

Short Paper: Towards Low-Cost Indoor Localization using Edge Computing Resources

Shweta Prabhat Khare

Janos Sallai

Abhishek Dubey

Aniruddha Gokhale

Institute for Software Integrated Systems, Vanderbilt University Nashville, TN 37221, U.S.A.

Abstract—Emerging smart services, such as indoor smart parking or patient monitoring and tracking in hospitals, incur a significant technical roadblock stemming primarily from a lack of cost-effective and easily deployable localization framework that impedes their widespread deployment. To address this concern, in this paper we present a low-cost, indoor localization and navigation system, which performs continuous and real-time processing of Bluetooth low Energy (BLE) and IEEE 802.15.4a compliant Ultra-wideband (UWB) sensor data to localize and navigate the concerned entity to its desired location. To keep deployment costs down, the indoor space in our solution is instrumented with (battery) as well as wired Edison devices, which provide both compute and BLE capabilities. Entities with managerial responsibilities in these scenarios can be equipped with both localization modalities: UWB tags and a BLE capable device (current generation smartphone or tablet), and are set up to maintain the BLE Received Signal Strength Intensity (RSSI) fingerprint map using the UWB positioning data as ground truth. The remaining entities rely exclusively on BLE RSSI fingerprinting-based localization using their smartphones.

I. INTRODUCTION

A significant technical roadblock that impedes the widespread deployment of smart applications such as smart parking, patient monitoring and indoor navigation systems stem from a general lack of an inexpensive entity localization solution. Existing localization solutions are either too costly (e.g. LIDAR, imaging sensors, radar), or incur high installation costs, requiring power and communications wiring for a large number of sensors. To address these technical difficulties, in this paper we present our low-cost system which performs continuous, on-line processing using a set of Edison compute nodes that can be dynamically deployed in the indoor space of interest. The Edison nodes provide Bluetooth Low Energy (BLE) messages, which when combined with a dynamic IEEE 802.15.4a compliant Ultra-wideband (UWB) sensor streams provide us the ability to create and process dynamic fingerprints allowing us to localize with a high accuracy with dynamic and low-cost edge devices. These finger prints rely on on BLE Received Signal Strength Intensity (RSSI).

Our approach is supported by prior research that suggests that fingerprint based localization using WiFi [1], [2], [3], [4], [5] or IEEE 802.15.4 compliant radios [6], [7] can yield an average positioning accuracy under 3m. However, it has been observed that BLE is substantially more susceptible to fast fading than IEEE 802.15.4 or WiFi [8], since BLE beaconing operates on a much narrower, 2 MHz wide channel as opposed to the 22 MHz wide channel used by WiFi and 5 MHz wide channel used by 15.4. Hence, our solution also

relies on the more accurate Ultra-wideband (UWB) Radio Frequency (RF) time-of-flight based ranging and positioning as the ground-truth for periodically updating the fast-fading BLE RSSI fingerprints. UWB is another indoor localization technology that has reached maturation and has broken a low price point (e.g., transceivers from vendors such as Decawave [9] are available in the sub-\$10 range). With direct line of sight between UWB beacons and a wireless tag which is to be localized, decimeter-scale accuracy has been reported indoors. However, since these UWB radios are not in current handheld devices, this technology alone is not sufficient for the smart applications that are easily available to the general public.

The approach we propose in this paper is to combine the above two modalities, i.e., BLE RSSI fingerprinting and UWB RF positioning using a computation architecture that allows us to fuse this information in real-time. We expect that frequent visitors, e.g. drivers with permits, or parking patrol vehicles, or nurses in hospitals or waiting staff in large restaurants or resorts will be equipped with both UWB tags and BLE-enabled Android devices, and contribute to building and maintaining the RSSI fingerprint map. For entities without UWB tags, we rely on indoor positioning by BLE RSSI fingerprint matching using the Android BLE APIs available today. Since RSSI fingerprint maps go “stale” quickly due to fast fading of the BLE channel characteristics, we rely on such dual-modality-equipped frequent occupants of these indoor spaces to keep the fingerprint map up-to-date. In this approach, we focus on running the computation at the edge (near the source of data) as opposed to a back-end cloud. Streaming all the sensor data to the cloud for processing will incur a prohibitively large latency which is highly undesirable for these applications. This is the key concept behind Fog Computing [10] and Mobile-Edge [11] computing paradigms.

Related Work: A wealth of research work exists in the field of indoor localization. Broadly, approaches for indoor localization can be categorized into [12]: angle-based, fingerprinting-based and ranging-based solutions. Angle-based approaches require expensive and specialized directional antenna which is unsuitable for widespread adoption [12]. Fingerprinting [13] and Radio RSS (Wifi, Bluetooth, ZigBee) based solutions [14], [15], [16], [17] are the least expensive solutions, requiring no specialized hardware, which makes them amenable for a wide variety of consumer applications. Prior efforts [17], [18], have explored the BLE transmission model- reception probabilities at varying distances and in the presence of obstruction- and have successfully demonstrated the use of BLE for localization.

Although, RSS-based methods are easily accessible, they are less-accurate, require prior-site information, cost-intensive fingerprinting and continuous updating to accommodate time varying channel characteristics. Similar to previous efforts which rely on crowd-sourcing [19], [20], we rely on frequent visitors like patrol vehicles in a parking garage, to keep the fingerprints up-to-date. Time-of-arrival (TOA) based ranging methods, such as UWB ranging [21], are much more robust and accurate. However, UWB technology is still under development and is not available on current handheld devices. Hence, our solution combines low-accuracy, easily-accessible BLE RSSI fingerprinting with high-accuracy UWB ranging.

II. RSSI FINGERPRINTING

BLE operates on a 2.4GHz ISM band which is divided into forty 2 MHz wide channels. Out of these 40 BLE channels, three are dedicated as advertising channels, and the rest are used for data exchange. BLE beacons are a special class of BLE devices which are limited to BLE transmission-only functionality and use the BLE advertising channels for periodic beaconing. However, we found that the BLE channel number of the incoming beacons is not exposed through any of the Android APIs, which makes RSSI fingerprinting-based localization challenging with Android devices.

Hence, we model the RSSI value of a beacon on a given BLE advertising channel as a Gaussian random variable. Therefore, as shown in Figure 1, the set of received RSSI values from a given beacon can be treated as a mixture of three Gaussians (one for each advertising channel). Collecting a sufficient number of RSSI observations at a given position allows us to characterize the empirical distribution of the received RSSI values. Intuitively, the measured distribution (i.e., the position and magnitude of peaks, etc.) will be characteristic of the given position, and will be sufficiently different even if the position changes by as little as a meter. Furthermore, by increasing the number of BLE beacon devices, we expect that the efficiency of distinguishing positions by BLE RSSI fingerprints will increase, as more data is available from a spatially diverse set of beacons.

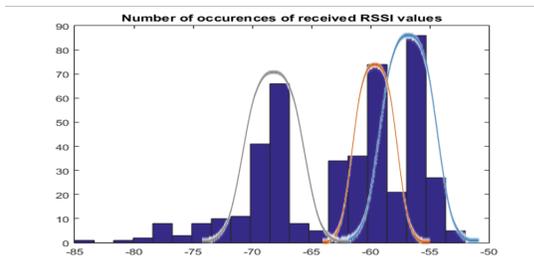


Fig. 1. A representative histogram of RSSI values received from a beacon node. The three peaks correspond to the mean RSSI value received on each of the three advertising channels, though the association between beacon messages and the corresponding channels is not available through the Android BLE API. By modeling the histogram as a mixture of three Gaussian kernels, the mean and variance of the RSSI values can be estimated on a per-channel basis.

Ideally, we would like the fingerprint map to remember the shape of each beacon’s empirical RSSI histogram at each grid point of the discretized coverage area. Storing the actual histograms, however, is not beneficial as it would have an unnecessarily large memory footprint. Moreover, the empirical histograms often contain transient peaks that are artifacts of fast fading, and capturing these would adversely affect localization accuracy. Hence, we have devised two methods for representing RSSI fingerprints: Expectation Maximization (EM) and Tercile-based methods.

In the EM-based method, we use the Expectation Maximization algorithm to find the parameters $\langle prior, \mu, \sigma \rangle$ for each of the three Gaussian kernel components, where *prior* is the weight of the kernel, μ and σ are the mean and standard deviation, respectively. If there are N BLE beacon devices deployed in the system, the fingerprint map for each grid position will comprise N , 9-field long vectors.

Since the EM algorithm is computationally expensive, and may not be suitable for devices with power or computational constraints (our Edison Nodes), we devised a lightweight fingerprinting solution relying on finding the terciles of the empirical RSSI distributions. Assuming that a beacon transmits equal number of messages on all three channels and that the standard deviation of the RSSI values is the same for each channel, the lowest third of the received RSSI values will correspond to the beacons received on the channel that has the highest path loss of the three. Similarly, the highest third will correspond to the channel with the lowest path loss, leaving the middle third coming on the third channel. Therefore, we use the medians of these three parts, i.e., the 1st 6-quantile, the 3rd 6-quantile (which is the median of the entire population), and the 5th 6-quantile as fingerprints. The tercile-based approach yields $3N$ -long feature vectors per grid position, where N is the number of beacons. The tercile-based approach trades off computational efficiency for accuracy. Unlike EM, this technique does not capture the variance of the RSSI values nor does it adapt to scenarios when the number of beacon messages received on different advertising channels is asymmetric.

For localization, it is sufficient to compare the Euclidean norm (L_2 distance) of the feature vectors of the target entity’s unknown position with the precomputed feature vectors of all grid positions, and return the grid position for which the L_2 distance is the smallest.

III. IMPLEMENTATION AND EXPERIMENTAL SETUP

We implemented our ideas and conducted a number of experiments to validate our claims. One of these experiments, reported in this paper was performed in a 6.4x5.5 meter area which was divided into a logical grid composed of 1x1 meter grid-cells. In this experiment, an Intel Nuc (with 1.6GHz Intel Core i5 processor, 3 MB cache and 16 GB memory) is used for processing of RSSI and UWB messages to carry out fingerprinting and localization. However, when the UWB devices are not in use we can do the localization on the edisons themselves.

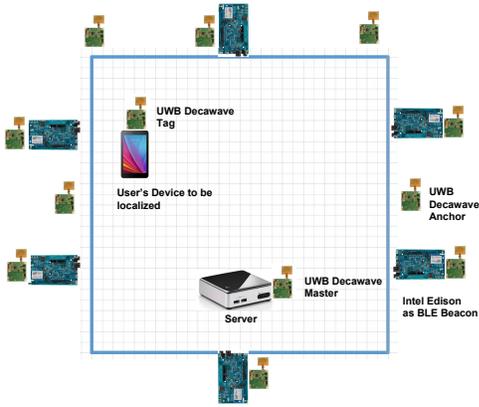


Fig. 2. Indoor localization experiment setup with 6 Intel Edison boards acting as BLE beacons; 12 UWB Decawave sensors; centralized server for sensor processing and the android tablet whose position is to be localized.

Six Intel Edison Arduino boards were used to act as BLE beacons in this experiment using open source *Bleno*¹ and *Noble*² libraries for beaconing. Additionally, 12 DWM1000 UWB sensors (i.e., Decawave) were used. We expect these receivers (battery operated) to be hung at different locations in the space where the localization solution is being implemented. An Android tablet was used to simulate both the entity with the decawave sensor (high accuracy line of sight) as well as the BLE beacon based localization (after turning off the decawave.). Figure 2 shows this experiment setup.

The Android tablet relays the RSSI messages received from all the beacons to the node doing the localization using Lightweight Communications and Marshalling (LCM) [22] publish-subscribe messaging library. LCM uses UDP Multicast for messaging and is well-suited for high-bandwidth and low-latency applications such as sensor-based localization. The localized tag's position information from the Decawave sensors, which we use as the ground truth for fingerprinting (used only when the decawave tag on the tablet was on) and assessing the performance of beacon based localization, is also received by the server on LCM. The server receives these RSSI messages and decawave based localized tag positions on two different LCM receiver threads. After every 300 milliseconds, a fingerprinting thread is scheduled by the server to re-compute and update the grid fingerprints. The fingerprinting thread recomputes both Expectation Maximization (EM) and Tercile-based grid fingerprints in parallel on a thread pool. The server also runs a localization thread every 30 milliseconds to perform beacon based localization of the android tablet.

During fingerprinting, the Decawave-based tag position is used as the ground truth to determine which grid-cell the tag/Android tablet is in. For a given grid-cell that the tablet is in (determined by the current Decawave-based position), 200 RSSI messages are collected per beacon (1 message is generated every 20 ms) before computing that grid-cell's fingerprint. Hence, 1,200 RSSI messages (6 beacons and 200

¹<https://github.com/sandeepmistry/bleno>

²<https://github.com/sandeepmistry/noble>

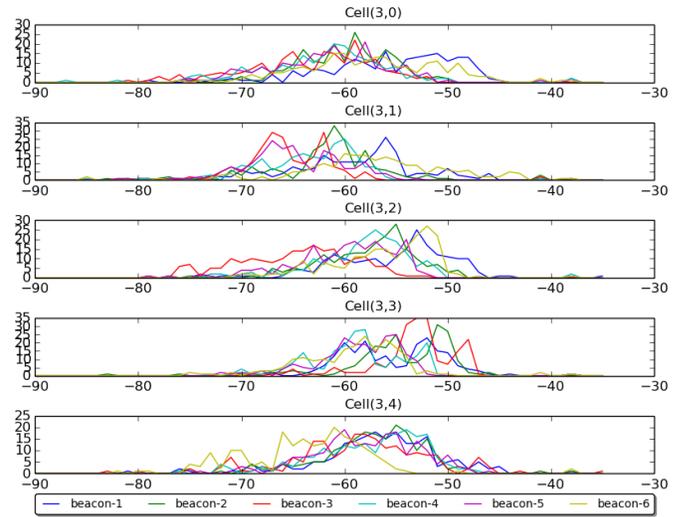


Fig. 3. Per-beacon RSSI histogram for adjacent grid cells in row 3.

messages per beacon) are used for computing a grid-cell's fingerprint. Figure- 3 shows how the RSSI distribution per beacon for adjoining grid-cells in row 3 differs from each other.

A grid-cell's RSSI messages for a beacon are maintained in a circular buffer, wherein new updates replace the old RSSI values in a FIFO order. When the fingerprinting thread comes up, it determines which grid-cell's RSSI values have been updated and uses the same snapshot of the updated grid-cell's RSSI values to recompute both EM and Tercile-based fingerprints in parallel. The fingerprint of this unknown location is then compared against the precomputed fingerprints of all grid-cells to find the best match.

For tercile based fingerprinting, the best match fingerprint is the one that minimizes the least squared distance from the unknown cell's fingerprint. The least squared distance for the tercile strategy is defined as sum of squares of the difference between corresponding tercile medians of unknown grid-cell u and fingerprinted grid-cell g for all six beacon. The equation has to be updated when we add more beacons to the system.

For EM based fingerprinting, the best match fingerprint is defined as the one that minimizes the least squared distance between the reconstructed Gaussian curves of unknown grid-cell u and fingerprinted grid-cell g for all six beacons. We reconstruct the Gaussian curve as represented by equation : $P(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/2\sigma^2}$ for RSSI values ranging from -90 to -35 using the mean and standard deviation provided by EM.

Experiment Results: We recorded the RSSI and UWB sensor data received while walking around the indoor location and this recorded sensor input was played back to perform multiple localization experiments. We changed the number of RSSI readings that are collected per beacon, also referred to as the "window" size for computing the fingerprint of the unknown grid location, so as to observe the effect of changing window sizes on the accuracy and computation time of fingerprinting. Each localization experiment was performed

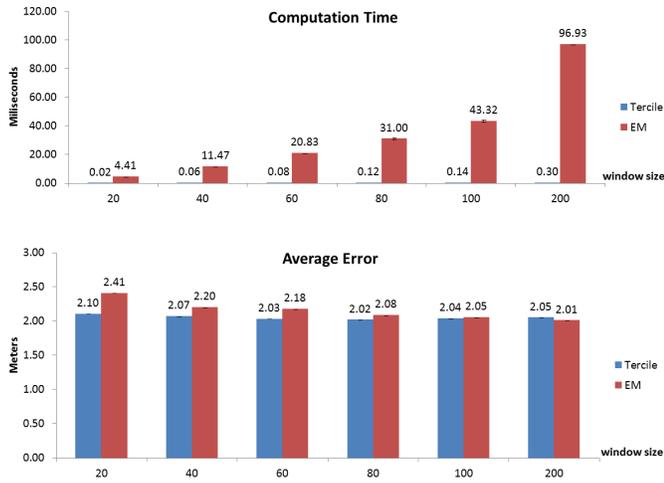


Fig. 4. Computation time and average error for EM and Tercile fingerprinting strategies with increasing window sizes.

for different window sizes: 20, 40, 60, 80, 100 and 200 by playing back the same recorded sensor data file. Figure 4 shows the computation time of tercile and EM approaches with increasing window sizes. Though EM’s computation time decreases with decreasing window sizes, EM is significantly more expensive than the tercile approach which has a computation time lower than 1 millisecond even for larger window sizes.

Figure 4 also shows the average error in meters for the two fingerprinting approaches. This error is the Euclidean distance between the localization result from EM or tercile strategy and the UWB based position information which is considered as the ground truth for assessing the accuracy of the two fingerprinting strategies. Since the average error for the two approaches across all window sizes is comparable (EM shows much higher error for window size=20), the tercile approach with sub millisecond computation time is preferable for fast beacon-based localization.

IV. CONCLUSION

In this paper, we described our solution for low-cost, multi-modal indoor localization based on Bluetooth Low Energy (BLE) RSSI fingerprinting and IEEE 802.15.4a compliant Ultra-wideband (UWB) RF time-of-flight based positioning. Currently, all sensor data processing happens at one server at a time. In future, we plan to distribute our localization and navigation algorithm to run on more than one server simultaneously. We are also working on fault-tolerant processing of sensor data to build a resilient smart indoor system.

Acknowledgement: This work has been funded under a research grant from Siemens, CT.

REFERENCES

[1] A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. S. Wallach, and L. E. Kavasaki, “Practical robust localization over large-scale 802.11 wireless networks,” in *Proceedings of the 10th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom ’04. ACM, 2004.

[2] T. Roos, P. Myllymaki, H. Tirri, P. Misikangas, and J. Sievanen, “A probabilistic approach to wlan user location estimation,” *International Journal of Wireless Information Networks*, 2002.

[3] A. Naguib, P. Pakzad, R. Palanki, S. Poduri, and Y. Chen, “Scalable and accurate indoor positioning on mobile devices,” in *Indoor Positioning and Indoor Navigation (IPIN), 2013 International Conference on*. IEEE, 2013, pp. 1–10.

[4] M. Youssef and A. Agrawala, “The horus wlan location determination system,” in *Proc. MobiSys*, 2005, pp. 205–218.

[5] S. Yoon, K. Lee, and I. Rhee, “Fm-based indoor localization via automatic fingerprint db construction and matching,” in *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys ’13. ACM, 2013.

[6] S.-Y. Lau, T.-H. Lin, T.-Y. Huang, I.-H. Ng, and P. Huang, “A measurement study of zigbee-based indoor localization systems under rf interference,” in *Proceedings of the 4th ACM International Workshop on Experimental Evaluation and Characterization*, ser. WINTECH ’09. New York, NY, USA: ACM, 2009, pp. 35–42.

[7] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen, “Zee: Zero-effort crowdsourcing for indoor localization,” in *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking*, ser. Mobicom ’12. New York, NY, USA: ACM, 2012, pp. 293–304.

[8] C. Frost, C. Jensen, K. Luckow, B. Thomsen, and R. Hansen, “Bluetooth indoor positioning system using fingerprinting,” ser. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering. Springer Berlin Heidelberg, 2012.

[9] [Online]. Available: <http://www.decawave.com/products/dwm1000-module>

[10] A. Ahmed and E. Ahmed, “A survey on mobile edge computing,” in *2016 10th International Conference on Intelligent Systems and Control (ISCO)*, Jan 2016, pp. 1–8.

[11] N. Fernando, S. W. Loke, and W. Rahayu, “Mobile cloud computing: A survey,” *Future Generation Computer Systems*, vol. 29, no. 1, pp. 84–106, 2013, including Special section: AIRCC-NetCoM 2009 and Special section: Clouds and Service-Oriented Architectures.

[12] K. Liu, X. Liu, and X. Li, “Guoguo: Enabling fine-grained indoor localization via smartphone,” in *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys ’13. ACM, 2013.

[13] M. Azizyan, I. Constandache, and R. Roy Choudhury, “Surroundsense: Mobile phone localization via ambience fingerprinting,” in *Proceedings of the 15th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom ’09. ACM, 2009.

[14] J. Yang and Y. Chen, “Indoor localization using improved rss-based lateration methods,” in *GLOBECOM 2009 - 2009 IEEE Global Telecommunications Conference*, 2009, pp. 1–6.

[15] C. Yang and H. r. Shao, “Wifi-based indoor positioning,” *IEEE Communications Magazine*, vol. 53, no. 3, March 2015.

[16] Z. Chen, H. Zou, H. Jiang, Q. Zhu, Y. C. Soh, and L. Xie, “Fusion of wifi, smartphone sensors and landmarks using the kalman filter for indoor localization,” *Sensors*, vol. 15, no. 1, pp. 715–732, 2015.

[17] P. Martin, B.-J. Ho, N. Grupen, S. Muñoz, and M. Srivastava, “An ibeacon primer for indoor localization: Demo abstract,” in *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*, ser. BuildSys ’14. ACM, 2014.

[18] A. S. Paul and E. A. Wan, “Rssi-based indoor localization and tracking using sigma-point kalman smoothers,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 3, no. 5, pp. 860–873, 2009.

[19] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen, “Zee: Zero-effort crowdsourcing for indoor localization,” in *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking*, ser. Mobicom ’12. ACM, 2012.

[20] K. Chintalapudi, A. Padmanabha Iyer, and V. N. Padmanabhan, “Indoor localization without the pain,” in *Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking*, ser. MobiCom ’10. ACM, 2010, pp. 173–184.

[21] C. Zhang, M. Kuhn, B. Merkl, A. E. Fathy, and M. Mahfouz, “Accurate ubw indoor localization system utilizing time difference of arrival approach,” in *Radio and Wireless Symposium, 2006 IEEE*. IEEE, 2006.

[22] “LCM: Lightweight communications and marshalling,” <https://lcm-proj.github.io/>.