Pulsar: Towards Ubiquitous Visible Light Localization

Chi Zhang
University of Wisconsin-Madison
czhang296@wisc.edu

Xinyu Zhang
University of California San Diego
xiz368@eng.ucsd.edu

ABSTRACT

The past decade's research in visible light positioning (VLP) has achieved centimeter location precision. However, existing VLP systems either require specialized LEDs which hinder large-scale deployment, or cameras which preclude continuous localization due to power consumption and short coverage. We propose Pulsar, which uses a compact photodiode sensor to discriminate existing ceiling lights 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's angle-of-arrival, and triangulates the device's 3D location and orientation. To facilitate ubiquitous deployment, we further develop a light registration mechanism that automatically registers ceiling lights' locations on a building's floor map. Our experiments demonstrate that Pulsar can reliably achieve decimeter precision with continuous coverage.

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 is the enabling technology behind a huge space of location-based services. Yet a recent study of state-of-the-art indoor localization systems [4] concluded that indoor localization persists as a grand challenge, as no existing solution simultaneously achieve *high accuracy, reliability* and *low cost*. Visible light positioning (VLP) has shown potential to fill the gap. VLP employs "smart LEDs" as landmarks, and photodiodes (PDs) [3, 6, 10] or cameras [2, 7, 9] as sensors. Dense overhead deployment of lights and multipath-free propagation enable VLP's high spatial resolution and resilience to environmental dynamics. Existing VLP technologies have achieved meter [3] or decimeter precision [2, 7, 9] with PD and camera, respectively.

S3'17, October 9–13, 2017, Snowbird, UT, USA

ACM ISBN 978-1-4503-5145-4/17/10…$15.00
https://doi.org/10.1145/3131348.3131357

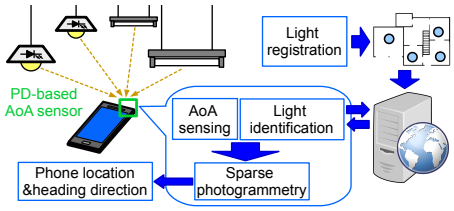

Figure 1: Pulsar's architecture.

However, two fundamental limitations are impeding the adoption of VLP: *(i)* Requirement for specially designed LED bulbs, which modulate the light signals to create unique beacons, hinders mass deployment. *(ii)* Dependence on regular light shape and angular response, as VLP systems utilizing received-signal-strength (RSS) models assume Lambertian emission [1, 3], which only applies to round-shaped light fixtures. Cameras unleashes VLP from RSS models with spatial resolution [2, 9], but its field-of-view (FoV), power consumption and processing latency preclude pervasive and continuous location tracking.

Pulsar overcomes these limitations by building on recent measurement observations [9], which captures flickering of fluorescent lights (FLs) at their characteristics frequencies (CFs), using cameras. Yet Pulsar targets ubiquitous location-based services: *(i)* working under both LEDs and FLs in existing buildings; *(ii)* achieving similar precision as cameras but with energy-efficient PDs; *(iii)* realizing continuous, seamless coverage; *(iv)* allowing easy integration and deployment.

Fig. 1 illustrates the main components of Pulsar. Pulsar uses PD as its light sensor, which reduces the power consumption by several folds compared to cameras. Owing to PDs' dynamic range, Pulsar can capture LEDs' CFs, which are much weaker than FLs', from more than 9m away, while camera-based solutions only work within 2.5m even for FLs [9]. This allows Pulsar to capture multiple lights simultaneously and obtain location fixes without blind spots.

Unfortunately, such benefits come with new challenges: *(i) Lack of spatial resolution.* Unlike cameras, PDs cannot provide spatial resolution or geometrical relations. *(ii) Near-far interference.* CFs of faraway lights are often masked by artifacts of closer ones, making it hard to distinguish them.

Pulsar's key solution lies in its *sparse photogrammetry*, which resolves the light's angle-of-arrival (AoA) using a compact light sensor. The sensor comprises two PDs with different FoVs[1], and hence different angular response patterns. The differential response between the two PDs follows a function with the AoA, which is fixed once manufactured. At runtime, Pulsar measures the differential RSS and map it back to the light's AoA. By using AoA instead of RSS, Pulsar circumvents the Lambertian model and works with arbitrary light fixtures. By combining the AoA of adjacent lights, Pulsar can pinpoint the device's 3D location and even heading

[1]Defined as the double-sided incident angle where RSS is halved.

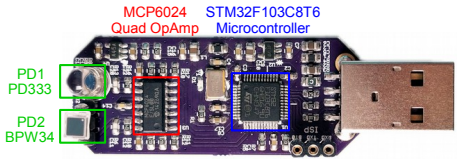

Figure 2: The Pulsar sensor prototype.

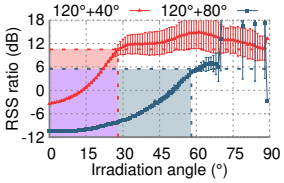

Figure 3: Composite angular response A_c. Shadowed area show ambiguity-free AoA range.

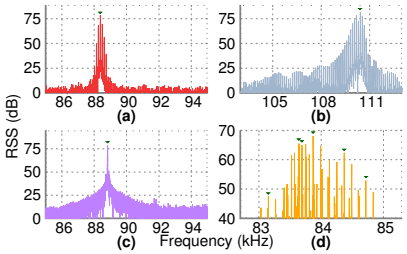

Figure 4: CF and side peaks from: (a) tube FL, (b) CFL, (c) LED, (d) multiple tube FLs. Triangles mark CFs.

direction through triangulation. To solve the near-far problem, we observe that all frequency components of the same light have the same AoA, which allows us to remove artifacts and extract a light's CF even when immersed in other lights' frequency components.

In addition, Pulsar includes a *light registration* mechanism to greatly simplify deployment: it only requires a surveyor to walks across the lights inside a building, while the lights' CFs are captured by Pulsar, and their locations automatically registered on a floor map stored in a database.

The Pulsar sensor is prototyped as a USB dongle (Fig. 2), plugged in an Android phone that runs the localization app. We expect Pulsar can readily augment new generations of smartphones, robots or low-power IoT devices. We have further implemented the light registration with a simple user interface. In controlled setups similar to many existing VLP systems [2, 3, 7], Pulsar can achieve 10cm of median location precision. While walking in practical building environments, Pulsar produces a median localization error of 0.6m and heading direction error of 4°. Our prototype incurs a response latency of 840ms, and a low power consumption of 150mW.

2 SYSTEM DESIGN

2.1 AoA Sensing

Cameras can convert geometries of multiple luminaries in an image into 2D AoA. This so-called *photogrammetry* technique [2] can arguably provide the most accurate and reliable VLP [1], albeit with higher power consumption and limited FoV. To overcome camera's limitations, Pulsar uses PDs to obtain AoA. The key innovation of Pulsar lies in a *dual-PD light sensor*. Pulsar uses the differential angular responses between the 2 PDs to obtain 1D AoA, which feeds Pulsar's triangulation algorithm.

Specifically, the RSS of each light received by the PDs follows the channel model [3]:

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

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. Since the 2 PDs are co-located 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 the linear RSS observed by the 2 PDs, all other factors related to the light will cancel out:*

$$\frac{\text{RSS}_1}{\text{RSS}_2} = \frac{A_{r1}(\phi)}{A_{r2}(\phi)} = A_c(\phi) \qquad \phi = A_c^{-1}\left(\frac{\text{RSS}_1}{\text{RSS}_2}\right) \tag{2}$$

The *composite angular response* A_c (Fig. 3) is essentially a function mapping AoA to the RSS ratio between the 2 PDs. Each PD's response A_{ri} is fixed, which allows us to derive A_c as a static lookup table.

Note that A_c is monotonic only within a certain range, beyond which the AoA cannot be uniquely determined by RSS ratio. Fortunately, Pulsar can detect such ambiguity by comparing the observed RSS ratio with the maximum monotonic value (dashed lines in Fig. 3, 56° and 114° for wide-FoV and narrow-FoV combinations, respectively). We also notice a tradeoff between FoV and spatial resolution for the sensor. The wide-FoV combination (120° + 80°) has lower spatial resolution than the narrow-FoV (120° + 40°) one, as the composite angular response A_c is flat at near-zero AoA region. Combining 3 or more PDs may provide a better tradeoff and provide even larger FoV.

2.2 Discriminating Light Landmarks

As in LiTell [9], Pulsar discriminates lights by their CFs. When capturing multiple lights simultaneously, LiTell can isolate the lights as their locations differ on the camera image. In contrast, the use of PDs brings two unique challenges to Pulsar, near-far interference and lack of spatial resolution. Here we describe how Pulsar meets these challenges.

CF extraction. To elucidate the near-far interference problem, we use a PD sensor to analyze the frequency-domain spectrum of representative luminaries. Fig. 4(a)-(c) show that each light manifests a peak CF, but with numerous side peaks. Fig. 4(d) further shows the mixed light spectrum in an office room. We can see that the faraway lights' CFs are immersed in the spurious peaks of close-by lights. Unlike camera-based system which only needs to extract the single strongest peak in the spectrum, we cannot extract multiple CFs by picking the strongest frequency components.

To overcome this challenge, note that we can compute the AoA of 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.* Thus, we 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).

Identifying multiple lights. Despite the diversity of the CF, occasionally there exist luminaries with barely distinguishable CF values. LiTell [9] mitigated such collisions by combining two *adjacent* lights, which greatly lowers possibility of the collision. But it requires the user to travel across lights, which precludes instantaneous location fix. Pulsar's PD-based sensor can resolve a group

of lights' CFs simultaneously, which enhances feature diversity without user intervention. However, PD has no spatial resolution and cannot recognize whether two lights are adjacent.

To overcome this barrier, for each observation, Pulsar uses all the CFs it extracts 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, a brute-force search among all possible groups is infeasible. However, we can prune the search space based on two observations: *(i) Frequency error of each light.* Existing measurement study [9] showed that the CF feature is highly diverse across a whole building. 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 in their registered locations are likely to be incorrect matches. We thus combine these two factors into a score that describes how likely a match is valid.

2.3 Localization

Once Pulsar identifies multiple lights and computes the AoA of each, it uses simple triangulation algorithm to derive its user's location.

Specifically, AoA ϕ_i of the i-th light L_i and the vector from the sensor $P = (x, y, z)$ to the light form an equation:

$$\cos\phi_i = (\mathbf{N} \cdot \mathbf{PL_i})/(|\mathbf{N}||\mathbf{PL_i}|) \quad (3)$$

where $\mathbf{N}$ is the normal vector of the sensor, measured by accelerometers and a compass *w.r.t.* gravity. Location of lights $L_i(x_i, y_i, z_i)$ are available in the database. Thus, with 3 lights, 3 such equations can be used to solve the 3 unknowns in the sensor's location. With 4 lights, even the heading direction can be solved, eliminating the need for the error-prone compass.

Occasionally, Pulsar may not be able to extract 3 lights simultaneously. Fortunately, it can gracefully degrade to cell-level accuracy with only one light in its FoV. When there are 2 observable lights, the AoA enables Pulsar to compensate for the PD's angular response, which in turn allows it to empirically identify the closer light.

2.4 Light Registration on Map

Ubiquitous deployment of VLP requires scalable landmark registration. Prior VLP systems assume *a priori* knowledge of the light fixtures' locations, but even the "smart" LEDs do not know their own locations. The lights' locations and identities are associated manually, either by programming the bulbs or via a separate database. Pulsar's *light registration* scheme solves this problem by automatically registers the lights' locations on the building's floor map, 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 map.

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 lights inside the target building. Meanwhile, the surveyor's walking trajectory is tracked using existing motion tracking methods. When Pulsar detects a light by an RSS peak, it records the light's CF and marks its 3D location the same as the surveyor's current 2D 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 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. The residual errors are small and can be compensated through manual adjustment on the map, considering the regular spacing of ceiling lights. The compensation can also be automated by computer vision in the future.

Note that the light registration should not be confused with the fingerprinting procedure in signature-based localization systems [5, 8]. 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.

3 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

This project was partially supported by the NSF under Grant CNS-1343363, CNS-1350039, CNS-1506657, CNS-1518728, and CNS-1617321.

REFERENCES

[1] 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.
[2] 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.*
[3] Liqun Li, Pan Hu, Chunyi Peng, Guobin Shen, and Feng Zhao. 2014. Epsilon: A Visible Light Based Positioning System. In *Proc. of USENIX NSDI.*
[4] 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.*
[5] 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.*
[6] Bo Xie, Guang Tan, and Tian He. 2015. SpinLight: A High Accuracy and Robust Light Positioning System for Indoor Applications. In *Prof. of ACM SenSys.*
[7] 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.*
[8] Zheng Yang, Chenshu Wu, and Yunhao Liu. 2012. Locating in Fingerprint Space: Wireless Indoor Localization with Little Human Intervention. In *Proc. of ACM MobiCom.*
[9] Chi Zhang. and Xinyu Zhang. 2016. LiTell: Robust Indoor Localization Using Unmodified Light Fixtures. In *Proc. of ACM MobiCom.*
[10] 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.*