

Pulsar: Towards Ubiquitous Visible Light Localization

Chi Zhang

University of Wisconsin-Madison
czhang296@wisc.edu

Xinyu Zhang

University of California San Diego
xyzhang@ucsd.edu

ABSTRACT

The past decade’s research in visible light positioning (VLP) has led to technologies with high location precision. However, existing VLP systems either require specialized LEDs which hinder large-scale deployment, or need cameras which preclude continuous localization because of high power consumption and short coverage. In this paper, we propose Pulsar, which uses a compact photodiode sensor, readily fit into a mobile device, to discriminate existing ceiling lights—either fluorescent lights or conventional LEDs—based on their intrinsic optical emission features. To overcome the photodiode’s lack of spatial resolution, we design a novel sparse photogrammetry mechanism, which resolves the light source’s angle-of-arrival, and triangulates the device’s 3D location and even orientation. To facilitate ubiquitous deployment, we further develop a light registration mechanism that automatically registers the ceiling lights’ locations as landmarks on a building’s floor map. Our experiments demonstrate that Pulsar can reliably achieve decimeter precision in both controlled environment and large-scale buildings.

CCS CONCEPTS

• **Computer systems organization** → **Special purpose systems**; • **Hardware** → **Signal processing systems**; **Sensors and actuators**; **Sensor applications and deployments**; **Sensor devices and platforms**;

KEYWORDS

Visible light sensing; Indoor localization; Signal processing; Sensor; Angle of arrival; Photodiode

1 INTRODUCTION

Indoor localization represents the enabling technology behind a huge space of location-based services, including human/robot navigation, targeted advertisement, inventory control, healthcare monitoring, augmented reality, *etc.* Recent research has led to solutions that deliver decimeter level accuracy under controlled experimental settings [24, 26]. Unfortunately, none of such solutions has widely penetrated the real-world. A recent pilot study of state-of-the-art indoor localization systems [33] concluded that indoor localization

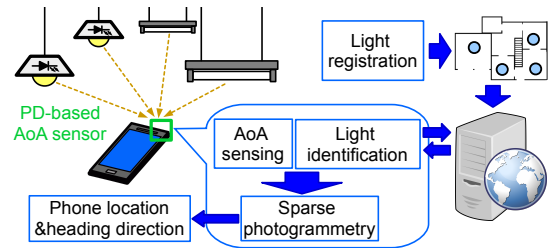


Figure 1: Pulsar’s architecture.

persists as a grand challenge, primarily because no existing solution simultaneously satisfies desired properties of *high accuracy*, *reliability* and *low cost*.

The visible light positioning (VLP) technology has shown potential to fill the gap. VLP employs “smart LEDs” as location landmarks, and photodiodes (PDs) [31, 55, 66] or cameras [26, 40, 58, 64] as sensors. The dense overhead deployment of light fixtures, and multipath-free propagation profile, together enable VLP’s high spatial resolution and resilience to environmental dynamics. Existing VLP technologies have achieved meter [31, 40] or decimeter precision [26, 58, 64], when using PD and camera, respectively.

However, two fundamental limitations have been impeding the adoption of VLP: (i) Requirement for specially designed smart LED bulbs, which modulate the light signals to create unique beacon IDs. Such bulbs are substantially more costly and bulkier than commercial LEDs or fluorescent lights (FLs) [48], which hinders near-term mass deployment. In fact, even the basic LEDs only make up around 7% of commercial/industrial lighting, and may take 10+ years towards wide penetration [49], mainly because FL fixtures have already dominated the market and retrofitting/updating cost is formidable. (ii) Dependence on regular light shape and angular response. VLP systems that employ RSS-based light propagation models commonly assume Lambertian radiation patterns [3, 31], which are applicable only to round-shaped LED bulbs (Sec. 2.1). A camera offers high spatial resolution and can unleash VLP from RSS models [26, 64]. A recent camera-based VLP system, LiTell [64], works even for conventional FLs. But the camera’s narrow field-of-view (FoV), high power consumption (~ 2 watts [64]) and high processing latency (~ 2 seconds [64]) prevent it from supporting pervasive, real-time, energy efficient location tracking. The limitation to FLs also makes LiTell a transition solution as LEDs will eventually take over.

In order to overcome the limitations of conventional VLP systems, we introduce Pulsar in this paper. Pulsar builds on recent measurement observations [64] which identified the intrinsic high-frequency flickering patterns inside FLs using cameras. Yet Pulsar targets more ambitious design goals that are conducive to ubiquitous location-based services: (i) working under both LEDs and incumbent FLs in existing buildings; (ii) using energy-efficient “single-pixel” PDs, but still able to resolve 3D location with similar

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom '17, October 16-20, 2017, Snowbird, UT, USA

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-4916-1/17/10...\$15.00

<https://doi.org/10.1145/3117811.3117821>

precision as the million-pixel cameras [64]; (iii) realizing *continuous, seamless coverage* as users roam around in practical building environment; (iv) allowing easy integration with mobile applications, and easy deployment with minimal bootstrapping overhead.

Fig. 1 illustrates the main components and workflow of Pulsar. Pulsar uses PD as its light sensor, which reduces the power consumption by several folds compared with camera-based VLP systems. In addition, owing to PDs' much higher dynamic range, Pulsar can capture LEDs' intrinsic frequencies that are much weaker than FLs', and even when the LEDs are more than 9m away (Sec. 3), in contrast to camera-based solutions [64] that only work within a range of 2.5m even for FLs. This salient property enables Pulsar to capture multiple lights simultaneously even when the lights are sparsely deployed on high ceilings, and obtain location fixes virtually without any blind spots.

Unfortunately, such benefits come with new challenges: (i) *Near-far interference*. The frequency features of faraway lights often mix with the weak side-peak artifacts of close-by lights and get masked, making it hard to distinguish collocated lights. (ii) *Lack of spatial resolution*. In contrast to the million-pixel camera, the PD is a single-pixel intensity sensor, which provides no spatial resolution and cannot easily identify its geometrical relation with different lights.

Pulsar's key solution principle lies in a novel mechanism, referred to as *sparse photogrammetry*, which resolves the light source's angle-of-arrival (AoA) using a compact light sensor. Pulsar's light sensor comprises two PDs with different FoVs¹, and hence different angular response patterns. The differential response between the two PDs follows a nonlinear function with the AoA, which can be calibrated and known at manufacturing time. At runtime, Pulsar can measure the differential response and map it back to the light source's AoA. By using AoA instead of RSS, Pulsar also circumvents the Lambertian model, enabling localization with lights of any shape. By combining the AoA of adjacent light sources, Pulsar can pinpoint the device's 3D location through a triangulation model. When more than 3 lights fall in its FoV, the Pulsar sensor can even compute its orientation angles. On the other hand, to solve the near-far problem, we observe that all frequency components of the same light have the same AoA, which allows us to single out a light's intrinsic frequency, even when it is immersed in other nearby lights' frequency components.

Full deployment of an indoor localization system requires knowing the landmarks' locations *a priori*, a problem often overlooked by existing localization systems. Pulsar uses a *light registration* mechanism to greatly simplify the issue: it only requires a one-time bootstrapping procedure, where a surveyor walks across the lights inside a building, while the lights' frequency features are captured by Pulsar, and lights' locations automatically registered on a publicly available floor map and stored in a database. Following any subsequent change of light fixtures, the database can be incrementally updated based on user devices' feedback.

We have conducted a full-fledged implementation of Pulsar. The Pulsar sensor is prototyped as a USB dongle, plugged in an Android phone that runs the localization algorithms. Owing to the use of

lightweight PDs, we expect Pulsar can readily augment new generations of smartphones, robots or low-power IoT devices that need location-based services [4, 25, 29, 36]. We have further implemented the light registration mechanism along with a simple user interface.

Our experiments verify Pulsar's high accuracy, robustness, and usability. In controlled setups similar to many existing VLP systems [26, 31, 40, 58], Pulsar can achieve 10 cm of median location precision. In practical building environments, Pulsar can easily detect 3 or more lights within its FoV, achieving a median localization error of 0.6 m and heading direction error of 4° while walking. More importantly, Pulsar is robust to typical usage dynamics, such as the device's orientation variation and ambient sunlight interference. Our prototype implementation incurs a response latency of 840 ms, ready for real-time location queries, and a low power consumption of 150 mW even without duty-cycling.

To summarize, our contributions in Pulsar include: (i) A *sparse photogrammetry* mechanism that enables single-pixel PDs to sense a light source's AoA, and then pinpoint their own locations and heading direction. (ii) A light identification framework that enables a PD to discriminate LEDs/FLs whose frequency features interfere with each other; (iii) Validation of Pulsar based on a full-fledged implementation including an automatic landmark registration scheme.

2 AUGMENTING PHOTODIODES WITH AOA SENSING

2.1 Why AoA?

Prior PD based VLP systems [3, 31] commonly adopt round LED bulbs whose angular emission pattern follows a Lambertian model, but commercial light fixtures present several key barriers to such RSS-based models: (i) The light fixtures are commonly tube/rectangular shaped, yielding different channel responses along different directions. So the RSS modeling will need the PD's location in the first place. (ii) The light fixture may have complicated reflective structures and covers which reshape the emission pattern [41]. (iii) Static and dynamic blockages such as ceiling decorators and people passing by may partially block the light, causing unpredictable deviation in RSS. In addition, due to aging and inevitable hardware deviation, different lights' emission power can vary, even among the same type of lights [5, 37, 41, 57]. All these factors render RSS propagation modeling unusable in real indoor environments.

To understand these challenges, we use our light sensor prototype (Sec. 6) to measure the RSS of light sources across different irradiance angles (defined in Fig. 2), at 1m distance. We measure two representative tube lights (with and without the diffuser cover, shown in Fig. 3), commonly used in commercial buildings. We then measure the light intensity with sensor facing towards the light, and normalize it against intensity at 0°. From the results in Fig. 4, we see that the Lambertian model deviates from the measured RSS by up to 6dB, which translates into a few meters of error when the PD is a few meters away from the light [31]. Worse still, the angular response of the same light can change with viewing angle, meaning that a precise RSS propagation model requires the location/orientation of the PD *a priori*.

In summary, for robustness and ubiquity, the VLP system should work independently from lights' emission characteristics such as angular response and intensity. However, this almost forbids the

¹Defined as the double-sided incident angle where RSS is halved in comparison to RSS at 0°. Unlike cameras, a PD can still sense light sources outside its FoV.

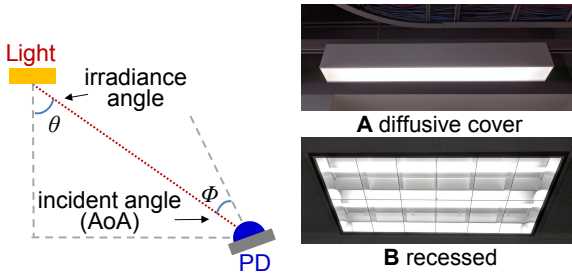


Figure 2: Channel model from light to PD.

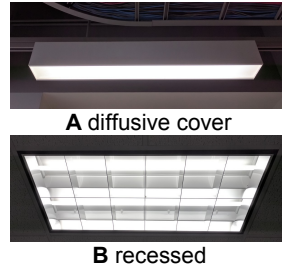


Figure 3: Two representative FLs.

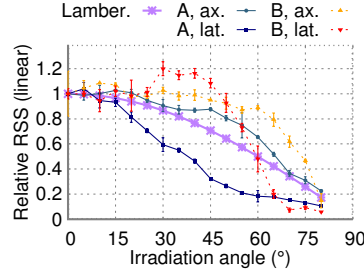


Figure 4: Axial/lateral Angular response of 2 tube fixtures, at 1m distance.

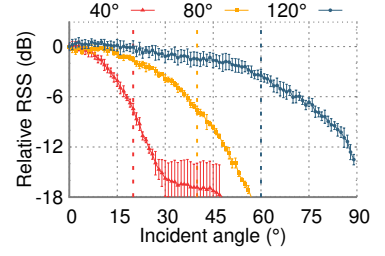


Figure 5: Angular response of 3 PDs. Dotted lines mark half-angle FoV from datasheets.

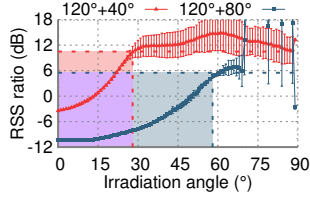


Figure 6: Composite angular response of 2 PD combinations.

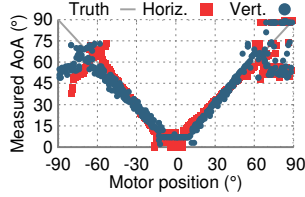


Figure 7: AoA accuracy for horizontally and vertically placed sensor.

use of RSS modeling. An AoA model can naturally avoid such problems, because it approximates the light as a point source, whose irradiation angle is unaffected by its shape, partial occlusion, and intensity variation.

2.2 Pulsar's AoA Sensing: Model and Principles

Harnessing the million-pixel sensors and high spatial resolution, a camera can easily extrapolate the physical geometries of multiple luminaries in an overhead scene image into 2-dimensional AoA information. This so-called *photogrammetry* technique [26] can arguably provide the most accurate and reliable geometrical information [3], albeit at the cost of higher power consumption, limited FoV, etc. In this section, we describe how Pulsar uses single-pixel PDs to obtain AoA while overcoming camera's inherent limitations (Sec. 1).

2.2.1 AoA Sensing Using Single-Pixel PDs. The key innovation of Pulsar lies in a *dual-PD light sensor*. Pulsar uses the differential angular responses between the 2 PDs to obtain 1-D AoA (*i.e.*, incident angle between incoming light and the PD's center axis, illustrated in Fig. 2), which will in turn feed Pulsar's sparse photogrammetry localization algorithm (to be discussed in Sec. 4).

More specifically, the RSS between each light source and the PDs follows a generic channel model [31, 63]:

$$\begin{aligned} \text{RSS}_1 &= P_t A_t(\theta_1) \alpha(r_1) A_{r1}(\phi_1) \\ \text{RSS}_2 &= P_t A_t(\theta_2) \alpha(r_2) A_{r2}(\phi_2) \end{aligned} \quad (1)$$

where P_t is the emission power of the light. $A_t(\theta)$ is the irradiance angular response at irradiance angle θ ; $\alpha(r)$ is the propagation loss at distance r , and $A_{ri}(\phi)$ is the incident angular response of the i -th PD at incident angle ϕ ². Since the 2 PDs are most likely co-located on a mobile device, with negligible separation compared with r , we can safely assume $\theta_1 = \theta_2$, $r_1 = r_2$ and $\phi_1 = \phi_2 = \phi$. By dividing

²Note that $A_{ri}(\phi) = \cos \phi$ only when the receiver follows the ideal Lambertian model and has 120° FoV.

the linear RSS observed by the 2 PDs, all other factors related to the emission/intensity characteristics of the light will cancel out:

$$\frac{\text{RSS}_1}{\text{RSS}_2} = \frac{A_{r1}(\phi)}{A_{r2}(\phi)} = A_c(\phi) \quad (2)$$

Inverting Eq. (2) and we can obtain the 1-D AoA or incident angle:

$$\phi = A_c^{-1} \left(\frac{\text{RSS}_1}{\text{RSS}_2} \right) \quad (3)$$

Note that RSS_i can be measured in real-time by each PD at each frequency. The *composite angular response* $A_c = A_{r1}/A_{r2}$ is essentially a function that maps AoA to the differential response between the 2 PDs. Each PD's response function A_{ri} is fixed and is known through data-sheet or a one-time factory calibration, which in turn allows us to derive A_c as a static lookup table.

The above model builds on three approximations: (i) For A_c^{-1} to be a valid function, A_c needs to be *monotonic*, at least within a range of AoA. (ii) The two PDs bear the same AoA relative to a light source. (iii) The light source can be approximated as a point source from the PDs' perspective in terms of area. We now validate these approximations through measurements.

2.2.2 Accuracy of AoA model. We first verify the monotonic relation between AoA and A_c , using different PD combinations. We test 3 PDs with different FoVs: 40° (Vishay BPV10), 80° (Everlight PD333) and 120° (OSRAM BPW34). Fig. 5 plots the actual angular responses of the PDs, measured by rotating the PD using a programmable motion controller [12] while illuminated by a LED. Note that the measured FoV often deviates from the nominal rating in the datasheet, which necessitates a one-time factory calibration of the angular response.

To verify the AoA model, we measure the composite angular response A_c of different narrow+wide PD combinations: BPW34 + BPV10 and BPW34 + PD333, respectively. Fig. 6 shows that A_c indeed grows monotonically with AoA up to a certain threshold, beyond which it starts to drop, making A_c^{-1} ambiguous. The threshold AoA equals 28° and 57°, translating to usable FoV of 56° and 114° for the two combinations, which is much wider than most smartphone cameras (around 70° [2]). For both combinations, the monotonic region is slightly larger than the FoV of the PD with narrower FoV. In practice, *to ensure a one-to-one mapping between AoA and A_c , Pulsar only uses A_c when it falls in the monotonic region.*

We also notice a tradeoff between FoV and spatial resolution for the dual-PD sensor. The wider-FoV combination BPW34 + PD333 has lower spatial resolution — the composite angular response A_c is flat at near-zero AoA region, which can induce ambiguities. This

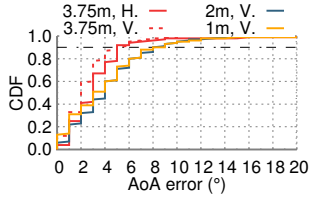


Figure 8: CDF of AoA sensing error.

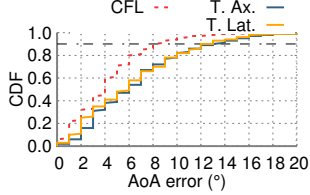


Figure 10: CDF of AoA error for tube FLs, along axial and lateral direction.

is because PDs with larger FoV have flatter angular response at lower AoA values. We expect combining 3 or more PDs can make a better tradeoff and provide even larger FoV. For simplicity, the rest of paper will use the BPW34 + PD333 combination, and leave other possibilities for future work.

2.2.3 Impact of PD Arrangement. Our AoA model assumes the two PDs are arbitrarily close and experience the same optical channel. But in practice, they have to be physically separated, inducing certain anisotropy. To evaluate the impact on AoA sensing, we place Pulsar’s dual-PD sensor 3 m below a ceiling light, and rotate it to 500 random but known directions using the motion controller. Since the PDs have a very small form-factor, they are placed approximately 5 mm apart. The scatter plot in Fig. 7 and CDF in Fig. 8 indicate that the 90-percentile errors are less than 5° within the ambiguity-free region for all cases. This implies the *dual-PD can be abstracted as a single sensor, receiving lights isotropically.*

2.2.4 Point-source Approximation. When using a PD to sense a large light fixture (e.g., tube lights in a rectangular housing) at close distance, the light spans across a large range of incident angles. The equivalent angular response will be a smoothed version of the original one, which may break the point-source approximation and introduce error in AoA.

To quantify the ultimate impact, we place a CFL lamp at various distances, leading to different levels of response smoothing. Using the same measurement methods as in Fig. 7, we found that even in the extreme case with 1m distance, the mean AoA error is only 4° (Figs. 8 and 9). And it reduces to 3° and 2° at 2m and 3.75m, respectively.

Further, we repeat the experiments with tube FLs (T8 type, Fig. 3A). The results (Figs. 10 and 11) show that when placed 2 m below the ceiling, Pulsar’s mean AoA sensing error is below 6° , and 90-th percentile below 13° . Interestingly, the error does not change with direction (Sec. 2.2.3), as the amount of smoothing in the PD’s angular response does not vary. Overall, *the point-source assumption well approximates reality.*

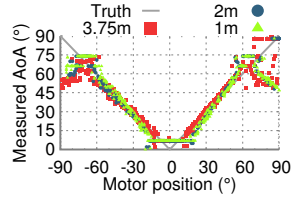


Figure 9: AoA accuracy vs light-to-sensor distance.

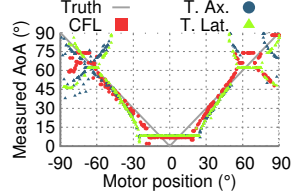


Figure 11: AoA sensing accuracy of tube FLs, along axial and lateral direction.

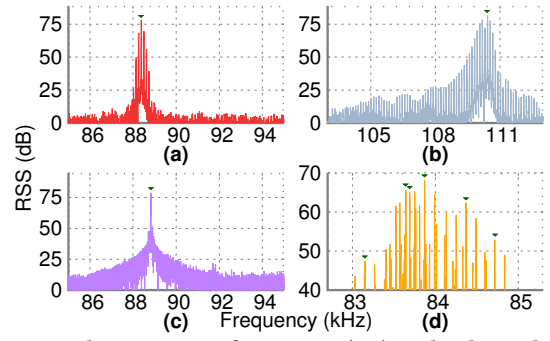


Figure 12: Characteristic frequency (CF) and side peaks for typical commercial lights: (a) T8 tube FL, (b) CFL bulb, (c) commercial LED, (d) signals from multiple tube FLs added together. Green triangles mark CFs.

3 DISCRIMINATING LIGHT LANDMARKS USING FREQUENCY FEATURES

Since 3-D localization requires 1-D AoA information from multiple landmarks, Pulsar’s sensor must be able to *simultaneously discriminate the lights and compute the AoA of each.* This requires adjacent light fixtures to emit uniquely identifiable signals that do not interfere with each other when captured by the PD. Most existing VLP systems use modified LEDs which modulate the optical signals to meet these requirements [26, 31, 58]. Meanwhile, LiTell [64] proposed a way to discriminate unmodified FLs, based on the fact that their driver works by oscillating [35, 43, 51], and manufacturing tolerance [45] leads to unique interference-free oscillating frequencies, referred to as *characteristic frequencies* (CFs).

As in LiTell, Pulsar aims to discriminate the lights using their CFs as features. When capturing multiple lights simultaneously, LiTell can isolate the lights’ signals because their locations differ on the camera image. In contrast, the use of single-pixel PDs brings two unique challenges to Pulsar. (i) *Near-far interference.* A light’s emission spectrum comprises not only the CF, but also various spurious frequencies. The spurious signals from nearby/brighter lights can overwhelm the CF features from further/dimmer ones, making them unidentifiable. (ii) *Lack of spatial resolution.* Previous work [64] utilizes relation between consecutively located lights to facilitate error correction, while Pulsar’s sensor cannot resolve spatial relation. This limits Pulsar’s light identifying accuracy when collision happens. In the following sections, we describe how Pulsar meets these challenges.

3.1 Light Feature Extraction

To elucidate the near-far interference problem, we use a PD sensor (Sec. 6) to sample and analyze the frequency-domain spectrum of representative luminaries offline, including tube FL, compact FL (CFL), and LED. From the plots in Fig. 12(a)-(c), we see that each type of light manifests a peak CF, but accompanied by numerous side peaks³. Fig. 12(d) further shows the mixed light spectrum in an office room, with tube FL fixtures separated 1.2m apart on the ceiling. The PD sensor has 120° FoV, faces upwards and sits 4m below the ceiling. The ground-truth CF is measured by turning on each light separately. We can see that the PD captures the CFs

³Diversity of CFs within each model of light has been validated by [64].

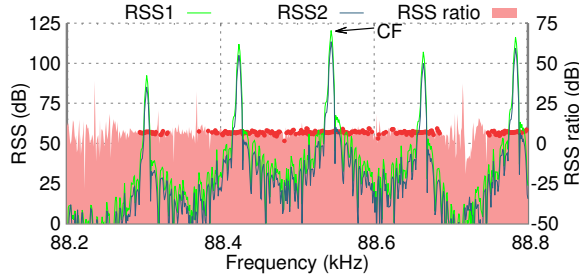


Figure 13: Artifacts in a single light’s spectrum and corresponding RSS ratio. Dots show ratio with sufficient RSS.

Algorithm 1: Algorithm for light extraction.

input : \mathcal{P} – spectrum peaks with sufficient RSS; ϵ_{AoA} – tolerance (minimum difference in AoA/RSS-ratio)
output : \mathcal{L} – CFs

sortByRSS(\mathcal{P})
 $\mathcal{L} \leftarrow \emptyset$
for $p \in \mathcal{P}$ **do**
 $\mathcal{L}.\text{add}(p)$, $\mathcal{P}.\text{remove}(p)$
 for $q \in \mathcal{P}$ **do**
 if $|q.\text{AoA} - p.\text{AoA}| < \epsilon_{\text{AoA}}$ **then**
 $\mathcal{P}.\text{remove}(q)$

of 6 lights simultaneously. However, the faraway lights’ CFs are immersed in the spurious peaks caused by close-by lights. Thus, unlike camera-based LiTell system [64] which only needs to extract the single strongest peak in the spectrum, we cannot single out every frequency by picking the strongest frequency components. In effect, there is even no easy way to identify the number of lights within the spectrum.

To overcome this challenge, note that we can use the model in Sec. 2.2 to compute the AoA of the signals corresponding to each individual frequency peak. We observe that *if different frequency peaks have the same AoA, it is likely that they originated from the same light source*. Therefore, we first sort the peaks in the spectrum by their AoAs, and then pick the strongest peak for each unique AoA, which corresponds to the true CF of one light. We repeat this until the remaining peaks become negligible (likely to be noise).

To verify the feasibility of this AoA-based light extraction, we set up one tube FL inside a darkroom. We then use our AoA sensor to receive the light signal, while moving the sensor up and down, just like walking. Fig. 13 shows that the RSS ratio is indeed consistent across different frequency components from the same light, as long as RSS’s from both PDs are sufficient.

Alg. 1 summarizes the procedure of CF extraction in Pulsar. Note that since the AoA value is derived from RSS difference from 2 PDs, a slight difference in frequency responses of the 2 amplifiers may cause slightly different AoAs to be observed at each frequency. As a result, we empirically set a tolerance to the RSS ratio (0.5dB in our experiments) when merging spectrum peaks with similar AoAs, according to experimental results for a single light in the dark room. Meanwhile, when 2 lights have similar AoAs, the weaker one will be removed. Such similar AoAs often come from co-located lights, and thus will not reduce the number of effective lights.

Occasionally, frequency components of 2 different lights may collide (fall within 2 FFT bins), which may distort the AoA. However, it will only result in spurious detection under the following additional conditions: (i) *Both lights have sufficient RSS and their perceived RSS ratio is distorted by more than 0.5dB*. Either their RSS at the colliding frequency is comparable⁴, or a different light dominates each PD. (ii) *RSS at the colliding frequency is higher than any other lights with similar perceived AoA*. In practice, we find this situation to be rare throughout our experiments.

3.2 Identifying Multiple Lights without Spatial Knowledge

Despite the diversity of the CF, occasionally there exist remotely separated luminaries with barely distinguishable CF values. LiTell [64] mitigated such CF collisions by combining two consecutively sampled lights, which provides a form of spatial error correction and greatly lowers possibility of CF feature collision. But it requires the user to travel across at least 2 lights, thus precluding instantaneous location fix without any user intervention. In contrast, Pulsar’s PD-based sensor can resolve a group of lights’ CFs simultaneously, which enhances feature diversity without user intervention, but entails a new challenge. Unlike camera [64], PD has no spatial resolution and cannot recognize whether two lights are adjacent.

To overcome this barrier, for each measurement observation of the light signals, Pulsar uses all the CFs it extracts (Sec. 3.1) as a feature collection, denoted as \mathcal{L} . It then searches for the set of lights in the database whose CFs are closest to \mathcal{L} in terms of Euclidean distance. Given the large number of lights in each building, a brute-force search among all possible group combinations is infeasible. However, we can prune the search space based on two observations: (i) *Frequency error of each light*. Existing measurement study [64] showed that the CF feature is highly diverse even among lights in the same building (99% of the lights have CFs separated more than 20Hz). It is also stable, varying only by tens of Hz across weeks. Therefore, for each element in \mathcal{L} , it suffices to put the few database candidates with closest CF into the search space. (ii) *Physical distance between each identified light*. Since within each observation the sensor should see lights which are near each other, groups that have candidates with large distance (computed with the candidates’ registered locations) are also likely to be incorrect matches. We thus combine these two factors into a score that describes how likely a match is valid. These two factors can be weighted and we empirically give them equal weight when values are in kHz and meters. In addition, we use the score as a *confidence* value for each light, which is later used for optimizing the reliability of localization (Sec. 4). Alg. 2 summarizes the above procedure which identifies the group of lights in one measurement observation.

3.3 Sensitivity and Coverage

Ideally, Pulsar needs to ensure multiple lights’ CF features are detectable to realize blind-spot free location sensing. Existing work re-purposed smartphone cameras to extrapolate the lights’ CFs [64]. Due to their narrow FoV [2] and low dynamic range, a camera can only capture a *single light* in office buildings with typical light

⁴If RSS’s of the 2 lights are within 9dB, the summed RSS would be higher than *both* lights by more than 0.5dB, and the error in RSS ratio may exceed 0.5dB.

Algorithm 2: Algorithm for group light matching.

input : \mathcal{L} – CFs; $\epsilon = 500\text{Hz}$ – maximum CF error; $N = 2$ – number of candidates; \mathcal{B} – CF database

output : \mathcal{E} – light identities

for $l \in \mathcal{L}$ **do**

$l.\text{candidates} \leftarrow \{b \in \mathcal{B} | \epsilon > |l.\text{CF} - b.\text{CF}|\}$
 $\text{sortByError}(l.\text{candidates})$
 $\text{truncate}(l.\text{candidates}, N)$

$C \leftarrow \text{candidateCombinations}(\mathcal{L})$

for $\mathcal{D} \in C$ **do**

$\text{score}(\mathcal{L}, \mathcal{D})$

$\mathcal{E} \leftarrow \{\mathcal{D} \in C | \mathcal{D} \text{ has lowest score}\}$

Procedure $\text{score}(\text{features } \mathcal{L}, \text{candidates } \mathcal{D})$

$\epsilon_d \leftarrow 0, \epsilon_f \leftarrow 0$

for $d \in \mathcal{D}$ **do**

$\epsilon_d \leftarrow \epsilon_d + \text{meanDistance}(d, \{d' \neq d \in \mathcal{L}\})$

for $l \in \mathcal{L}, d \in \mathcal{D}$ **do**

$\epsilon_f \leftarrow \epsilon_f + |l.\text{CF} - d.\text{CF}|$

$\mathcal{D}.\text{score} = (\epsilon_d + \epsilon_f) / |\mathcal{L}|$

density/height [64]. In contrast, Pulsar’s PD sensor, owing to wider FoV and higher dynamic range (typically 30dB higher than CMOS cameras [38, 61]), should easily observe more than 3 lights in the same environment (see measurement results in Sec. 3.3). We now conduct experiments to verify how close Pulsar can be to this ideal.

Impact of distance. How far need the light/PD to be separated to ensure detectable CF? We answer the question by experimenting with two PDs of different FoVs (40° and 120°), and a GE CFL with 15W power, which is conservative compared with most commercial lighting fixtures (20W - 78W [50]). We also set up a Philips 8.5W LED bulb, with the same 800-lumen output as the CFL. To isolate the impact of incident angle, we pose the PD and lights so that they face each other directly. Then we measure the signal-to-noise ratio (SNR) between the CF component and the noise floor on the received light spectrum. The results in Fig. 14 show that the CFL’s SNR decreases over distance, but remains above 26dB even at a distance of 8m. Considering the much higher emission power and high density of commercial tube FLs [50], we expect Pulsar to easily capture multiple lights’ CF features simultaneously in real-world buildings. Note that the narrower PD has higher SNR due to higher angular gain.

The LEDs’ CF feature is relatively weaker, but still 16dB above the noise floor at 8m distance. The weaker feature is due to the fact that LEDs are driven by DC and manufacturers usually employ filters on outputs of LED drivers to suppress flicker. As a contrast, we note that the camera-based CF detection system [64] only works for FL, and its SNR drops close to 0dB even at a short distance of 2.5m. In addition, we have followed the methodology in [64], and tested the CF diversity in a college building with 179 LEDs, which is the only large-scale LED deployment nearby. There are 2 different light models. The first covers 120 LEDs whose CFs roughly follow a distribution $\mathcal{N}(89.092, 0.831)$ kHz, while the CFs of the other model are scattered between 18 and 77 kHz. We found the LEDs exhibit

diverse CFs just as the FLs measured in [64], albeit at weaker SNR which is not discernible by cameras.

Joint impact of ceiling height and light density. The actual number of lights Pulsar can capture in practical indoor environments depends on not only the lights’ height/density, but also relative angle to the PD sensor. Since it is infeasible to test all buildings, we synthetically generate different height/density configurations, and use the SNR-to-distance trace in the foregoing experiment to compute the number of CFL bulbs visible to a PD. In addition to distance, the signal intensity is also scaled by the angular gains due to the light’s irradiation angle (from the bulb) and incidental angle (into the PD), assuming the PD faces upwards and each light source has roughly 20 W power. For simplicity but without loss of generality, we only simulate round-shaped CFL with known gain patterns. A bulb is considered visible to the PD if the signal corresponding to its CF exceeds an SNR of 3dB.

From the results (Fig. 16), we observe that the number of observable lights increases with higher light density (smaller spacing). It also grows with ceiling height despite the increase in propagation loss. By analyzing the data, we observe that as long as the PD sensor has $> 80^\circ$ FoV (which is quite common for commodity PDs), it can observe 3+ lights in most practical settings (spacing $< 4\text{m}$, height $> 2\text{m}$).

Impact of ambient light interference. Sunlight may affect VLP when it operates near the window or under skylights. To test whether Pulsar can work under strong ambient light, we place its sensor dongle (Sec. 6) by a west-facing window on a sunny day, 2m below a tube FL. We started measurements at 1:30 PM when the sunlight starts to shine through the window, and stop at 7:00 PM close to sunset. The sensor dongle has its TIA gain configured to a low value of $1\text{k}\Omega$ to avoid saturation, and the second stage gain is set to $300\times$. Fig. 15 shows that the sensor remains at high SNR and does not saturate even under direct sunlight (approx. 35 kLux). Since the sunlight has no flicking frequency, it does not interfere Pulsar as long as saturation does not occur. This is salient property is harnessed in Pulsar’s 2-stage amplifier design (Sec. 6) that is only sensitive to flickering light.

4 LOCALIZATION WITH SPARSE PHOTOGRAMMETRY

Once Pulsar identifies multiple lights within its sensor’s FoV (Section 3) and computes the AoA relative to each light (Section 2), it uses simple geometrical models to derive the sensor’s location. We now introduce how the models work, under common light density (3+ lights within FoV) and extremely sparse deployment.

4.1 3D Location Fix with 3+ Lights

Without loss of generality, we assume Pulsar’s dual-PD to be installed on a smartphone’s front-face, just like the current light sensor and front-camera. Denote the phone’s location as $K = (x, y, -h)$, where (x, y) is its 2D coordinate and h the vertical distance to the ceiling. The light i ’s location $L_i = (x_i, y_i, 0)$ is known from Pulsar’s light registration mechanism (Sec. 5), and stored in an online database. Further, denote the unit normal vector of the phone’s surface as $\mathbf{N} = (a, b, c)$, which can be measured by standard smartphone motion sensors (compass, gyroscope, and accelerometer) [67]. The

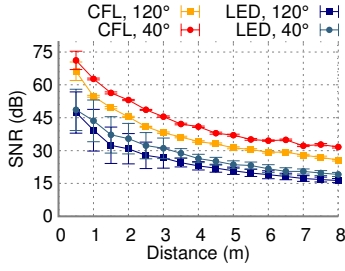


Figure 14: RSS vs distance for PDs with wide and narrow FoV.

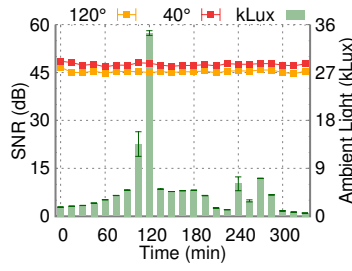


Figure 15: Impact of ambient light on SNR for PDs with narrow and wide FoV.

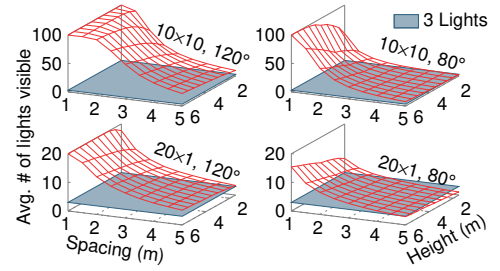


Figure 16: Number of lamps visible to the PD under different ceiling height and lamp density, with matrix (10 × 10) and array (20 × 1) configurations.

line vector from phone to light $\mathbf{PL}_i = (x_i - x, y_i - y, h)$ is parallel to the direction of incoming light beam. By definition, the AoA ϕ_i is the angle between the light beam and the normal vector of the phone's surface (Fig. 2), so we have:

$$\begin{aligned} \cos \phi_i &= \frac{\mathbf{N} \cdot \mathbf{PL}_i}{|\mathbf{N}| |\mathbf{PL}_i|} \\ &= \frac{a(x_i - x) + b(y_i - y) + ch}{\sqrt{(x_i - x)^2 + (y_i - y)^2 + h^2}} \end{aligned} \quad (4)$$

since $|\mathbf{N}| = 1$ by design.

Eq. (4) is the core equation that enables Pulsar's sparse photogrammetry localization. To resolve the 3 unknowns representing the phone's location, we need AoA measurements from at least 3 lights to form a system of equations. When $n > 3$ lights are sensed, we can turn the *equation solving* into a *minimization problem*:

$$K = \underset{h > 0}{\operatorname{argmin}} \sum_{i=0}^n w_i \left(\frac{\mathbf{N} \cdot \mathbf{PL}_i}{|\mathbf{N}| |\mathbf{PL}_i|} - \cos \phi_i \right)^2 \quad (5)$$

which essentially finds the location vector that best fits the AoA measurements to all lights. Here w_i represents the matching confidence for light i (Sec. 3.2), which acts as a weight to prevent occasionally mismatched lights from disturbing the location computation.

4.2 Circumventing Unreliable Compass Information Using 4+ Lights

Compasses can measure the phone's azimuth angle, but are notorious for its poor accuracy indoor due to magnetic interferences from metal structures [13, 42, 60]. To overcome this limitation, Pulsar can treat azimuth as a variable in the system of Eq. (4), and use 4 or more lights to solve for both 3D location and azimuth. To this end, we note that the normal vector \mathbf{N} can be rewritten as a function of azimuth α , roll β and pitch γ based on rotation matrix⁵:

$$\begin{aligned} \mathbf{N} &= (\cos \alpha \sin \gamma + \cos \gamma \sin \beta \sin \alpha, \\ &\quad -\sin \gamma \sin \alpha + \cos \gamma \cos \alpha \sin \beta, \\ &\quad \cos \gamma \cos \beta) \end{aligned} \quad (6)$$

The pitch γ and roll β angles are measured by the accelerometers relative to the direction of gravity. They have proven to be reliable in indoor environment [42, 67]. So we only solve for the azimuth variable α along with x, y, h , following the same minimization problem in Eq. (5). More generally, given 6 or more lights, we

may even resolve the γ and β without any motion sensors, achieving both 3D location and 3D orientation sensing. But this entails ultra-high density light deployment and will not be covered in our experiments.

4.3 Identifying the Nearest of 2 Lights

Pulsar can fall back to coarse-grained light-cell level localization when insufficient number (≤ 2) of lights are simultaneously visible to populate the system of Eq. (4). This fallback mode requires Pulsar to identify the light closest to the phone, which can still afford meter-scale location granularity, depending on the light density.

Far-away lights have larger irradiance angle and hence lower angular gain, as well as higher propagation loss. However, the nearest light does not necessarily yield the highest RSS. Ambiguities can be introduced by variations of lights' output power and PD's angular response, as evident from the channel model in Eq. (1). The former effect is negligible compared to the distance attenuation [57], especially because most buildings prefer homogeneous light sources. To combat the latter variation (*i.e.*, RSS uncertainty due to PD's orientation), we leverage the PD's known angular response *w.r.t.* the measured AoA ϕ_i of the light i . More specifically, we can compensate for the angular response to obtain the true RSS: $RSS' = RSS / A_{r1}(\phi_i)$. Then the light with the highest RSS' always has the shortest distance to the PD.

In Pulsar, the cell-level location is also used to seed the initial point of the nonlinear minimization problem in Eq. (5) and accelerate its convergence. It also helps identify each overhead lights in the light registration process (Sec. 5).

5 LIGHT REGISTRATION ON MAP

Ubiquitous deployment of VLP requires a scalable way of surveying the landmark locations in the first place. Prior VLP systems [3, 26, 31, 40] assume *a priori* knowledge of the light fixtures' locations, but actually, even the "smart" LEDs do not know their own locations. The lights' locations and identities have to be associated manually, either by programming the bulbs or via a separate database. Pulsar's *light registration* scheme solves this problem. Simply put, it requires a one-time survey of the lights inside a building, which automatically registers the lights' locations on the building's floor plan image, and stores the results in Pulsar's database, eliminating the need for physically measuring the location of each light. The lights' locations can be translated into physical locations given the scale of the floor map.

⁵The rotation matrix has many forms depending on definition of the axis and order of the rotation. Here we use the version that matches the Android API [19].

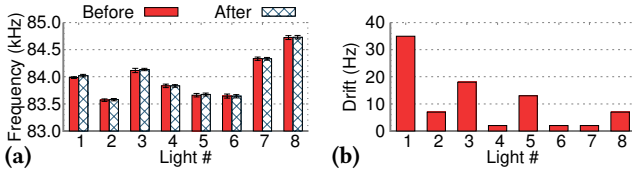


Figure 17: (a) Monthly mean and *std.* and (b) drift of CF before and after 1 year for 8 office tube FLs.

More specifically, a surveyor needs to start under a light with known location (on the map), hold a Pulsar receiver and walk across each of the ceiling lights inside the target building. Meanwhile, the surveyor’s walking trajectory is tracked using existing inertial-tracking methods. When Pulsar detects a light (manifested by an RSS peak), it records the light’s CF and marks its 2D location the same as the surveyor’s current location (ceiling height is assumed known).

Our Pulsar prototype employs Google’s Tango tablet to track the surveyor’s trajectory. Tango combines cameras and inertial sensors to achieve decimeter precision [16] with its visual-inertial odometry (VIO). Due to inevitable drift of the inertial sensors, Tango’s precision degrades over long distances. So we partition the surveying process into sections, each starting from a light with known location on map. The residual errors are small and can be compensated through manual adjustment on the map, considering the ceiling lights have highly regular separations. The compensation can also be automated through geometrical fitting and computer vision, but this is beyond the scope of our current work. We note that Pulsar is not tied to the Tango device — advances in robotics have enabled high-precision VIO using commodity smartphone cameras and motion sensors [8, 28].

We emphasize that the VIO methods work under several assumptions that are tolerable for controlled one-time surveying, but cannot substitute Pulsar’s real-time localization: certain anchor points exist whose locations are known to the user device; the surrounding environment scene is static; the walking speed is low and needs to be regular. In addition, the camera operations incur high power consumption and are unsuitable for continuous localization battery-powered devices.

Further, note that the CF features depend on the lights’ hardware and have shown to be highly stable over multiple weeks [64]. We have also conducted a year-long measurement of 8 randomly sampled lights in an office building. We found that the CFs are also stable at annual scale, with monthly and year-to-year average drift less than 10 Hz for most of the lights, as shown in Fig. 17, which is stable enough for reliable long-term identification [64]. In the longer term, a light’s CF feature may change due to failure, aging or replacement. Such changes will manifest as a consistent single-light mismatch when users run the light identification algorithm (Section 3). Therefore, Pulsar can detect such outage events, and update its light database automatically, keeping the overhead of re-registration to a minimum.

Finally, the light registration should not be confused with the fingerprinting procedure in signature-based localization systems [44, 59]. The fingerprinting entails extensive survey of all *locations*, whereas Pulsar only needs a one-time walk-through of all light *landmarks* which are much sparser.

6 IMPLEMENTATION

Dual-PD light sensor. We design a custom light sensor that can be plugged into mobile devices with USB ports (Fig. 18). The sensor contains an STM32F103C8T6 microcontroller, which samples the signals from 2 PDs (BPW34 and PD333). The MCP6024 quad operational amplifier is used to form 2 transimpedance amplifiers and 2 voltage amplifiers, which condition the signals before they enter the microcontroller. The entire prototype costs around \$30, while a production version built into the phone (Sec. 8) will be significantly cheaper.

The firmware running on the microcontroller uses DMA and timer to drive 2 internal ADCs, which sample 2 PDs simultaneously at 300Ksps with 8-bit depth. Since the ADCs are directly timed by the hardware timer, jitter and artifacts in sampling are kept to a minimum, ensuring highly accurate spectrum analysis. The sampled data are buffered in a 16 KB ring buffer before being transferred to the mobile device via 12 Mbps USB 2.0 full speed. We use a rigid USB OTG adapter so Euler angles of the device and the dongle are kept the same.

Android app and registration software. As illustrated in Fig. 19, the Pulsar Android app comprises a few modules: UI (map/location display), 3-D position solver, database, light identification, signal processing and dongle communication. The light identification first runs FFT on the samples (streamed from the dongle) to obtain the spectrum, and then simultaneously extracts multiple lights’ CFs and match their identities, following Sec. 3. The app then looks up a local database that maps each light’s CF to its location. Then, the RSS of each light’s CF feature feeds the differential channel response, which is mapped to the light’s AoA (Sec. 2). Finally, Pulsar computes the sensor’s 3D location and heading direction (Sec. 4), and presents the information in the UI.

We implement the dongle driver part in native C code with JNI to allow best real-time performance and avoid discontinuities in samples, which can introduce severe artifacts in spectrum analysis. We use a modified version of libusb [14] together with Android’s USB Host API to provide access to the USB dongle in C code without requiring root privilege. Some of the signal processing and the location solver are also implemented in C since they require functions such as nonlinear minimization from a numerical library. We pick GNU Scientific Library (GSL) [17] as our numerical library and have ported it to Android. The database and PDs’ calibration data are stored in pure text, providing the lights’ identification number, CF, and physical locations. All other components are implemented in Java with standard Android API.

The inertial tracking for the light registration process (Sec. 5) is provided by the Tango tablet [20] since it has a complete and ready-to-use visual-inertial odometry (VIO) implementation. Many other monocular VIO techniques emerged recently in the robotic navigation field [8, 16], which have shown comparable tracking precision as Tango. But the implementations are either kept in private or constrained to a specific smartphone model. We have also implemented a graphical application on PC that allows the surveyor to fit the motion trace to floor plans and save the results as the database (Fig. 20).

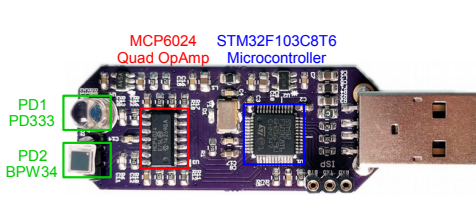


Figure 18: Photo and diagram of the customized AoA sensor.

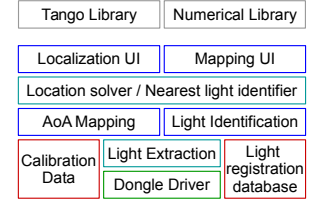
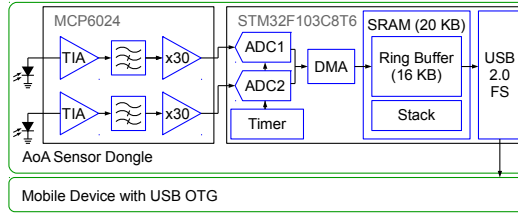


Figure 19: Architecture of the Pulsar Android app.

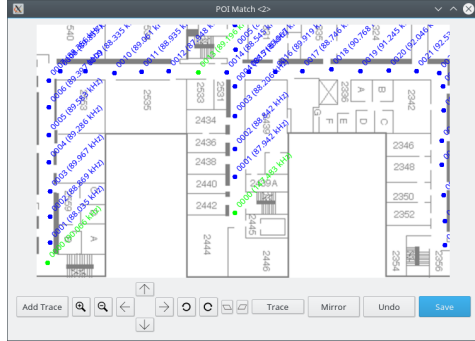


Figure 20: Light registration application running on a PC, generating the light database.



Figure 21: Photos of test venues: aisles with tube FLs (upper); an atrium with CFLs (lower).

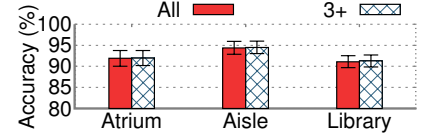


Figure 22: Light identification accuracy for 3 different types of areas, error bars show *std*.

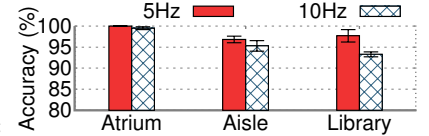


Figure 23: Light identification accuracy for smart bulbs.

Table 1: A comparison between LiTell and Pulsar.

	Ceiling height ^a	Coverage ^b	ID accuracy ^c	Battery life ^d	Latency	Light registration	Orientation	Requires tube FLs
LiTell	≈ 1m	57.1%	80.5%	2.45 hrs	1.62s ^e	Manual	Compass and light	Yes
Pulsar	> 9m	99.8%	94.4%	6.67 hrs	0.84s ^f	Automatic	Compass or light	No

^a From phone to a 15W cabinet tube FL. ^b By area in an office building aisle. ^c For single measurement with tube FLs. ^d Nexus 5, full battery cycle, screen turned off, maximum update rate. Active idling alone last ≈11 hrs. ^e 7 pictures, ≈ 5 Hz spectral resolution. ^f 100,000-point FFT, 3 Hz resolution.

7 PERFORMANCE EVALUATION

We evaluate Pulsar’s performance both through controlled experiments like existing VLP systems [26, 31], and through field tests in two environments: the aisles of an office building and the atrium of an activity center (Fig. 21). The aisles have tube FLs arranged linearly, 4.5m apart on 3m ceiling, while the atrium has 2 rings of CFLs spaced 1.5m apart on a 4m ceiling. For the light identification test, we include an additional library setting with densely deployed tube FLs on 2.8m ceiling. All computations are performed on the smartphone, with the self-contained Pulsar App. For the light registration mechanism, its landmark location precision depends on the external inertial tracking mechanism and can be manually adjusted during bootstrapping (Sec. 5). So we omit the tests for it.

7.1 Accuracy of Light Identification

We run Pulsar’s multi-light identification algorithm over databases built for the test venues with light registration. The aisle, atrium, and library involve 64, 110 and 157 lights, each with 2-3 models of lights ⁶, respectively. For each venue, we hold the sensor, move to 1000 random locations, and compare the identification results with the database ground-truth. Further, we simulate the identification for smart LEDs with randomly allocated frequencies, with frequency drifts following $N(0, 5)$ Hz and $N(0, 10)$ Hz distribution, since we could not find a building with smart-bulbs deployed in

⁶ CF range: aisle 88~93 and 137~148 kHz; atrium 81~86, 103~114 and 129~134 kHz; library 58~65 and 88~90 kHz.

large scale. We observed that the number of lights that Pulsar captures within its FoV varies from 3 to 6 in each measurement. In addition, with incumbent lights, in 92.5% of the time, Pulsar’s light identification mechanism matches every single light correctly (Fig. 22, label ‘All’). In 92.6% of the cases, Pulsar can identify at least 3 lights correctly (Fig. 22, label ‘3+’). If smart-bulbs are employed, the accuracy further rises to 93% - 100% (Fig. 23), depending on the number of lights and allocated bandwidth. The identification accuracy also shows high consistency across measurements, primarily owing to the homogeneous light layout in most buildings. Note that the wrongly matched lights in one capture are usually deemphasized by low weights (Sec. 4.1), which lessens their impact on the final localization result.

7.2 Accuracy of Sparse Photogrammetry

To receive from multiple lights, Pulsar’s sensor must be able to extract frequency and AoA information from sources with highly diverse RSS simultaneously. Therefore, the sensor must work under a wide range of RSS. Its FoV and ambiguity-free range (Sec. 2) also needs to guarantee that it can “see” enough number of lights under low ceiling, sparse deployment scenarios. In the following experiments, we set up a controlled environment similar to existing VLP systems [26, 31] to evaluate Pulsar’s performance under different settings. Specifically, we arrange 23W CFLs in a linear array, whose spacing and height (*w.r.t.* the sensor) can be flexibly adjusted. Unless

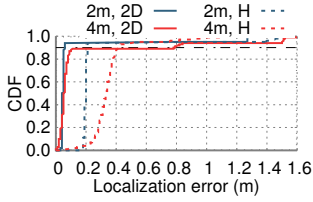


Figure 24: CDF of localization error vs. ceiling height.

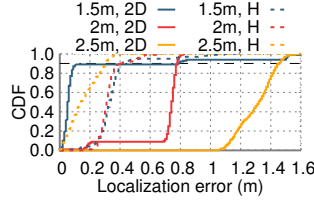


Figure 25: CDF of localization error vs. light spacing.

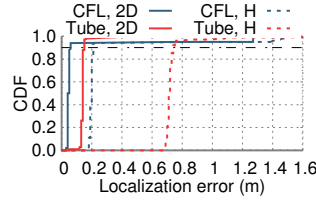


Figure 26: CDF of localization error vs. fixture type.

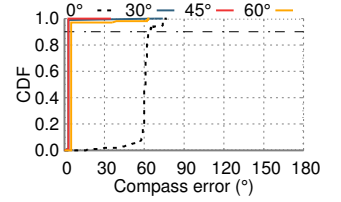


Figure 27: CDF of screen orientation error.

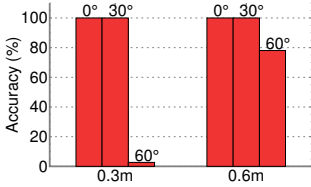


Figure 28: Accuracy of nearest light identification.

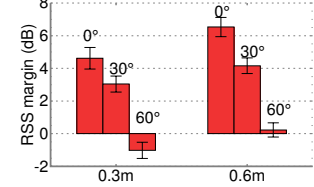


Figure 29: Margin of RSS' between near/far lights.

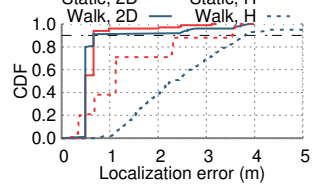


Figure 30: Localization error of static and walking user.

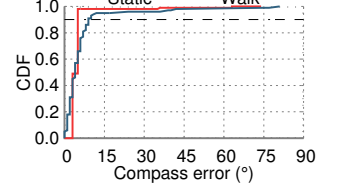


Figure 31: Orientation error of static and walking user.

noted otherwise, we repeat the experiment 500 times under each configuration.

3D localization accuracy. Fig. 24 shows Pulsar’s localization performance across different ceiling heights, with 4 CFLs placed 1.5m apart and sensor placed flat. Pulsar achieves very high precision for both 2D and height localization (label ‘H’), thanks to its AoA-based sensing scheme. The median error is 5cm for 2D localization (x, y) and 20cm for height (h), when the sensor is held 2m below the ceiling. The result is highly consistent across experiments, with 90-th percentile error close to the mean, due to the determinism and multipath-free propagation profile of visible lights. Under a high ceiling height of 4m, due to weaker RSS, the median error increases slightly, to 6cm and 31cm, for (x, y) and h , respectively. In an extreme test where the height grows to 6m, we observe a large systematic error—the estimated locations are relatively tightly packed, but deviate from ground truth for more than 1m. However, note that typical commercial lights [50] use much higher emission power than our CFLs. Therefore, we expect Pulsar’s sparse photogrammetry can provide accurate sub-light level localization under common luminaires and ceiling height settings.

Fig. 25 shows Pulsar’s localization accuracy under different separations of the 4 CFLs, while height is maintained at 4m. We notice that Pulsar’s localization error in h remains roughly the same, whereas the 2D localization precision degrades as the separation increases (median error 0.07m, 0.72m, and 1.27m, for 1.5m, 2m, and 2.5m separation, respectively). Further analysis shows the major source of error comes from the lateral direction *w.r.t.* the light array, as the RSS becomes much weaker from lights near the tail.

To verify that Pulsar behaves consistently for different types of lights, we compare Pulsar’s localization accuracy between 4 CFLs and 4 T8 tube FLs. The tube FLs are enclosed in diffusive covers, arranged with 1.2m separation in between. Pulsar’s sensor is placed 2m below the ceiling. Fig. 26 shows that Pulsar’s 2D localization performance degrades by only about 10cm when we switch from CFL to tube FL, although the area of the light is larger by at least 10 \times . This echoes the micro-benchmark in Sec. 2.2.4. It further verifies that Pulsar’s point source assumption holds in practical scenarios and does not become the bottleneck in location accuracy.

Pulsar’s 3D localization mechanism requires at least 3 lights. Ideally, adding more lights should improve accuracy. However, since the visible light channel is very deterministic and AoA sensor suffers from very little interference, we observed that increasing the number of lights does not improve the accuracy significantly across our experiments.

Accuracy of heading direction estimation. When there are 4+ lights available within FoV, Pulsar can derive the sensor’s normal vector (same as the smartphone’s), which can substitute the compass and be used to derive the user’s heading direction (Sec. 4.2). To evaluate the accuracy of heading solving, we hold the sensor in $\sim 0^\circ$ (roughly flat), 30° , 45° and 60° inclinations under one row of 0.45m-long tube FLs spaced 0.8m apart, with a sensor-to-ceiling distance of 1.5m. We take 200 samples and compare the error with ground truth obtained relative to the floor-plan. Fig. 27 shows that the 90-th percentile error is only $3^\circ \sim 5^\circ$, when the sensor is not laid flat. When the sensor is paralleling to the ceiling, however, orientation of the smartphone no longer affects AoA, which means the orientation can be resolved to arbitrary value, yielding huge errors, while 3D localization results are correct. Since users normally will not hold the smartphone completely parallel to the ceiling, we expect this to be a corner case.

Accuracy of nearest light identification. To evaluate the effectiveness of Pulsar’s nearest light identification, we place 2 CFLs 1.5m apart, and the sensor 2m below. As a stress test, we put the sensor near the center of the 2 lights’ (x, y) projection, but slightly biased to one of them by a short distance of 0.3m to 0.6m. We then gradually rotate the sensor to the other light and see how often it is mistaken for the nearest one. To evaluate the reliability in discrimination, we use the margin in the compensated RSS (RSS’ in Sec. 4.3), *i.e.*, difference between the compensated RSS of the nearest light and that of the other light. Each configuration takes 75 trials. From Figs. 28 and 29, we notice that Pulsar can identify the nearest light with close to 100% reliability when both lights fall roughly within the sensor’s ambiguity-free range (Sec. 2.2.2). Once the nearest light is far out of the ambiguity-free range, error in identification is likely to happen, especially when the sensor is close to the center of the 2 lights. Since such extreme case should remain minority, and localization error caused by wrong identification is limited since

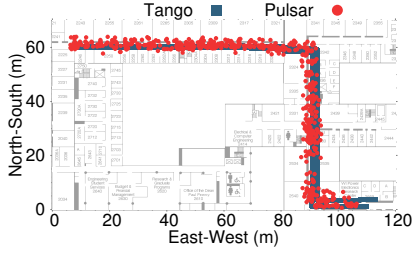


Figure 32: Localization accuracy within aisles of an office building.

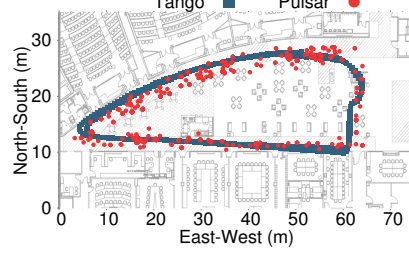


Figure 33: Localization accuracy within an atrium.

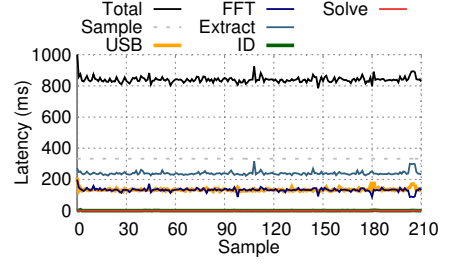


Figure 34: Pulsar's localization latency.

the receiver is close to the center between 2 lights, we conclude that Pulsar's nearest light identification is sufficient as a fallback mechanism to achieve cell-level location granularity.

Impact of human motion. We evaluate impact of user's movement with a compact setup, which includes one row of 0.45m-long tube FLs spaced 0.8m apart, placed 1.5m above the sensor. The sensor is first mounted on a tripod with 30°, 45° and 60° inclinations, and then is held by a user walking in the same spot, with device naturally rotated while walking. Fig. 30 shows that Pulsar's 2D localization accuracy does not degrade significantly under user's motion, while height estimation does show increased error. In addition, 90-th percentile error of heading estimation only increased slightly, from 5° to 9°. We thus conclude Pulsar is robust against normal user motion, as long as during each sampling period the AoA remains relatively constant.

7.3 Large-Scale, Real-time Field Tests

In this section, we verify Pulsar's localization performance in uncontrolled environment in real-time. For each test venue, we found the light registration process takes only about 10 to 20 minutes per 100 lights, since we just need to sample each light once. We walk across each venue following a marked route, and then compare the location traces generated by Pulsar and by Tango's visual-inertial odometry algorithm (shown to have sub-meter precision [16]). For a fair comparison, both run on the same Tango tablet simultaneously.

Fig. 32 and Fig. 33 depict Pulsar's performance in the aisles of the office building and the atrium, respectively, with background showing the floor-plan annotated by Pulsar's light registration application. Pulsar shows overall median and 90-th percentile accuracy of 0.6m and 1.6m, respectively. Its real-world performance is comparable to the VIO method which, despite the high precision, runs costly image processing and can only track relative movement (Sec. 5). In addition, unlike VIO, Pulsar does not suffer from sensor drift, thanks to deterministic light propagation and robust AoA sensing. Note that a few missing spots exist in the location traces, because of occasional shadowing or location solver failure. These problems can be avoided by enlarging the FoV and ambiguity-free range of the sensor (Sec. 2.2.2) and employing better minimization algorithms (Sec. 4.1) in future.

We have also measured Pulsar's processing latency within its Android app. Fig. 34 shows that the end-to-end localization latency is around 840ms across most of the location queries, which allows real-time 3D location fix. USB transfer and FFT all contribute to latency substantially, each making up 16% of total latency. Latency in extraction process is dominated by peak finding. Both the light

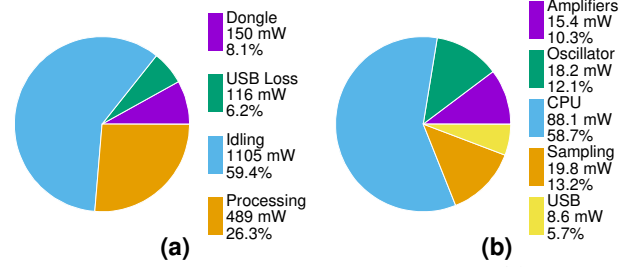


Figure 35: Power consumption breakdown for (a) the whole system and (b) the sensor dongle alone.

extraction and location solving run efficiently, costing only up to tens of milliseconds. Integrating the sensor into smartphones and eliminating the USB interface, along with sparse spectrum recovery (Sec. 8) can potentially reduce 1/3 of the latency.

7.4 Power Consumption

We measure Pulsar's end-to-end energy consumption/battery life through the sysfs interface [39], and further break the power consumption down by disabling/disconnecting each component. Fig. 35(a) shows a system-wide breakdown, where "USB Loss" is caused by converting the battery voltage up to the USB's 5V and then back down to the dongle's 3.3V; and "Idling" means active idling with adb over Wi-Fi running. Fig. 35(b) further analyzes the sensor dongle, where "Sampling" includes ADC, timing, and DMA, which copies sampled data to the dongle's buffer. The microcontroller CPU of the dongle is responsible for copying data to the USB interface, whereas it busy waits when the sampling is in progress since we have not implemented any power-management. The dongle only consumes about 45mA of current when sampling (less than 150mW since it takes 3.3V power), which is even lower than WiFi's idle listening power [65]. From an end-to-end perspective, it is clear that Pulsar's PD-based solution avoids energy-hungry camera sensors and image processing, which leads to a power consumption of less than half of camera's power end-to-end (Table 1), despite that there still exist substantial spaces for optimization (Sec. 8).

8 DISCUSSION

Integrating the AoA sensor into COTS phones. With the micro-lens technology that is widely used in CMOS cameras [15], it is possible to integrate multiple PDs onto a compact and low-cost chip, much like the ambient light and proximity sensors used on today's smartphones. These sensors usually contain at least 2 PDs [53], with the proximity sensor already operating at tens to hundreds

of kHz [23]. Unfortunately, due to hardware design and sometimes merely firmware/software limitation, fine-grained samples from these sensors are usually not available to applications. We envision future smartphones can easily integrate such AoA sensors much like the ambient light sensor, enabling more accurate localization as well as many other applications.

Further improving energy efficiency. Although the Pulsar prototype is significantly more energy efficient than camera-based counterparts [26, 64, 68], its high sampling rate (300 Ksps) provides a large space for optimization. To achieve reasonable (3 Hz) spectral resolution, Pulsar runs FFT with a long sample length (100,000 points), which wastes energy. Sparse spectrum recovery techniques such as bandpass sampling and sFFT [21] can significantly reduce computational loads, resulting in less power consumption. With such optimizations, Pulsar may be able to run entirely on energy-efficient microcontrollers.

Pulsar’s hardware prototype also has space for polishing, such as integrating the sensor into smartphones to eliminate the costly USB interfaces, and adopting better hardware components. For example, a newer version of the microcontroller used [46] consumes less than 1/3 of the power while providing more computational power. With proper hardware implementations, Pulsar’s sensor can be extremely low power and support continuously localization for battery-operated devices.

9 RELATED WORK

Existing localization schemes broadly fall into two categories: (i) Fingerprinting approaches that associate each location spot to a set of features provided by images [8, 16, 18], acoustics [47], magnetic field [11], RF signals [9, 62], or a mix of them [6]. Real-world tests showed that such features lack stability in the presence of human activities and environmental changes, and are often not discriminative enough to offer fine location granularity [32]. More crucially, they require blanket fingerprinting of all location spots, which entails huge bootstrapping and maintenance overhead. (ii) Model-driven approaches that calculate the distance between user devices and infrastructure landmarks, through models of received signal strength (RSS) [7, 10], phase [24, 56] or propagation time [34]. To achieve high location precision, such models often require specialized hardware like antenna arrays; they tend to be disturbed by multipath reflections from human bodies and ambient objects [1, 32].

Motion sensors (accelerometer, gyroscope, and magnetometer) can track a user’s relative movement via step counting [30], but need to be calibrated by other approaches that provide absolute location fixes [52]. State-of-the-art robotic systems often integrate such sensors with a camera to realize visual-inertial odometry, which can simultaneously map the environment while tracking user movement [16]. But the performance suffers in homogeneous environment (e.g., office hallways) and dynamic environment (e.g., retail stores) [8]. The use of cameras entails heavy sensing and computational power, hence unsuitable for long-term mobile usage.

Compared with these approaches, VLP has several inherent advantages, especially in its dense overhead anchoring points and line-of-sight propagation profile that evade interferences from ambient environment. Most existing VLP systems employ *modulated*

LEDs as anchor points to emit beacons. The receiver can either be a PD or a camera. The key challenge for PD based VLP lies in low spatial resolution. Existing work either uses multi-light trilateration or requires controlled user movement to resolve the spatial ambiguities [31]. Trilateration schemes assume LEDs’ Lambertian radiation pattern, which no longer holds when collimating or diffusing optics are used for uniform illumination [54]. On the other hand, controlled device movement disturbs user experience and precludes continuous, real-time localization operations. Camera-based VLP systems [26, 27] can achieve decimeter precision through *photogrammetry*, which maps the geometrical relation between LED luminaries on a camera image to the camera’s physical location/orientation [26]. Other systems [55, 58] use physical rotation or electronic polarization to create spatial diversity and simplify the sensor-side workload. All these modulated VLP systems require customized LEDs with a beaconing circuitry, which adds to the already high retrofitting cost of LEDs compared with FLs.

Our work was inspired by LiTell [64], which obviates the need for specialized LEDs. Yet LiTell only works for FLs whose CF features can be extrapolated from camera images. It is not applicable to LEDs, which have much weaker CFs, or light fixtures with small areas that will not leave enough samples in the image. More critically, cameras’ low dynamic range limits LiTell’s working range to around 2.5 meters, and their narrow FoV causes many blind spots where the ceiling lights are not directly visible. All these factors, along with other limitations (Table. 1), hinder continuous localization services. IDyLL [57] harnesses the regular separation of ceiling lights to calibrate motion sensors, thus improving dead-reckoning accuracy. But it needs users to walk in straight lines, and can only track relative movement. Fiatlux [5] extracts location fingerprints from minor intensity differences among incumbent lights, yet the fingerprints vary over time and are sensitive to ambient light and the sensor’s orientation.

Despite more than a decade of research [22], VLP has not reached real-world users. By overcoming the high cost [26, 31, 40, 58], high power consumption [26, 64], and intermittent coverage [26, 64] in traditional VLP technologies, Pulsar can significantly accelerate the penetration of VLP.

10 CONCLUSION

To our knowledge, Pulsar represents the first VLP system that employs incumbent FLs/LEDs and lightweight PDs to achieve continuous 3D localization with sub-meter precision. It marks a new step in enabling the ubiquitous deployment of VLP, by solving three problems: discriminating co-located LEDs/FLs, sensing AoA and hence 3D location using single-pixel PDs, automatically registering light landmarks. We believe Pulsar’s salient advantages will enable a wide range of location-based services, including robotic navigation, physical analytics and augmented reality.

ACKNOWLEDGEMENT

We would like to thank the anonymous reviewers and our shepherd Xia Zhou for their valuable comments and feedback. This project was partially supported by the NSF under Grant CNS-1343363, CNS-1350039, CNS-1506657, CNS-1518728, and CNS-1617321.

REFERENCES

- [1] S. Adler, S. Schmitt, and M. Kyas. 2014. Path Loss and Multipath Effects in a Real World Indoor Localization Scenario. In *11th Workshop on Positioning, Navigation and Communication (WPNC)*.
- [2] AnandTech. 2013. Understanding Camera Optics & Smartphone Camera Trends. <http://www.anandtech.com/show/6777/>. (2013).
- [3] J. Armstrong, Y. A. Sekercioglu, and A. Neild. 2013. Visible light positioning: a roadmap for international standardization. *IEEE Communications Magazine* 51, 12 (2013), 68–73.
- [4] Luigi Atzori, Antonio Iera, and Giacomo Morabito. 2010. The internet of things: A survey. *Computer networks* 54, 15 (2010), 2787–2805.
- [5] Martin Azizyan, Ionut Constandache, and Romit Roy Choudhury. 2007. Fiatlux: Fingerprinting Rooms Using Light Intensity. In *Proc. of Pervasive*.
- [6] Martin Azizyan, Ionut Constandache, and Romit Roy Choudhury. 2009. SurroundSense: Mobile Phone Localization via Ambience Fingerprinting. In *Proc. of ACM MobiCom*.
- [7] P. Bahl and V.N. Padmanabhan. 2000. RADAR: an In-Building RF-Based User Location and Tracking System. In *Proc. of IEEE INFOCOM*, Vol. 2.
- [8] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian D. Reid, and John J. Leonard. 2016. Simultaneous Localization And Mapping: Present, Future, and the Robust-Perception Age. *CoRR* abs/1606.05830 (2016).
- [9] Yin Chen, Dimitrios Lymberopoulos, Jie Liu, and Bodhi Priyantha. 2012. FM-based Indoor Localization. In *Proc. of ACM MobiSys*.
- [10] Krishna Chintalapudi, Anand Padmanabha Iyer, and Venkata N. Padmanabhan. 2010. Indoor Localization Without the Pain. In *Proc. of ACM MobiCom*.
- [11] Jaewoo Chung, Matt Donahoe, Chris Schmandt, Ig-Jae Kim, Pedram Razavai, and Micaela Wiseman. 2011. Indoor Location Sensing Using Geo-magnetism. In *Proc. of ACM MobiSys*.
- [12] Cinetics. 2017. Axis360 Pro. <https://cinetics.com/kit/axis360/>. (2017).
- [13] Ionut Constandache, Xuan Bao, Martin Azizyan, and Romit Roy Choudhury. 2010. Did You See Bob?: Human Localization Using Mobile Phones. In *Proc. of ACM MobiCom*.
- [14] Kuldeep Singh Dhaka. 2017. libusb (modified version). <https://github.com/kuldeepdhaka/libusb/tree/android-open2>. (2017).
- [15] Abbas El Gamal and Helmy Eltoukhy. 2005. CMOS image sensors. *IEEE Circuits and Devices Magazine* 21, 3 (2005), 6–20.
- [16] Christian Forster, Luca Carlone, Frank Dellaert, and Davide Scaramuzza. 2016. On-Manifold Preintegration for Real-Time Visual-Inertial Odometry. *IEEE Transactions on Robotics* (2016).
- [17] Free Software Foundation, Inc. 2017. GNU Scientific Library. <https://www.gnu.org/software/gsl/>. (2017).
- [18] R. Gao, Y. Tian, F. Ye, G. Luo, K. Bian, Y. Wang, T. Wang, and X. Li. 2016. Sextant: Towards Ubiquitous Indoor Localization Service by Photo-Taking of the Environment. *IEEE Transactions on Mobile Computing* 15, 2 (2016).
- [19] Google. 2017. Android Sensor Manager. <https://developer.android.com/reference/android/hardware/SensorManager.html>. (2017).
- [20] Google Inc. 2017. Tango. <https://get.google.com/tango/>. (2017).
- [21] Haitham Hassanieh, Piotr Indyk, Dina Katabi, and Eric Price. 2012. Simple and practical algorithm for sparse Fourier transform. In *Proc. of ACM-SIAM SODA*. 1183–1194.
- [22] Shuhei Horikawa, Toshihiko Komine, Shinichiro Haruyama, and Masao Nakagawa. 2004. Pervasive Visible Light Positioning System using White LED Lighting. *Technical report of IEICE. DSP* 103, 719 (2004).
- [23] Intersil Americas LLC. 2014. Digital Ambient Light Sensor and Proximity Sensor with Interrupt Function. <http://www.intersil.com/content/dam/Intersil/documents/isl2/isl29011.pdf>. (2014).
- [24] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. 2015. SpotFi: Decimeter Level Localization Using WiFi. In *Proc. of ACM SIGCOMM*.
- [25] Tomáš Krajník, Vojtěch Vonásek, Daniel Fišer, and Jan Faigl. 2011. AR-drone as a platform for robotic research and education. In *International Conference on Research and Education in Robotics*. 172–186.
- [26] Ye-Sheng Kuo, Pat Pannuto, Ko-Jen Hsiao, and Prabal Dutta. 2014. Luxapose: Indoor Positioning with Mobile Phones and Visible Light. In *Proc. of ACM MobiCom*.
- [27] Hui-Yu Lee, Hao-Min Lin, Yu-Lin Wei, Hsin-I Wu, Hsin-Mu Tsai, and Kate Ching-Ju Lin. 2015. RollingLight: Enabling Line-of-Sight Light-to-Camera Communications. In *Proc. of ACM MobiSys*.
- [28] Lenovo. 2017. Lenovo Phab 2 Pro. <http://www3.lenovo.com/in/en/tablets/android-tablets/tablet-phab-series/Lenovo-Phab-2-Pro/p/WMD00000220>. (2017).
- [29] J. J. Leonard and H. F. Durrant-Whyte. 1991. Mobile robot localization by tracking geometric beacons. *IEEE Transactions on Robotics and Automation* 7, 3 (1991), 376–382.
- [30] Fan Li, Chunshui Zhao, Guanzhong Ding, Jian Gong, Chenxing Liu, and Feng Zhao. 2012. A Reliable and Accurate Indoor Localization Method Using Phone Inertial Sensors. In *Proc. of ACM Conference on Ubiquitous Computing (UbiComp)*.
- [31] Lique Li, Pan Hu, Chunyi Peng, Guobin Shen, and Feng Zhao. 2014. Epsilon: A Visible Light Based Positioning System. In *Proc. of USENIX NSDI*.
- [32] Hongbo Liu, Yu Gan, Jie Yang, Simon Sidhom, Yan Wang, Yingying Chen, and Fan Ye. 2012. Push the Limit of WiFi Based Localization for Smartphones. In *Proc. of ACM MobiCom*.
- [33] Dimitrios Lymberopoulos, Jie Liu, Xue Yang, Romit Roy Choudhury, Vlado Handziski, and Souvik Sen. 2015. A Realistic Evaluation and Comparison of Indoor Location Technologies: Experiences and Lessons Learned. In *Proc. of ACM/IEEE IPSN*.
- [34] Andreas Marceletti, Maurizio Rea, Domenico Giustiniano, Vincent Lenders, and Aymen Fakhreddine. 2014. Filtering Noisy 802.11 Time-of-Flight Ranging Measurements. In *ACM International on Conference on Emerging Networking Experiments and Technologies (CoNEXT)*.
- [35] Maxim Integrated Products, Inc. 2004. Why Drive White LEDs with Constant Current? <https://www.maximintegrated.com/en/app-notes/index.mvp/id/3256>. (2004).
- [36] Daniele Miorandi, Sabrina Sicari, Francesco De Pellegrini, and Imrich Chlamtac. 2012. Internet of things: Vision, applications and research challenges. *Ad Hoc Networks* 10, 7 (2012), 1497–1516.
- [37] N. Narendran and Y. Gu. 2005. Life of LED-based white light sources. *Journal of Display Technology* 1, 1 (2005), 167–171.
- [38] OSI Optoelectronics. 2013. Basic PIN Photodiode Characteristics. <http://www.osioptoelectronics.com/technology-corner/frequently-asked-questions/basic-pin-photodiode-characteristics.aspx#06>. (2013).
- [39] Kevin Parker. 2011. Getting the battery current values for the Android Phone. <https://stackoverflow.com/a/7031032/6520528>. (2011).
- [40] Niranjini Rajagopal, Patrick Lazik, and Anthony Rowe. 2014. Visual Light Landmarks for Mobile Devices. In *Proc. of ACM/IEEE IPSN*.
- [41] Julian Randall, Oliver Amft, Jürgen Bohn, and Martin Burri. 2007. LuxTrace: Indoor Positioning Using Building Illumination. *Personal Ubiquitous Computing* 11, 6 (2007).
- [42] Nirupam Roy, He Wang, and Romit Roy Choudhury. 2014. I Am a Smartphone and I Can Tell My User's Walking Direction. In *Proc. of ACM MobiSys*.
- [43] Reid Iwao Sasaki. 1994. *The impact of electronic ballast compact fluorescent lighting on power distribution systems*. Purdue University School of Electrical Engineering.
- [44] Souvik Sen, Božidar Radunovic, Romit Roy Choudhury, and Tom Minka. 2012. You Are Facing the Mona Lisa: Spot Localization Using PHY Layer Information. In *Proc. of ACM MobiSys*.
- [45] Robert Spence and Randeep Singh Soin. 1988. *Tolerance design of electronic circuits*. World Scientific.
- [46] STMicroelectronics. 2017. STM32L433CB. <http://www.st.com/en/microcontrollers/stm32l433cb.html>. (2017).
- [47] Stephen P. Tarzia, Peter A. Dinda, Robert P. Dick, and Gokhan Memik. 2011. Indoor Localization Without Infrastructure Using the Acoustic Background Spectrum. In *Proc. of ACM MobiSys*.
- [48] The Wall Street Journal. 2015. More Companies Are Outfitting Warehouses With Smart Lights. (2015).
- [49] U.S. Department of Energy. 2014. Energy Savings Forecast of Solid-State Lighting in General Illumination Applications. (2014).
- [50] U.S. Department of Energy. 2015. Adoption of Light-Emitting Diodes in Common Lighting Applications. (2015).
- [51] H. van der Broeck, G. Sauerlander, and M. Wendt. 2007. Power driver topologies and control schemes for LEDs. In *IEEE APEC*.
- [52] He Wang, Souvik Sen, Ahmed Elgohary, Moustafa Farid, Moustafa Youssef, and Romit Roy Choudhury. 2012. No Need to War-drive: Unsupervised Indoor Localization. In *Proc. of ACM MobiSys*.
- [53] Elliott Wen, Winston Seah, Bryan Ng, Xuefeng Liu, and Jiannong Cao. 2016. UbiTouch: ubiquitous smartphone touchpads using built-in proximity and ambient light sensors. In *Proc. of ACM UbiComp*.
- [54] A. J. W. Whang, Y. Y. Chen, and Y. T. Teng. 2009. Designing Uniform Illumination Systems by Surface-Tailored Lens and Configurations of LED Arrays. *Journal of Display Technology* 5, 3 (2009).
- [55] Bo Xie, Guang Tan, and Tian He. 2015. SpinLight: A High Accuracy and Robust Light Positioning System for Indoor Applications. In *Proc. of ACM SenSys*.
- [56] Jie Xiong and Kyle Jamieson. 2013. ArrayTrack: A Fine-grained Indoor Location System. In *Proc. of USENIX NSDI*.
- [57] Qiang Xu, Rong Zheng, and Steve Hranilovic. 2015. IDyLL: Indoor Localization Using Inertial and Light Sensors on Smartphones. In *Proc. of ACM UbiComp*.
- [58] Zhice Yang, Zeyu Wang, Jiansong Zhang, Chenyu Huang, and Qian Zhang. 2015. Wearables Can Afford: Light-weight Indoor Positioning with Visible Light. In *Proc. of ACM MobiSys*.
- [59] Zheng Yang, Chenshu Wu, and Yunhao Liu. 2012. Locating in Fingerprint Space: Wireless Indoor Localization with Little Human Intervention. In *Proc. of ACM MobiCom*.
- [60] Zheng Yang, Chenshu Wu, Zimu Zhou, Xinglin Zhang, Xu Wang, and Yunhao Liu. 2015. Mobility Increases Localizability: A Survey on Wireless Indoor Localization Using Inertial Sensors. *ACM Comput. Surv.* 47, 3 (2015), 54:1–54:34.

- [61] Timothy York and Raj Jain. 2011. Fundamentals of Image Sensor Performance. <http://www.cse.wustl.edu/~jain/cse567-11/ftp/imgsens/index.html>. (2011).
- [62] Moustafa Youssef and Ashok Agrawala. 2005. The Horus WLAN Location Determination System. In *Proc. of ACM MobiSys*.
- [63] Chi Zhang, Joshua Tabor, Jialiang Zhang, and Xinyu Zhang. 2015. Extending Mobile Interaction Through Near-Field Visible Light Sensing. In *Proc. of ACM MobiCom*.
- [64] Chi Zhang, and Xinyu Zhang. 2016. LiTell: Robust Indoor Localization Using Unmodified Light Fixtures. In *Proc. of ACM MobiCom*.
- [65] Chi Zhang, Xinyu Zhang, and Ranveer Chandra. 2015. Energy Efficient WiFi Display. In *Proc. of ACM MobiSys*.
- [66] Jialiang Zhang, Chi Zhang, Xinyu Zhang, and Suman Banerjee. 2016. Towards a Visible Light Network Architecture for Continuous Communication and Localization. In *Proc. of ACM VLCS*.
- [67] Pengfei Zhou, Mo Li, and Guobin Shen. 2014. Use it free: Instantly knowing your phone attitude. In *Proc. of ACM MobiCom*.
- [68] Shilin Zhu and Xinyu Zhang. 2017. Enabling High-Precision Visible Light Localization in Today's Buildings. In *Proc. of ACM MobiSys*.