Radio Map Optimization Through Unsupervised Learning for Indoor Localization

Jens Trogh*, Wout Joseph*, Luc Martens* and David Plets*

* Department of Information Technology, IMEC - Ghent University, Belgium, jens.trogh@ugent.be

Abstract—A major burden of signal strength based fingerprinting for indoor positioning is the generation and maintenance of a radio map, also known as a fingerprint database. Model-based radio maps are generated much faster than measurement-based radio maps but are generally not accurate enough. This work proposes a method to automatically construct and optimize a model-based radio map. The method is based on unsupervised learning, i.e., random walks for which the ground truth locations are unknown, serve as input for the optimization, along with a floor plan and a location tracking algorithm. No measurement campaign or site survey, which are labor-intensive and time-consuming, or inertial sensor measurements, which are often not available and consume additional power, are needed for this approach. Experiments in a large office building, covering over 1100 m^2, resulted in a median accuracy of 2.2 m, or an improvement of 60% with less than 5 minutes of unlabeled training data.

Index Terms—localization, tracking, positioning, fingerprinting, rss, indoor environment, unsupervised learning

I. Introduction

Localization and tracking in indoor environments is important for a wide range of location-aware applications, e.g., navigation in a shopping mall, finding your car in a parking garage, or asset tracking in the industrial sector. Most positioning systems, in GPS-denied environments, rely on signal strength measurements from existing wireless network infrastructures due to their simplicity and availability, e.g., WiFi, ZigBee or Bluetooth Low Energy (BLE) compatible devices. These Received Signal Strength (RSS) measurements can be translated to a location by making use of a path loss model and the well-known multilateration method [1]. Alternatively, with a fingerprinting technique, the position of an unknown user or object is estimated by looking for the closest match in the fingerprint database (online phase) [2], [3]. The radio map, or fingerprint database, is a signal space that links RSS values to positions in a building. This database is constructed in an offline training phase by making use of a ray-tracer simulator or an elaborate measurement campaign, also known as war-driving [4]. The first approach is simulation-based, hence much faster but will generally lead to less accurate location estimations. The second approach consists of manually performing RSS measurements at known locations (grid points) and needs to be redone each time the wireless network or even the office layout undergoes changes. Other localization systems use special-purpose hardware and infrastructure like ultra-wideband (UWB), radio-frequency identification (RFID) or acoustic ranging [5]–[7]. These systems can be very accurate but the initial deployment cost will generally be much higher.

Another technique is pedestrian dead-reckoning which uses inertial sensors that are typically embedded in smartphones, e.g., accelerometers, magnetometers, and gyroscopes [8]. The positions are calculated based on a previous position and the estimated movement of a user, by detecting steps, estimating stride lengths and the direction of motion. These systems are typically prone to drift, i.e., the positioning error accumulates over time because of noise in the inertial sensor measurements. In this work, an unsupervised learning method to automatically construct, maintain and optimize fingerprint databases, without the need for inertial sensor units, calibration or extensive measurements, is proposed.

II. Related work

In the past, other techniques for indoor localization without the need for pre-deployment efforts, e.g., site survey, measurement campaign or device calibration, have been proposed [9]–[13]. The EZ algorithm is a configuration-free indoor localization scheme that uses a genetic algorithm and occasionally available GPS locks, e.g., at the entrance or near a window, to localize mobile devices [9]. Another technique that bypasses war-driving is UnLoc [10], which uses dead-reckoning, urban sensing, and WiFi-based partitioning. A dead-reckoning scheme is used to track a user's smartphone between so called internal landmarks of a building, e.g., a distinct pattern on a smartphone's accelerometer or an unusual magnetic fluctuation in a specific spot. In [11], WiFi and inertial sensor information are combined with constraints imposed by a map of the indoor space of interest and an augmented particle filtering is used to estimate the position concurrently with other variables such as the stride length. A joint indoor localization and radio map construction scheme is presented in [12]. This scheme employs manifold alignment to produce a projection between a source spatial correlation preserving data set and a limited number of calibration fingerprints. A crowd sourcing-based scheme to construct a probabilistic radio map based on parametric fitting is presented in [13]. This technique describes location signatures by transforming RSS into signal envelopes but relies on an additional localization mechanism and a large amount of RSS samples.

Our approach does not rely on any calibration measurements, GPS fixes or inertial sensor units, e.g., accelerometers, gyroscopes, or magnetic compasses. To the best of our knowledge, this is the first unsupervised learning approach that relies solely on floor plan information and unlabeled RSS data. No manual calibration, measurement campaign, or additional
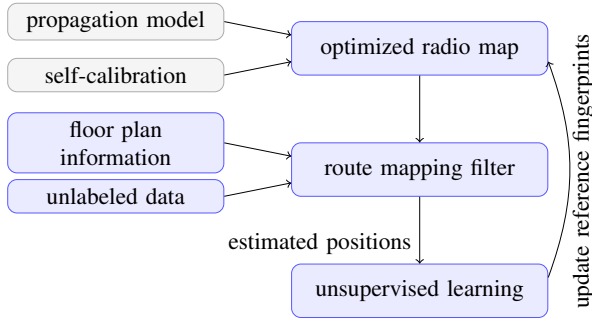
Fig. 1. Flow graph of the proposed radio map optimization technique. The initialization and regular flow are indicated in gray and blue, respectively.

sensor information are needed to construct, maintain, and optimize radio maps for indoor localization, e.g. to make model-based databases more accurate or to automatically cope with changes in an office layout.

III. METHODOLOGY

Our approach consists of an initial model-based radio map, a self-calibration technique to match a user device with this radio map, an unsupervised learning technique to optimize this radio map, and a route mapping filter to reconstruct the most likely trajectory of unlabeled training data by including floor plan information. Figure 1 shows a flow graph of the proposed radio map optimization technique.

A. Initial radio map

The initial fingerprint database is a simulated radio map, that can be based on any propagation model. Our approach uses a theoretical model for an indoor environment that includes wall and interaction losses [14]. Here, RSS values are modeled as:

$$RSS_{ref} = \underbrace{RSS_0 - 10\gamma \log_{10}\left(\frac{d}{d_0}\right)}_{\text{distance loss}} - \underbrace{\sum_i L_{W_i}}_{\text{cumulated wall loss}} - \underbrace{\sum_j L_{B_j}}_{\text{interaction loss}} + X_\sigma \quad [dB]$$

$$(1)$$

RSS_{ref} [dB] denotes the received signal strength, RSS_0 [dB] is the received signal strength at a reference distance d_0 [m], γ [-] is the path loss exponent, d [m] is the distance along the path between transmitter and receiver. These two terms represent the path loss due to the traveled distance. The cumulated wall loss represents the sum of all wall losses L_{W_i} when a signal propagates through a wall W_i. The interaction loss represents the cumulated losses L_{B_j} caused by all propagation direction changes B_j along the path between sender and receiver, and X_σ [dB] is a log-normally distributed variable with zero mean and standard deviation σ, corresponding to the large-scale shadow fading.

B. RSS self-calibration

Reported RSS values tend to be very different, depending on a user's device. Therefore, a self-calibration method is used to obtain a good mapping between the measured RSS values and the reference values from the fingerprint database [15]. This method relates the RSS histogram from the reference radio map to the RSS histogram of a user's device, which requires no user intervention or ground truth location data, and results in the following bias:

$$RSS_{bias} = med(F_{RSS_{ref}}^{-1}(y) - F_{RSSmeas}^{-1}(y)), \quad y \in \{0.1, 0.2, ..., 0.9\} \quad (2)$$

RSS_{bias} is the bias between the measurements of the user device and the fingerprint database, and is equal for all access points. $F_{RSS_{ref}}$ is the empirical cumulative distribution functions (cdf) of the reference fingerprints and $F_{RSS_{meas}}$ is the empirical cdf of the measurements from the user's device. For our unlabeled training data, the value of RSS_{bias} stabilized after 15 seconds of measurement data (the sending rate was fixed at 5 Hz). This RSS_{bias} value is added to all reference RSS values of our initial fingerprint database (RSS_{ref} in equation 1).

C. Unsupervised learning

An experiment showed that measurements of neighboring locations are similar and deviations from the fingerprint database tend to be correlated per room and per access point. This experiment is done in a regular office building where all rooms shaped as rectangles, except for the corridor (Figure 2). The differences between measurements and reference fingerprints have a standard deviation of 7.8 dB for the whole building, whereas the average standard deviation is only 2.1 dB if they are grouped by room and per access point. Under the assumption that a user's trajectory can be roughly reconstructed, the deviations per room and access point can be learned, resulting in a fingerprint database that matches the actual measurements more closely. Consequently, this optimized fingerprint database can increase the positioning accuracy of the trajectories or static locations of other users or objects. In our approach, the trajectories of unlabeled training data are first reconstructed with a route mapping filter. Next, the differences (RSS_{diff}) between measured RSS values and the corresponding reference fingerprint are grouped per room and per AP, based on the timestamps of the measurements and the estimated location at that timestamp. Then, the reference fingerprints, for which there are measurements in its room, are updated by adding the median value of these RSS_{diff}.

$$RSS_{new}^{i,j} = RSS_{ref}^{i,j} + med(RSS_{diff}^{i,room_j}) \quad (3)$$

$RSS_{new}^{i,j}$ and $RSS_{ref}^{i,j}$ are the new and old reference RSS value for access point i and grid point j. (A grid point is a location on the floor plan for which there is a reference RSS value in the fingerprint database and the grid size was set to 50 cm.) The median difference between all measurements from access point i that are labeled (estimated) to be in the same room as grid point j and the correspondent reference fingerprints, is represented by $med(RSS_{diff}^{i,room_j})$. This process can be applied to new unlabeled data or even iteratively on the same random walk more then once, because the estimated trajectories tend to become more accurate in the next iteration.

D. Route mapping filter

The trajectory of unlabeled training data is reconstructed with a route mapping filter that is based on the Viterbi path, a technique related to hidden Markov models [16], [17]. It uses an additional motion model and floor plan information to determine the most likely path (i.e., sequence of locations) instead of only the most likely current position. These constraints ensure that no unrealistically large distances are traveled within a given time frame and no walls are crossed. By processing all available data at once, previous estimated locations can be corrected by future measurements (similar to backward belief propagation). This route mapping filter makes it possible to optimize the fingerprint database because the estimated positions, along the reconstructed trajectory, are generally assigned to the correct room. Hence, the discrepancies between reference fingerprints and real measurements can be learned, and the radio map quality and positioning accuracy can be improved. This is less likely with stateless positioning techniques, where consecutive estimated positions can fluctuate between different rooms because of measurement noise and outliers.
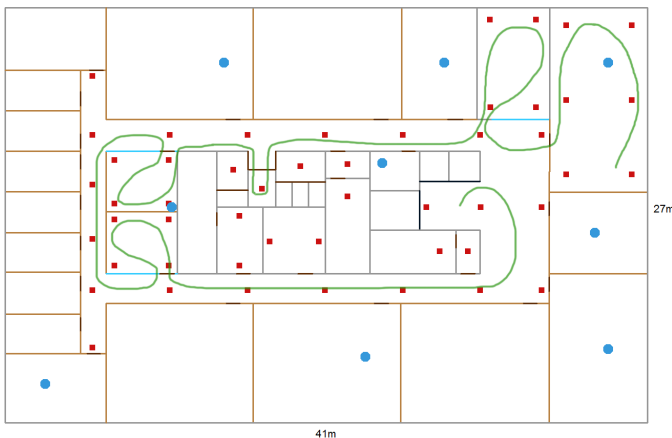
IV. EVALUATION

Fig. 2. Floor plan with location of the access points (blue dots), validation test points (red squares) and training data (roughly indicated with a green line).

The experiments are conducted on a wireless testbed, located on the ninth floor of an office building in Ghent, covering over 1100 m^2 (41 m by 27 m, see Figure 2). The inner structure of the building is made of thick concrete walls (gray) and the meeting rooms, offices, and kitchen have plaster walls (amber) and wooden doors (brown). The wireless network consists of 9 fixed sensor nodes that are installed at a height of 3 m and are indicated with a blue dot in Figure 2. These sensor nodes are based on the Zolertia RE-mote platform, which is based on the Texas Instruments CC2538 ARM Cortex-M3 system on chip, with an on-board 2.4 GHz IEEE 802.15.4 RF interface, running at up to 32 MHz with 512 KB of programmable flash and 32 KB of RAM, bundled with a Texas Instruments CC1200 868/915 MHz RF transceiver to allow dual band

operation [18]. The mobile tag is based on the same platform and broadcast 5 packets per second which are received by the fixed nodes. Every second a location update is generated (the average RSS values of the packets received within this second are used as input for the location tracking algorithm).

The fifty validation points (test data) are indicated with a red square and are used as separate testing locations, i.e., no tracking algorithm is applied, only positioning by searching for the closest match in the (optimized) fingerprint database. To diminish the effect of temporal fading, the RSS values for the validation points are averaged over a 20 seconds measurement interval, i.e. 100 packets if there is no packet loss (sending rate is fixed at 5 Hz). The training data consists of a random walk along the corridor, kitchen, and meeting rooms, and is roughly indicated with a green line. The exact positions are not known and not needed for the learning phase (hence, unsupervised) but are indicated to give an idea of the covered area. During this walk, that lasts only 200 seconds, 7800 measurements are logged, which corresponds to a packet loss of 13%, due to fading or interference.

V. RESULTS

The experimental validation consists of four scenarios and are summarized in Table I. The first scenario uses neither the self-calibration method nor the unsupervised learning technique. The second and fourth scenario use unsupervised learning to optimize the fingerprint database. The third and fourth scenario use self-calibration to match a user's device with the fingerprint database.

TABLE I
VALIDATION POSITIONS: ACCURACY PER SCENARIO

Scenario	μ [m]	σ [m]	50th [m]	95th [m]
No learning + no calibration	6.58	3.49	5.66	12.92
Unsupervised learning + no calibration	5.83	3.89	4.51	13.91
No learning + self-calibration	3.61	2.5	3.41	8.77
Unsupervised learning + self-calibration	2.64	1.70	2.24	5.54

The initial accuracy, without unsupervised learning or self-calibration, is 6.58 m on average. Applying the unsupervised learning technique on the training data, without calibration of the user's device, improves the mean and median accuracy of the test data to 5.83 m and 4.51 m, which corresponds to a relative improvement of 11% and 20%, respectively. However, the 95th percentile error deteriorates with 8% because some differences between the real measurements and the reference fingerprints are learned incorrectly. This happens when the estimated trajectory from the training data deviates too much from some of the real locations, which causes the fingerprint database to learn RSS values measured in other rooms.

Applying the self-learning calibration technique on the training data, improves, the mean and median accuracy of the test data to 3.61 m and 3.41 m, which corresponds to a relative improvement of 45% and 40%, respectively. Combination of the unsupervised learning technique after the self-calibration results in the largest improvements and does not result in a higher 95th percentile error due to incorrect learned RSS

values. The mean, standard deviation, median, and 95[th] percentile value are 2.6 m, 1.7 m, 2.2 m, and 5.5 m, respectively, which correspond to a relative improvement of 60%, 51%, 60%, and 57%, respectively. Compared to the scenario without unsupervised learning but with self-calibration (third scenario), the improvements are 27%, 29%, 34%, and 37%, for the mean, standard deviation, median, and 95[th] percentile value, respectively. The latter being a better comparison to show the effectiveness of solely the proposed unsupervised learning technique.

VI. CONCLUSIONS

This work presents an unsupervised learning technique to optimize fingerprint databases for indoor positioning systems, e.g. to make model-based databases more accurate or to automatically cope with changes in an office layout. The proposed technique does not rely on extensive measurements, device calibration or additional, power consuming, inertial measurement units but instead uses random, unlabeled training data, a self-calibration method, and a route mapping filter. The premise of this work is that deviations between real measurements and reference RSS values from a model-based radio map tend to be correlated per room and per access point. These differences can be learned, even by roughly reconstructing the random walks that a typical person does. This results in reference fingerprints that match the real measurements more closely, and hence, will lead to better positioning accuracies for every user. The absolute median accuracy improved from 5.7 m to 2.2 m, which corresponds to a relative improvement of 60%. The mean and standard deviation showed similar improvements. Future work will include, but is not limited to, test and training data with multiple, simultaneously active users, influence of the access point density, covering multiple floors, and ability to recover from worse fingerprinting maps or from physical changes in the environment.

ACKNOWLEDGMENT

REFERENCES

[1] Federico Thomas and Lluís Ros. Revisiting trilateration for robot localization. *IEEE Transactions on robotics*, 21(1):93–101, 2005.

[2] Paramvir Bahl and Venkata N Padmanabhan. Radar: An in-building rf-based user location and tracking system. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 775–784. Ieee, 2000.

[3] Paramvir Bahl, Venkata N Padmanabhan, and Anand Balachandran. Enhancements to the radar user location and tracking system. *Microsoft Research*, 2(MSR-TR-2000-12):775–784, 2000.

[4] Ionut Constandache, Shravan Gaonkar, Matt Sayler, Romit Roy Choudhury, and Landon Cox. Enloc: Energy-efficient localization for mobile phones. In *INFOCOM 2009, IEEE*, pages 2716–2720. IEEE, 2009.

[5] Abdulrahman Alarifi, AbdulMalik Al-Salman, Mansour Alsaleh, Ahmad Alnafessah, Suheer Al-Hadhrami, Mai A Al-Ammar, and Hend S Al-Khalifa. Ultra wideband indoor positioning technologies: Analysis and recent advances. *Sensors*, 16(5):707, 2016.

[6] Lionel M Ni, Dian Zhang, and Michael R Souryal. Rfid-based localization and tracking technologies. *IEEE Wireless Communications*, 18(2), 2011.

[7] Zaher Merhi, Mohamed Elgamel, and Magdy Bayoumi. A lightweight collaborative fault tolerant target localization system for wireless sensor networks. *IEEE Transactions on Mobile Computing*, 8(12):1690–1704, 2009.

[8] Antonio R Jimenez, Fernando Seco, Carlos Prieto, and Jorge Guevara. A comparison of pedestrian dead-reckoning algorithms using a low-cost mems imu. In *Intelligent Signal Processing, 2009. WISP 2009. IEEE International Symposium on*, pages 37–42. IEEE, 2009.

[9] Krishna Chintalapudi, Anand Padmanabha Iyer, and Venkata N Padmanabhan. Indoor localization without the pain. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, pages 173–184. ACM, 2010.

[10] He Wang, Souvik Sen, Ahmed Elgohary, Moustafa Farid, Moustafa Youssef, and Romit Roy Choudhury. No need to war-drive: Unsupervised indoor localization. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*, pages 197–210. ACM, 2012.

[11] Anshul Rai, Krishna Kant Chintalapudi, Venkata N Padmanabhan, and Rijurekha Sen. Zee: Zero-effort crowdsourcing for indoor localization. In *Proceedings of the 18th annual international conference on Mobile computing and networking*, pages 293–304. ACM, 2012.

[12] Sameh Sorour, Yves Lostanlen, Shahrokh Valaee, and Khaqan Majeed. Joint indoor localization and radio map construction with limited deployment load. *IEEE Transactions on Mobile Computing*, 14(5):1031–1043, 2015.

[13] Qideng Jiang, Yongtao Ma, Kaihua Liu, and Zhi Dou. A probabilistic radio map construction scheme for crowdsourcing-based fingerprinting localization. *IEEE Sensors Journal*, 16(10):3764–3774, 2016.

[14] David Plets, Wout Joseph, Kris Vanhecke, Emmeric Tanghe, and Luc Martens. Coverage prediction and optimization algorithms for indoor environments. *EURASIP Journal on Wireless Communications and Networking*, 2012(1):1–23, 2012.

[15] Christos Laoudias, Robert Piché, and Christos G Panayiotou. Device self-calibration in location systems using signal strength histograms. *Journal of Location Based Services*, 7(3):165–181, 2013.

[16] Jens Trogh, David Plets, Luc Martens, and Wout Joseph. Advanced real-time indoor tracking based on the viterbi algorithm and semantic data. *International Journal of Distributed Sensor Networks*, 501:271818, 2015.

[17] Jens Trogh, David Plets, Arno Thielens, Luc Martens, and Wout Joseph. Enhanced indoor location tracking through body shadowing compensation. *IEEE Sensors Journal*, 16(7):2105–2114, 2016.

[18] Zolertia online resources and documentation, 2018 (accessed April 20, 2018). https://github.com/Zolertia/Resources/wiki.