Yuming Hu, Feng Qian *University of Minnesota, Minneapolis, MN, USA*
Zhimeng Yin *City University of Hong Kong, Hong Kong* **Zhenhua Li** *Tsinghua University, Beijing, China*
Zhe Ji, Yeqiang Han, Qiang Xu *XYZ10 Technology* **Wei Jiang** *State Grid Corporation of China*

LOCALIZE YOURSELF IN MALLS:
An Anatomy of a Commercial Localization System with One-million Users

Excerpted from "Experience: Practical indoor localization for malls," from *MobiCom '22: Proceedings of the 28th Annual International Conference on Mobile Computing And Networking Conference* with permission. https://dl.acm.org/doi/abs/10.1145/3495243.3517021 ©ACM 2022

The past two decades have witnessed the surge of smartphones and mobile applications. When going out, people will naturally open the map application (such as Google Maps) to find their current location (i.e., localization) and the way to their destination (i.e., navigation). At the same time, many people rely on their smartphones for localization when they walk into shopping malls, especially in large shopping centers. In outdoor environments, we usually leverage GPS to estimate our location. However, satellite signals are oftentimes weak inside buildings, resulting in large positioning errors and thus a poor navigation experience for users.

Therefore, many efforts have been put into achieving accurate indoor localization. A Google Scholar search of "indoor localization" yields more than 596,000 results. State-of-the-art solutions can even achieve sub-meter positioning accuracy [1] in the laboratory. In sharp contrast, large-scale deployment of indoor localization systems is lagging far behind in the real world. There are very limited reports on commercial indoor localization deployments [2,3], among which few offer deep insights into their deployment experiences.

In this work, we aim to fill the above gap by reporting our experiences of developing, deploying, and evaluating *MLoc*, a smartphone-based localization system for indoor

Photo, istockphoto.com

malls (commercial complex buildings) typically with tens or hundreds of retail stores. MLoc helps customers find paths to stores (e.g., the nearest Starbucks) by providing accurate, easy-to-use localization and store-level navigation. Since its debut in 2018, MLoc has been used by more than 1 million customers in China.

For more than one year, we conducted extensive evaluations at 35 malls in 7 cities in China, covering 152km^2 localization areas. Our findings show that MLoc yields a median location tracking error of 2.4m (10-th and 90-th percentile: 0.8m and 7.3m). MLoc is also a promising marketing platform: it can distribute targeted advertisements based on customers' realtime location. Through a sales event co-organized by MLoc and a large mall, we observed an ad conversion rate of 22%, significantly higher than those of online advertising. More technical details of MLoc can be found in our recently published MobiCom paper [4], which was the recipient of the best community award in ACM MobiCom 2022.

To benefit the research community, we published the whole 43GB dataset, containing the fingerprints and the localization ground truth. The data can be found at https://mloc.umn.edu/.

CHALLENGES & DECISIONS

Developing MLoc is very different from building an indoor localization prototype in the lab. It faces unique challenges, involves additional constraints, and requires judicious decisions considering numerous technical and non-technical factors, as elaborated below.

First, an important decision is to select the appropriate physical signals as localization fingerprints. Even if WiFi APs are ubiquitously deployed in today's malls, we reject using WiFi signals as location fingerprints. Android devices usually have very low WiFi scanning frequency (e.g., every 20s), iOS devices do not even offer public APIs for querying WiFi APs' RSSI, and commercial WiFi APs may periodically change MAC addresses for security considerations. Instead, MLoc uses the Bluetooth Low Energy (BLE) RSSI and geomagnetic field (GMF) strength as the location fingerprints. BLE requires a light infrastructure consisting of cheap, small, battery-powered beacons, whereas GMF is infrastructure-free. We find their synergy can lead to an accuracy adequate for store-level navigation.

Second, MLoc adopts a landmark-based outsourcing approach (i.e., hiring human collectors to survey a few predefined landmarks) to collect BLE/GMF fingerprints and the ground truth location data. However, we note that the hired collectors are quite distinct from the knowledgeable collectors in academic research – they can easily miss certain landmarks, but meanwhile would like to move existing landmarks or even suggest adding new ones. To this end, we enhance the common approach by strategically restricting the landmark visiting paths; in addition, we respect collectors' on-site opinions by allowing them to improve the predefined landmarks (calculated based on imperfect floor plan information).

Third, we choose not to build MLoc from scratch given the rich literature. Our main challenge thus lies in how to pick the suitable building blocks from existing works. We find that many sophisticated algorithms in the literature aim at dealing with challenging cases in various indoor environments. In our domain of in-mall localization, surprisingly perhaps, we observe only several generic challenging areas (e.g., atrium, corridor dead ends, corridor connectors, and elevators) despite the malls' complex layouts, based on extensive field studies at 35 malls with different scales. This observation greatly simplifies our algorithm design. Encouragingly, we find that classical algorithms can be enhanced by simple yet strategic customizations such as fingerprint preprocessing, weight adjustment, and lightweight AI to tackle the challenging areas.

OFFLINE TRAINING

The MLoc system consists of two phases: *offline training*, where pairs (fingerprint, location) are collected to build a localization model, and *online inference*, where a user's smartphone collects fingerprints, uploads them to the edge, and obtains the location and/or navigation guidance in real time. This section focuses on the offline training component, and online inference will be discussed in the next section.

Hardware Deployment

Obtaining BLE signatures requires installing BLE beacons. Due to aesthetic considerations, installation of additional power and networking cables is forbidden in shopping malls. Also, according to our agreements with the malls, we are only allowed to deploy the beacons in shared areas (as opposed to the gross leasable areas) in each mall. Therefore, as shown in Figure 1, we mount small-sized, battery-powered beacons on the ceiling of the corridor or on the surrounding edges in atrium (open space) areas. The typical distance between two beacons is from 10 to 15m, depending on the specific layout of the area. Compared to corridors, BLE fingerprints in atrium areas are more likely to cause confusion. Thus, we reduce the inter-beacon distance in the challenging areas to 6m to ensure good localization accuracy.

FIGURE 1. BLE beacons are deployed on the ceiling.

Data Collection

MLoc adopts an outsourcing approach (i.e., hiring human workers) for collecting BLE/GMF fingerprints and the ground truth location data. The entire mall's localization areas are divided into equal-sized grids (1m × 1m for BLE and 0.1m × 0.1m for GMF). It is labor-intensive to collect the fingerprints of all the grids in a mall. Thus, to reduce the data collection overhead, we developed a custom mobile app (Figure 2), using the floor plan to automatically generate the vertices of the stores' bounding boxes as landmarks. The collector needs to visit each landmark and collect the corresponding fingerprint. Then, the fingerprint of each grid is calculated by performing linear interpolation of nearby landmarks. After that, we obtain a grid map of fingerprints that will be used for online localization and tracking.

We observed two major issues in our pilot deployment using the above design. First, many collectors do not visit the landmarks efficiently; they frequently miss landmarks and must return, increasing the walking distance. Second, the floor plans are not detailed enough, making some automatically generated landmarks not visually recognizable. To tackle these issues, the data collection app generates suggested paths for the collector based on the landmarks. It uses a strategically selected zigzag path for each corridor and a circular path for each atrium. The app also allows collectors to modify the landmarks through its UI. In this way, more visually recognizable landmarks could be identified by the collectors and incorporated into future collections.

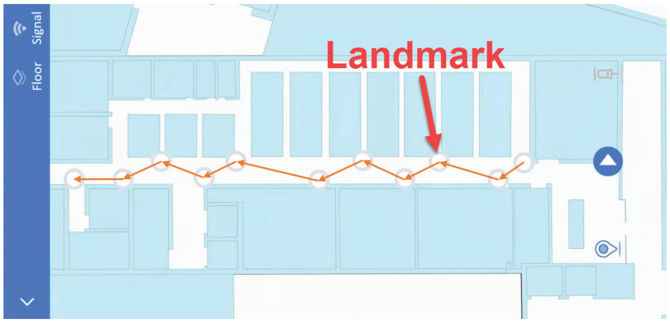

FIGURE 2. Data collection app.

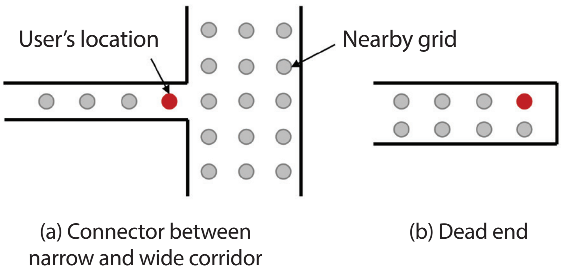

FIGURE 3. Two types of challenging areas identified in malls.

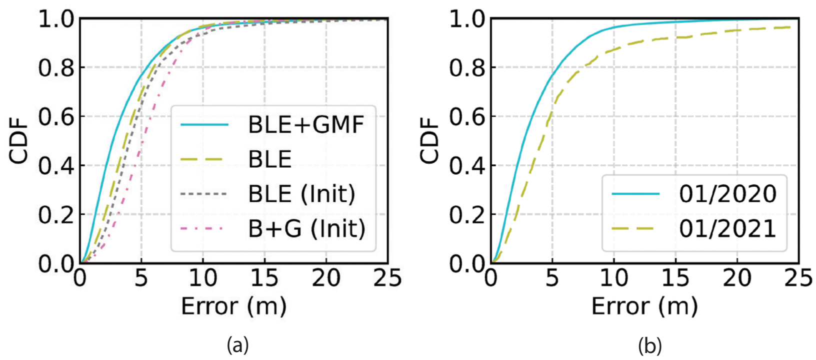

FIGURE 4. Localization and tracking accuracy.

THE DATA USED IN THIS PAPER IS FULLY OPEN SOURCED AT https://mloc.umn.edu/

ONLINE INFERENCE

During the online inference phase, MLoc computes the user's location by comparing the fingerprints scanned by their smartphone with the dataset we collected in offline training. Specifically, online inference involves two stages: initial positioning and location tracking. First, MLoc finds a user's initial location by adopting the k-nearest neighbors (KNN) algorithm, which takes the weighted average of k most likely locations as the estimated position. Second, MLoc leverages particle filtering (PF [5]) to track and refine the target's location as they are walking. PF randomly samples points on the floor map as candidates of the target's location and keeps narrowing down the candidate locations as the target moves until the estimation converges. We leverage KNN and PF to locate our target because they are lightweight, robust, and easy to implement. We demonstrate that with careful adaptations as detailed below, these traditional algorithms can lead to satisfactory localization accuracy in malls.

Fingerprint Preprocessing. Note that the initial positioning stage only uses BLE fingerprints because GMF is much noisier. We exclude the BLE beacons whose signal strengths are weaker than a threshold (empirically set to -95dB), because weaker signals are less sensitive to distance changes and introduce more noises. In addition, we observe that the fingerprint readings across different smartphone brands/models often exhibit disparities (i.e., Model A's RSSI reading is always slightly higher than Model B). To overcome this issue, MLoc adopts a simple yet effective method: it normalizes the BLE fingerprints by subtracting from each RSSI reading the average RSSI across all the samples of all the beacons collected by the same device. Given that a device can sense a large number of beacons with diverse RSSI readings, the mean value of these readings provides a good per-device "baseline". The normalization is applied to both the training and testing fingerprints.

Handling Challenging Areas. According to our field experiments, there exist two types of challenging areas. The first is illustrated in Figure 3a where the user is near the connection point between a narrow corridor and a wide corridor (or an open area). The second case is shown in Figure 3b where the user is at the dead end of a corridor. In this case, the weighted sum of the nearest neighbors will be shifted to the open end. Note that these two types are generic and representative, observed in almost all the malls we have studied. Both cases in Figure 3 are attributed to the non-uniform distribution of the grids imposed by the mall's layout. We thus modify the weight of each grid based on the floor layout information. More details can be found in our MobiCom paper [4].

Floor Detection. By default, MLoc adopts a simple floor detection algorithm, which performs a majority vote of the floors associated with the 5-strongest BLE beacon signals captured over a 5-second window. This straightforward approach works well in most cases, yielding a median accuracy of more than 97%. However, in the atrium areas, MLoc suffers from large floor errors (39%), severely influencing the users' experience. We employ a simple Deep Neural Network (DNN) [6] with 12 layers to boost the accuracy to 96%. Note that we only use DNN when necessary (i.e., in atrium areas).

OPERATION & EVALUATION

MLoc adopts the edge computing paradigm. We deploy a centralized server as the gateway to our localization service, and each mall has its own edge server, provided by us or the mall itself. When the user launches the app, the gateway server identifies which mall the user is at; then, the remaining localization and tracking tasks are handed over to the local edge server. All the computation tasks are performed on the edge. The encrypted client-edge communication is over the Internet (in-mall WiFi or cellular).

From November 2019 to January 2021, we conducted extensive evaluations at 35 malls in 7 cities in China: Hangzhou, Shanghai, Wuxi, Wuhan, Guangzhou, Tianjin, and Shenyang (up to 2200 km apart). The evaluation dataset consists of 4.3K paths, with a total walking distance of 215 km. We find that MLoc yields a median location tracking error of 2.4m (detailed results are shown in the next paragraph).

Overall Accuracy

The positioning accuracy is plotted in Figure 4a, in which the four curves correspond to the errors in the {initial positioning, tracking} stage using {BLE only, BLE+GMF}. As expected, initial positioning gives a low accuracy (median error 4.1m for BLE and 5.0m for BLE+GMF), which is significantly improved in the tracking stage (median error 2.4m for BLE+GMF and 3.5m for BLE). Note that in daily operation, MLoc uses BLE for initial positioning and BLE+GMF for location tracking. Meanwhile, we observe that for 3% of the landmarks, the location tracking error is higher than 10m. This is caused by various factors such as failed BLE beacons, fingerprint noises, erroneous floor detection, and low smartphone scanning frequency. In addition, positioning accuracy is affected by many factors, such as beacon broadcast interval, beacon spacing, and smartphone brand. Some smartphone vendors throttle the BLE scanning frequency to save energy. This may severely impact the performance of localization/tracking applications.

Temporal Stability

Once deployed, MLoc may experience a change of fingerprints. Therefore, we conducted a separate long-term experiment in three malls in Hangzhou to assess the temporal stability of fingerprints. The training data was collected in December 2019. Since then, no maintenance was performed. Then we launch two test campaigns involving the same paths, one in January 2020, and the other in January 2021. As shown in Figure 4b, after one year, the median localization error increases from 2.4m to 4.1m. This is attributed to two factors: (1) the change in the physical environment such as store renovations and various events held in atrium areas, and (2) the failure of BLE beacons, including hardware issues and the falling of beacons from ceilings. After one year of usage without replacing failed beacons (5%) or updating the training data, the error remains at an acceptable level. In the long run, our maintenance overheads mostly come from replacing fallen beacons due to glue failure, and periodically replacing the beacon battery. The maintenance can be handled by shopping malls' management teams.

MLoc as a Marketing Platform

We also showcase that MLoc can offer promising business opportunities. From November to December 2020, we co-organized a sales event with a large mall in Wuxi in Eastern China. For each participating customer, when their location estimated by MLoc is close to a store, the mall pushes the corresponding e-coupons to the customer. This event has achieved high engagement, obtaining 11K navigation sessions from 7K customers. Among them, 73% of the navigation sessions from the 3K customers involve at least one coupon retrieval. Among the customers who have retrieved coupons, 22% of them actually used the coupon(s) in stores. Such a conversion rate is significantly higher than those of traditional online advertising [7].

LESSONS AND CONCLUDING REMARKS

Our experiences of developing, deploying, and evaluating MLoc suggest that high localization accuracy is only one of the multiple objectives of MLoc, which also needs to carefully balance the tradeoffs between accuracy, human labor, infrastructure complexity, usability, and maintenance overhead, to name a few. Our key lessons learned include the following.

- In complex malls, it is feasible to survey a small number of landmarks and use them to generate a fingerprint grid map. To reduce the data collection overhead, not only the landmarks but also their visiting paths should be pre-generated.
- Despite the complex floor layouts, there are only several generic types of challenging areas (e.g., atrium, corridor dead ends, corridor connectors, and elevators) based on our extensive field studies. They can be tackled by classical algorithms, such as KNN and PF, enhanced by simple yet strategic customizations (e.g., weight adjustment and lightweight AI).
- BLE and GMF are complementary. BLE is accurate, having low resolutions, and slow to scan; GMF is noisy, having high resolutions, and quick to collect.
- MLoc can be used as a promising marketing platform that can distribute targeted advertisements based on customers' real-time location.

The data used in this paper is fully open sourced at https://mloc.umn.edu/. We hope our insights and dataset can boost future efforts on transforming the two-decade research on indoor localization into commercial products. ■

Yuming Hu is a PhD student in the Computer Science and Engineering Department at University of Minnesota. His research interests focus on network protocols, indoor localization, and VR/AR systems.

Feng Qian is an associate professor in the Computer Science and Engineering Department at University of Minnesota. His research interests cover the broad areas of intelligent mobile systems, 5G/6G networks, VR/AR/MR systems, network protocols, and real-world system measurement.

Zhimeng Yin is an assistant professor in the Department of Computer Science at City University of Hong Kong. He has a broad interest in cyber-physical systems (CPS) or Internet of Things (IoT), including the embedded, wireless, and low-power sensor systems, mobile computing, localization, and wireless security.

Zhenhua Li is an associate professor with the School of Software, Tsinghua University. His research areas cover mobile networking/emulation and cloud computing/storage. He received his PhD in computer science from Peking University.

Zhe Ji is the Chief Scientist of XYZ10. His research areas cover AI applications and algorithms. He received his MS from Zhejiang University.

Yeqiang Han is the founder of XYZ10. His research interests are around embedded systems and hardware production. He received his PhD from Zhejiang University.

Qiang Xu is the founder of XYZ10. His research interests are broadly in networking systems and big data/AI applications in cloud computing. He received his PhD from the University of Michigan in 2013.

Wei Jiang is the Chief of the New Technology Section in the State Grid of China. He mainly focuses on delivering/applying AI technologies to the energy industry. He received his PhD from Zhejiang University.

REFERENCES

[1] D. Vasisht, S. Kumar, and D. Katabi, D. 2016. {Decimeter-Level} Localization with a Single {WiFi} Access Point. *Proceedings of the 13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, 165–178.

[2] Airport Benchmarking. 2016. Airlines and Airports are Beaconizing. https://www.airportbenchmarking.com/airlines-and-airports-are-beaconizing/

[3] Locatify. 2020. BLE Beacon Museum Guide with real-time user location. https://locatify.com/blog/case-studies/eldheimar-museum/.

[4] Y. Hu, F. Qian, Z. Yin, Z. Li, Z. Ji, Y. Han, Q. Xu, and W. Jiang. 2022. Experience: Practical indoor localization for malls. *Proceedings of the 28th Annual International Conference on Mobile Computing and Networking*, 82–93.

[5] P.M. Djuric, J.H. Kotecha, J. Zhang, Y. Huang, T. Ghirmai, M.F. Bugallo, and J. Miguez. 2003. Particle filtering. *IEEE Signal Processing Magazine*, 20(5), 19-38.

[6] M. Wardat, W. Le, and H. Rajan. May 2021. Deeplocalize: Fault localization for deep neural networks. IEEE. 2021 *IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, 251-262.

[7] WebFX. 2020. What Is a Good Conversion Rate (and How Does Yours Compare)? https://www.webfx.com/blog/marketing/what-is-a-good-conversion-rate/.