells in the left-bottom corner is stronger. This is because the resolution of the TFT LCD is higher, making it more difficult to precisely deploy the receiver in a specific cell, i.e., the origin in our setting. As a result, the measured interference map is shifted a little bit due to this small deployment error. However, both types of LCDs generate nearly no interference, i.e., -10dB or less, to a neighboring cell 2 cm away. This confirms that the directionality of light beams from multiple pixels of an LCD is the key to make the success of CELLI's fine-grained positioning.

5.1.3 Transmittance

We now check how the intensity of the illuminating light source degrades as passing through the transmitter's LCD.

Experiment setup: We configure the backlight LED to transmit a 1000 Hz carrier wave. A receiver is placed in different cells at distance 225 cm. All the pixels of the LCD set their states to 0° such that the receiver would receive the maximum intensity from the light passing through the LCD. It then compares the intensity of the received signals when the LCD is installed and is removed, respectively. The received intensity is then normalized by the maximum intensity among all the measurements, which is then converted to decibel (dB).

Results: Fig. 15(a) and Fig. 15(b) illustrate the normalized intensity for the STN LCD and TFT LCD, respectively. The x-axis of the figure is the x-axis position of the receiver with respect to the origin of the coordinate system. The results show that the trans-

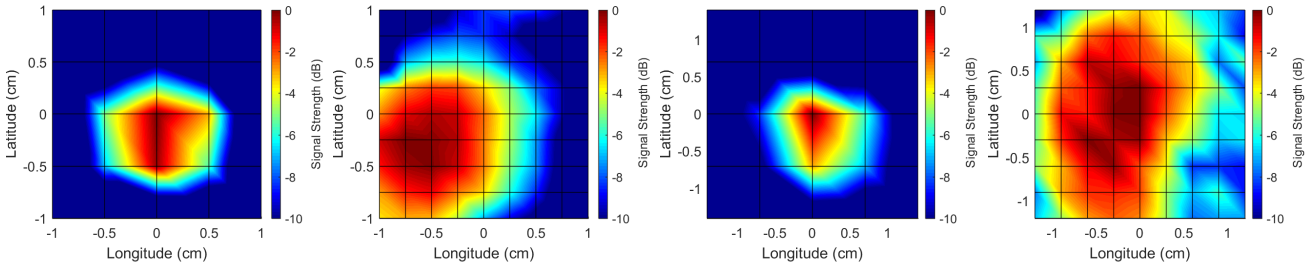

(a) STN LCD at 175cm **(b)** TFT LCD at 175cm **(c)** STN LCD at 225 cm **(d)** TFT LCD at 225 cm

Figure 14: Cross-cell leakage interference. Both types of LCDs generate nearly no interference, i.e., -10 dB or less, to a cell 2 cm away.

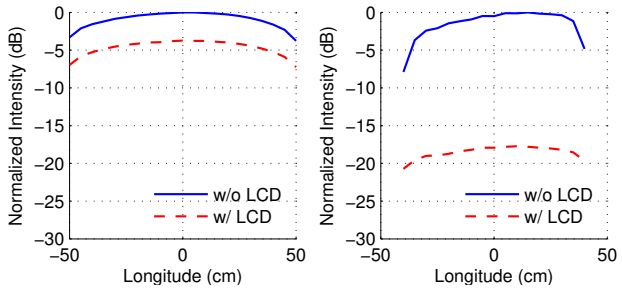

(a) STN **(b)** TFT

Figure 15: Transparency comparison. An TFT LCD locks a large amount of light. While the TFT LCD provides a better positioning accuracy, some additional light sources for illumination purpose might be required.

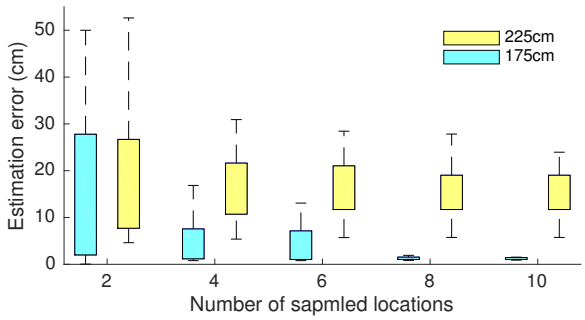

Figure 16: Fingerprinting cost. A small fingerprinting overhead allows us to train a very accurate regression model for height estimation.

mittance of the STN LCD is higher than that of the TFT LCD. By supporting CELLI's positioning service with the STN LCD, the intensity reduction is -3.7 dB, implying that the light intensity for illumination is reduced to around 43%. On the other hand, the TFT LCD introduces a much larger intensity reduction. Due to this intensity blockage effect, when the TFT LCD is used, it is likely that the transmitting light cannot be simultaneously used as illumination sources, and the environment would need some additional light sources for illumination purposes.

5.1.4 *Fingerprinting cost*

When CELLI is used only for 2D positioning, no prior training is required. For 3D positioning, as mentioned in Sec. 3.3, CELLI uses regression to learn the coefficients of the linear function converting a coordinate to the absolute position. This task needs to be performed for the transmitter only if the manufacturing error of lens configuration is not negligible, and can be done as part of the manufacturing process. To reduce the training cost, we now test how many location samples are sufficient for learning an accurate regression model.

Experiment setup: We adopt the TFT LCD in this experiment. A few positioning cells at heights 175cm and 225 cm, respectively, are sampled to train the regression model. An intuitive training method is to directly use the measured values $(x_1 - x_2)$ of randomly-selected fingerprint locations to train the regression model. However, noisy training data might degrade the accuracy of regression. Note that, as explained in Sec. 3.3, a receiver, in theory, observes the same coordinate difference $(x_1 - x_2)$ in all the cells at a given height. With this property, for each height, we instead find the median value of $(x_1 - x_2)$ of the sampled fingerprint locations to filter out the noisy data. Then, only the median values of $(x_1 - x_2)$ at the two heights are used to train the regression model. We compare the prediction accuracy as various numbers of fingerprint locations are sampled. For each setting, we repeat 50 rounds of random sampling and learn a regression model for each round of sampling. We then place the receiver in other cells to test the accuracy of the regression model.

Result: Fig. 16 shows the box plots of the height estimation errors when the number of training samples per height varies from 2 to 10. The results demonstrate that our regression-based estimation can achieve a very high accuracy as we only sample eight out of 320×240 cells to train the linear coefficient. It confirms that the fingerprint cost of learning the linear function is quite small. The results further show that the estimation error at height 225 cm is slightly larger than 175 cm. The main reason is that the dimension of each cell on the projected plane at 225 cm is larger than that at 175 cm. As a result, an error on coordinate detection translates to a larger height estimation error.

5.1.5 *Impact of PPM Cycle Duration*

We finally check how the duration of state switching T_{switch}, and thereby the positioning cycle, affects the positioning accuracy.

Experiment setup: We use the TFT LCD at the transmitter side and place the receiver in a randomly-selected cell at 225 cm. The switching duration T_{switch} is varied from 0.5 to 4 milliseconds, corresponding to a positioning cycle from 0.3 to 2.4 seconds. The pre-switching parameter δ should be updated according to T_{switch} such that the duration of the 90° state, δT_{switch}, can still be roughly the same with the LCD response time. For each setting, we repeat the experiment with different randomly-selected cells for 40 times.

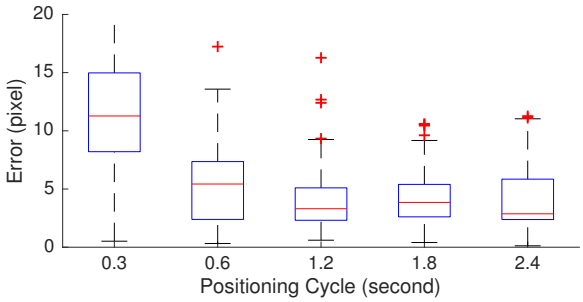

Figure 17: Impact of positioning cycle. With CELLI's pre-switching design, we reduce the positioning cycle to 0.6 s, but maintaining a comparable positioning accuracy.

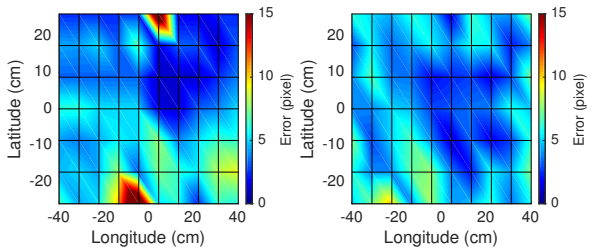

(a) Bit-pattern **(b)** PPM

Figure 18: Comparison between bit-pattern and PPM. The estimation error of bit-pattern and PPM are 4.18 pixels and 4.27 pixels, respectively, at height 225 cm.

Result: Fig. 17 shows the box plot of the coordinate estimation errors of all the measures for each tested positioning cycle. The results demonstrate that, though the transmitter ensures that the duration of the 90° state keeps constant, the positioning error still increases slightly when the positioning cycle is shorten. The main reason is that, while we are able to change the sampling rate of the transmitter, i.e.,$1/T_{\text{switch}}$, the sampling rate at the receiver, called r_{Rx}, is, however, a constant. As a result, the same error of identifying the sample index of the deep received signal will correspond to different coordinate estimation errors when T_{switch} varies. Consider an example that the true deep is the k-th received sample, but the receiver incorrectly detects the $(k-\epsilon)$-th received sample as the deep. Then, since CELLI's PPM uses the time-slot of the deep signal to represent the coordinate, the constant error time gap ϵ/r_{Rx} corresponds to more time-slots of the transmitter's state switching and, thereby a larger coordinate estimation error, when T_{switch} is smaller. However, we can see that the increase in coordinate detection errors is fairly small when we reduce the positioning cycle from 2.4 seconds to 0.6 seconds. Therefore, in our later evaluation, we set the positioning cycle to 0.6 seconds so as to achieve a balance between the positioning latency and accuracy.

5.2 Positioning Accuracy

We now conduct several experiments to evaluate the overall positioning performance of CELLI under different scenarios. The experiments are designed to answer the following questions:

- What is the positioning error in various cells at different heights, i.e., on different projected planes?
- How the types of LCDs (STN and TFT) affect the position accuracy?
- Does a receiver suffer from arbitrary rotation?

5.2.1 *Comparison between Bit-pattern and PPM*

We first compare the accuracy of the bit-pattern positioning signals and our PPM positioning signals.

Experiment Setup: We place a receiver in all the different cells at height 225 cm. Since the response time of the STN LCD is too long, a positioning cycle of the bit-pattern scheme would be more than 3 seconds, which is too long to realize practical positioning applications. We hence only evaluate the performance of the TFT LCD in this experiment. Since the bit-pattern scheme does not allow pre-switching, we cannot reduce the duration of state switching, T_{switch}, which is hence set to the inverse of the refresh rate of the TFT LCD, i.e., 100 ms. For our PPM scheme, we then reduce T_{switch} to 1 ms, but set the pre-switching parameter to $\delta = 12$. For each possible positioning cell, we collect ten measurements and output the average error.

Result: Fig. 18(a) and (b) plot the average 2D positioning errors of all the cells for bit-pattern and PPM, respectively. The median errors of the bit-pattern and PPM approaches are 1.22 cm (4.18 pixel) and 1.25 cm (4.27 pixel), respectively. Though the two schemes offer similar positioning accuracy, they, however, introduce a very different positioning latency. In particular, the positioning cycle of PPM is only 0.6 second, while the bit-pattern approach requires every pixel to take 1.8 seconds to deliver its coordinate.

5.2.2 *Impact of the LCD types*

We next check the positioning accuracy when different types of LCDs are applied in CELLI.

Experiment Setup: We place a receiver in all the different cells at height 175 cm and 225 cm, respectively, to receive the PPM signals from the transmitting light. Based on Fig. 12, the average response time of the STN LCD and the TFT LCD are 5.3 ms and 212 ms, respectively. Hence, for the STN LCD, we set T_{switch} to 5 ms and set the pre-switching parameter δ to 15. We choose these parameters to limit the entire positioning cycle within 2 seconds, which is the timing requirement of most practical positioning applications, while keeping δT_{switch} large enough to achieve an acceptable intensity contrast. As for the TFT LCD, we set $(T_{\text{switch}}, \delta) = (1\text{ ms}, 12)$, as mentioned in Sec. 5.1.5, to keep the positioning period short. Each pixel of the LCD is assigned a distinct coordinate. Therefore, the dimensions of each coordinate for the STN LCD and the TFT LCD are different, depending on their resolutions. The receiver repeats the positioning procedure 10 times in each cell and calculates the average positioning error.

Result: Fig. 19 plots the average 2D positioning errors of all the cells for the both LCDs in height 175 cm and 225 cm, respectively. The results show that both LCDs achieve a fairly accurate positioning with an median error lower than 3 cm, regardless of the positioning height. The TFT LCD achieves a higher accuracy since the TFT LCD offers a much higher resolution than the STN LCD. For the STN LCD, the errors in the top area are slightly larger due to the imperfection of the optical components. In particular, since the intensity of the light passing through the projection lens is usually not uniform, this non-uniform transmittance leads to non-uniform errors in both height configurations. The problem can be alleviated by using better optical components.

5.2.3 *Impact of Receiver Rotation*

We next check whether a receiver with rotation can still find its location.

Experiment setup: In this experiment, the transmitter uses the TFT LCD. We place a receiver in 10 different cells at height 225 cm. In each cell, the receiver rotates to different orientations, and ap-

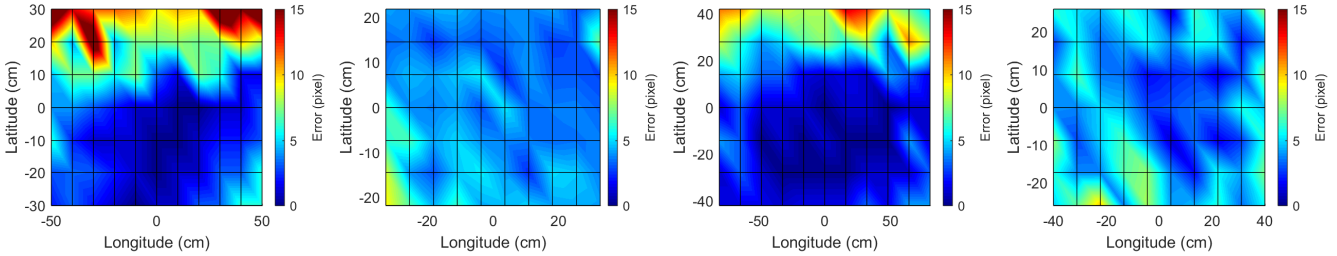

(a) STN LCD at 175cm **(b)** TFT LCD at 175cm **(c)** STN LCD at 225 cm **(d)** TFT LCD at 225 cm

Figure 19: Performance of both LCD types. Both LCDs achieve a fairly high positioning accuracy with a median error lower than 3 cm.

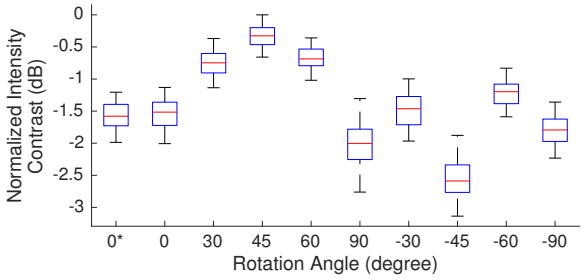

Figure 20: Impact of rotation. By dynamically adapting the receiver's polarizer direction, the achievable intensity contrast becomes similar for any receiver orientation. (The index 0^* means the case where both the orientation and the polarizer direction of a receiver is 0° and the LC adaptation is disabled.)

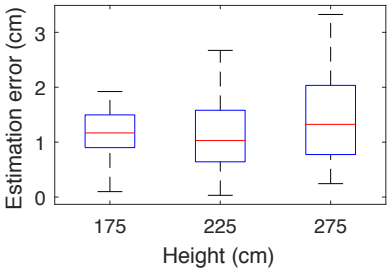

Figure 21: Impact of positioning height. The estimation errors of all heights in common indoor environment are all very small, with median error less than 1.5 cm.

plies the algorithm mentioned in Sec. 3.4 to adjust the polarization direction of its LC. We also measure the received intensity when the receiver has no rotation and does not apply the LC as baseline for comparison. The effectiveness of LC adjustment is evaluated in terms of the intensity contrast.

Result: The results in Fig. 20 show that, by testing a few sampled LC angles, the receiver can automatically configure its polarization direction to achieve a intensity contrast comparable to that of the baseline no-rotation scenario. The median reduction of the intensity contrast is no more than 2dB. A receiver can, hence, reliably recover the PPM signals and estimate its coordinate even if it rotates to an arbitrary orientation.

5.2.4 *Impact of Positioning Height*

We next check whether the positioning height affects the positioning accuracy.

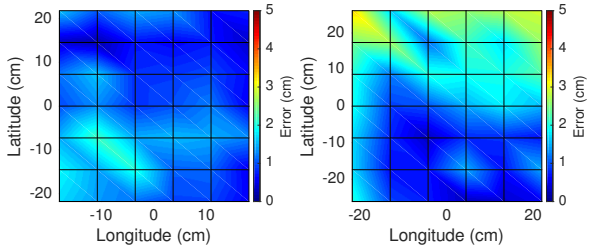

(a) height=175 cm **(b)** height=225 cm

Figure 22: 2D positioning error. CELLI achieves a median error 1.07 cm and 1.59 cm at height 175 cm and 225 cm, respectively.

Experiment setup: In this experiment, the transmitter uses the TFT LCD. Other than the heights of 175 and 225 cm, we further test a more challenging case of height 275cm, which is almost the longest distance in typical indoor environments. We place the receiver in 12 randomly selected cells at these three heights.

Result: The results in Fig. 21 show that CELLI can provide extremely accurate positioning at all common positioning distances in indoor enviroments. The median positioning errors at height 175, 225, and 275 cm are 1.17, 1.03, and 1.32 cm respectively. This is because CELLI utilizes PPM to modulate the positioning information, which could tolerate low SNR and hence can perform reliably even when the height increases and the intensity gets smaller in the 275cm case.

5.2.5 *Accuracy of 2D and 3D Positioning*

We finally check CELLI's positioning accuracy.

Experiment setup: In this experiment, the transmitter uses the TFT LCD with the pre-switching parameter $\delta = 12$ and the switching duration $T_{\text{switch}} = 1$ ms, while the receiver is deployed at heights 175 cm and 225 cm, respectively. The receiver applies the proposed LC-based polarization adaptation and the dual-lens coordinate-to-position conversion schemes to find its global location.

Result: We first demonstrate the 2D positioning accuracy in Fig. 22. The results show that the 2D positioning error is 1.6 cm on average and 3.3 cm at worst. This implies that, assuming the height of a position area is given (which is a reasonable assumption for some IoT applications, such as smart factories), CELLI can provide extremely fine-grained positioning. Fig. 23 illustrates the 3D positioning errors when the height of a receiver is unknown and is estimated based on Eq. 6. The results show that the positioning errors is higher as compared to 2D positioning. The root cause is that the x-axis and y-axis positions are estimated based on only Eq. 5. However, the z-axis coordinate P_z is estimated from both Eq. 5 and

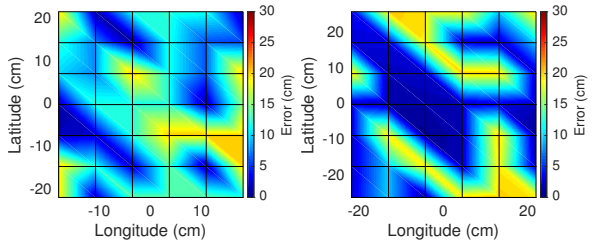

(a) height=175 cm **(b)** height=225 cm

Figure 23: 3D positioning error. Though the error along the z-dimension is slightly larger, CELLI can still provide sub-meter accuracy. The median errors at height 175 cm and 225 cm, respectively, are 11.81 and 2.65 cm respectively.

Eq. 6. Hence, the aggregated errors of the coefficients in the both equations caused by regression increase the positioning errors in the z-dimension, which have been confirmed in Fig. 16. Even with a slightly higher height estimation error, the accuracy of CELLI's overall 3D positioning is still fairly high. The median 3D positioning errors at heights 175 cm and 225 cm are 11.8 cm and 2.7 cm, respectively.

6. RELATED WORK

We categorize the existing works on indoor visible light positioning to camera-based and photodiode-based approaches.

Light-to-Camera Positioning: PIXEL [24] uses a single-pixel LCD and a disperser in front of a light source to disperse the polarized lights of different wavelength, i.e., different colors, along different light oscillation directions. By attaching the second polarizer to a receiving camera, the camera can observe different colors form different light sources and estimate the Angle-of-Arrival (AoA) of each light according to the received color. By collecting the AoA of more than three light sources, CELLI leverages conventional AoA-based positioning algorithms to identify its own location. However, the positioning error of AoA-based positioning algorithms is typically up to tens of centimeters. Some other works [12, 14, 15, 25] utilize computer vision and triangulation techniques to estimate the location of a camera based on the location of multiple lights captured in an image. The above approaches all require multiple light sources, increasing the deployment overhead and cost. Moreover, those solutions usually introduce a significant overhead on image processing, and the estimation may fail in complex environments where the light sources in the image cannot be reliably located. More importantly, the positioning resolution of the above camera-based solutions heavily rely on the resolution of a captured image. They may perform worse for low-resolution cameras or even cannot be applied on single-pixel photodiodes. Also, camera-based technologies may be suitable for human positioning, but not object positioning, in which energy consumption and computational overhead are major concerns.

Light-to-Photodiode Positioning: Some works [11, 27] migrate the idea of RF-based triangulation to optical channels. In [11], a receiver estimates its location based on the signal strength received from at least three transmitting lights operating in different frequencies. Benefiting from the independent, clear and low multipath channel characteristics of visible light, the positioning error can be much smaller than RF-based triangulation, e.g., 2 centimeters. It, however, requires a 3-LED transmitting array in each 60 cm by 60 cm cell to perform positioning. In addition, the work does not evaluate how the interference from neighboring cells affects the positioning performance. SpinLight [23] covers the infrared light source by a 3D-printed spinning lampshade with holes. By explicitly designing the patterns of printed holes, the light source can only pass through those holes and projects bright-dark patterns onto a positioning service area such that a receiving photodiode can estimate its location based on the received light intensity. However, this design can only use infrared, which is imperceptible to human eyes, but not common LED lights. Otherwise, human eyes will visually see the patterns shaded by the lampshade. In addition, this system can only allow a receiver to estimate its coordinate, instead of the global 3D location. By contrast, CELLI is a pioneer work that needs only a single light source and a single-pixel photodiode to enable 3D positioning.

7. DISCUSSION

This section briefly discusses the limitations of this work and potential directions for future research.

Range (throw-distance). Our experimental results report the positioning accuracy within typical throw distances of 175, 225, and 275 cm, respectively corresponding to ceiling to user's shoulder, chest, and floor distances in the common case of 3 meter ceiling height. Although not evaluated in this paper, in less common settings with larger ceiling height, despite lower received signal strength, we expect similar positioning accuracy due to the energy efficient PPM scheme. The output power of the backlight can also be increased to compensate for the larger range, which could also be required for illumination purposes.

Service area. Our implementation has a service area of 160 cm by 84 cm (with 225 cm throw distance). If a larger service area is desired, a projection lens with smaller focal length can be used, allowing CELLI to produce a larger projection, at the cost of decreased signal strength. An alternative approach is to use multiple CELLI transmitters such that positioning can be performed at any of their service areas. Different transmitters use different preambles and data frequencies such that a receiver can demodulate the signal at possible frequencies and locate itself correctly even if some service areas overlap.

RX and TX instrumentation. As mentioned in Sec. 4, we built our prototypes with off-the-shelf components. In other words, many components were not customized for the system and far from optimized in both size and efficiency. In real application scenarios, a transmitter could be further compressed into a light bulb size by carefully choosing the guiding lens of proper focal length, while a receiver could be implemented by only a chip composed of low-cost photodiode, a MCU, and a small battery.

Comparison to RFID. Past works have shown that RFID can be used for fairly accurate indoor positioning in both line-of-sight and non-line-of-sight cases, and its tags do not require an energy source. CELLI, on the other hand, also has a few advantages over RFID. First, multipath propagation is less prominent for optical signals, and therefore CELLI can have good accuracy in a wider range of environments. Second, CELLI utilizes the lighting infrastructure, and potentially can reduce the cost to deploy transmitters compared to deploying RFID readers.

Some future works could overcome CELLI's limitations. CELLI receiver can be modified to use a small solar panel [13] to provide its power, eliminating the need for an energy source. In addition, although the optical signal processing scheme in CELLI receiver currently runs on an off-the-shelf microprocessor, its low complexity provides opportunities for easy generation of an application-

specific integrated circuit (ASIC), enabling a low-cost, energy-efficient receiver solution comparable to passive RFID tags.

8. CONCLUSION

In this paper, we present CELLI, an accurate indoor positioning system using only a single custom light as the transmitter and a simple light sensor as the receiver, making our designs applicable even for energy-constrained small objects. By utilizing the high spatial resolution of an LCD, CELLI can emit a large number of narrow but interference-free light beams that project distinct positioning information to different cells in a service area. CELLI combines a novel pre-switching mechanism with PPM to significantly shorten the positioning cycle, and furthers leverage a dual-lens design to support 3D positioning even using only a single transmitting light. We demonstrate via testbed experiments that CELLI achieves a centimeter-level 3D positioning accuracy, and the error of 2D positioning can be 3.3 cm at worst.

9. ACKNOWLEDGEMENT

This research was supported in part by the Ministry of Science and Technology of Taiwan (MOST 106-2633-E-002-001, MOST 104-2628-E-009-014-MY2), National Taiwan University (NTU-10 6R104045), Intel Corporation, and Delta Electronics.

10. REFERENCES

[1] 4D Systems | uLCD-28PTU. http://www.4dsystems.com.au.

[2] Good Display | GDC8811. http://www.good-display.com.

[3] NHD-C12864GG-RN-GBW. http://www.newhavendisplay.com.

[4] OSRAM Opto Semiconductors | SFH 206 K | PIN Photodiodes in Through hole. http://www.osram-os.com/os.

[5] OSRAM Opto Semiconductors | SFH 2430 | SMT PIN Photodiodes . http://www.osram-os.com/os.

[6] STM32F405RG - High-performance foundation line. http://www.st.com/content/st_com/en.html.

[7] Thorlabs | PDA100A Si Switchable Gain Detector. https://www.thorlabs.com.

[8] J. Armstrong, Y. A. Sekercioglu, and A. Neild. Visible light positioning: a roadmap for international standardization. *IEEE Communications Magazine*, 51(12):68–73, December 2013.

[9] J. Hamkins. Pulse position modulation. *Handbook of Computer Networks: Key Concepts, Data Transmission, and Digital and Optical Networks, Volume 1*, pages 492–508, 2007.

[10] A. Jha. *A Textbook of Applied Physics*, volume 1. IK International Pvt Ltd, 2009.

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

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

[13] J. Li, A. Liu, G. Shen, L. Li, C. Sun, and F. Zhao. Retro-VLC: Enabling Battery-free Duplex Visible Light Communication for Mobile and IoT Applications. In *Proceedings of the 16th International Workshop on Mobile Computing Systems and Applications*, HotMobile '15, pages 21–26, New York, NY, USA, 2015.

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

[15] M. S. Rahman, M. M. Haque, and K.-D. Kim. Indoor positioning by led visible light communication and image sensors. *International Journal of Electrical and Computer Engineering*, 1(2):161, 2011.

[16] N. Rajagopal, P. Lazik, and A. Rowe. Visual light landmarks for mobile devices. In *Proceedings of the 13th international symposium on Information processing in sensor networks*. IEEE Press, 2014.

[17] T. Scheffer and J. Nehring. Supertwisted nematic (stn) liquid crystal displays. *Annual Review of Materials Science*, 27(1):555–583, 1997.

[18] L. Shangguan and K. Jamieson. The design and implementation of a mobile rfid tag sorting robot. In *ACM MobiSys*, 2016.

[19] L. Shangguan, Z. Li, Z. Yang, M. Li, and Y. Liu. Otrack: Order tracking for luggage in mobile rfid systems. In *2013 Proceedings IEEE INFOCOM*, 2013.

[20] J. Wang, F. Adib, R. Knepper, D. Katabi, and D. Rus. Rf-compass: Robot object manipulation using rfids. In *ACM MobiCom*, 2013.

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

[22] J. Wang, D. Vasisht, and D. Katabi. Rf-idraw: Virtual touch screen in the air using rf signals. In *ACM SIGCOMM*, 2014.

[23] B. Xie, G. Tan, and T. He. Spinlight: A high accuracy and robust light positioning system for indoor applications. In *ACM SenSys*, 2015.

[24] Z. Yang, Z. Wang, J. Zhang, C. Huang, and Q. Zhang. Wearables can afford: Light-weight indoor positioning with visible light. In *ACM MobiSys*, 2015.

[25] M. Yoshino, S. Haruyama, and M. Nakagawa. High-accuracy positioning system using visible LED lights and image sensor. In *IEEE Radio and Wireless Symposium*, pages 439–442, 2008.

[26] Y. Zhao, Y. Liu, and L. M. Ni. Vire: Active rfid-based localization using virtual reference elimination. In *2007 International Conference on Parallel Processing (ICPP 2007)*, 2007.

[27] Z. Zhou, M. Kavehrad, and P. Deng. Indoor positioning algorithm using light-emitting diode visible light communications. *Optical Engineering*, 51(8):085009–1–085009–6, 2012.