

CRAFT Reducing the Effort for Indoor Localisation

Paul Crane, Zhiyi Huang, Haibo Zhang
Computer Science, University of Otago
Dunedin, New Zealand
Email: {pcrane,hzy,haibo}@cs.otago.ac.nz

Abstract—Indoor localisation systems have slowly become more and more accurate. Each localisation system needs tuning to affect reasonable performance. In this paper we propose CRAFT, a crowd sourced approach to constructing a WiFi fingerprint database. The method uses a temporarily deployment of a small number of anchor nodes to roughly locate the position of the WiFi sample. Through thorough experiments in a real-world building, CRAFT's error is 2.2m a decrease of 25% when compare to other published results.

Keywords—Indoor environments, Crowdsourcing, Received signal strength indicator, Indoor Navigation

I. INTRODUCTION

There are multiple indoor localisation techniques that use Received Signal Strength Indicator (RSSI) to locate a device. Trilateration, often requires a small amount of data to train the system at the cost of inaccurate distance measurements due to the variability of RSSI. An alternative approach, fingerprinting, can offer higher accuracies at the expense of large amounts of training data. Other researchers have used Pedestrian Dead-Reckoning (PDR) to locate the device [1], and rely on accelerometers, magnetometers, or gyroscopes (or a combination of these sensors) [2], and often require detailed floor plan information [3, 4, 5, 6, 7] which may not always be available, accessible, or useful (such as airports).

In this paper we combine the advantages of trilateration and fingerprinting and explore localisation in an environment where the floor plan information is inferred by the deployment of anchor nodes. By deploying a temporary localisation system and performing trilateration with crowdsourced data, locations of the WiFi fingerprints can be estimated. Once the fingerprint database has been completed the temporary localisation system can be removed.

CRAFT a novel technique that uses crowdsourcing to reduce the collection effort needed to construct a comprehensive and extensive fingerprint database. The ability of CRAFT to reduce the burden of fingerprint construction without sacrificing location accuracy is demonstrated through thorough experiments conducted in a real world building.

The rest of the paper is organised as follows: Section II gives an overview of recent work on training the fingerprint database for RSSI based indoor localisation systems. Section III motivates the approach by examination of the work needed to deploy a fingerprint based indoor localisation in a real world environment with the approach to crowdsourced construction described in Section IV. The testing procedures, environment, and evaluation of the approach are discussed Section V. Finally, the conclusions are presented in Section VI.

II. RELATED WORK

Three approaches for the construction of the fingerprint database are: supervised, where each fingerprint is labelled with a location; semi-supervised, where some fingerprints are labelled and other fingerprint labels are inferred; and finally, unsupervised, where none of the fingerprints are labelled [8].

An unsupervised approach is attractive as it minimises the amount of manual effort needed [3, 9, 10, 11]. However, accurate modelling, taking a wide variety of factors into account, is needed to ensure good radio propagation modelling. These factors consist of static and dynamic components, such as positions and orientations of walls, doors, windows, furniture, as well as their radio propagation properties (which are often difficult to obtain). Even if these details were easily obtainable, modeling the radio propagation is a complex process [12].

While a completely manual approach is an option, it is expensive in terms of the time needed to sample the space. A suggestion to reduce the time taken is to reduce either the density of sampling, or reduce the number of samples collected per point [13]. Both of these suggestions will result in a increase in error by either omitting some of the potential locations or by increasing the fingerprint's susceptibility to outliers. Neither of these options are desirable as they produce a poor-quality fingerprint database. These problems are explored in further detail in Section III.

A solution to the problem lies in a semi-supervised approach. One method to achieve the goal of minimal effort radio mapping is to interpolate between the sparsely collected sample points [14, 15]. While this reduces the number of sample points needed, these methods are sensitive to the interpolation method used. For example, using a naïve linear method is problematic as RSSI is not linear between two points.

Distributing the task of sampling across multiple people avoids the problems associated with either interpolation or modelling. A number of proposed indoor localisation systems already employ the power of crowdsourcing to build a fingerprint database, such as, leveraging opportunistic Global Positioning System (GPS) signals, or, PDR techniques both with and without floor plan information.

The PDR techniques, without floor plan information, generally suffer from sensor drift over periods of time and can produce vastly inaccurate estimates of location. For example, the mean error is from 7.6 m to 8.9 m in some indoor experiments [16]. Given the fingerprint database is built upon location estimates that can be highly inaccurate, approaches using PDR without floor plan information are too inaccurate. On the other hand, with the availability of floor plan information, PDR is used to compute the likely location which can then be refined using floor plan information [17, 18, 19, 20]. LIFS




Fig. 1: Areas and sample points used in the evaluation.

[5] uses the floor plan to generate a uniform grid of sample points which is then subjected to Multi-dimensional Scaling (MDS). The result is a stress-free floor plan. The RSSI data are collected along with the number of steps between subsequent measurements. Once similar fingerprints are merged the step counts are used to produce a distance matrix. This distance matrix is processed using the same MDS technique resulting in fingerprint space. The mapping between the stress-free floor plan and the fingerprint space can be calculated and used to locate devices. Another approach, Sub-area [6], splits the floor plan into different sub-areas. The crowdsourced WiFi data is then clustered using MDS and the clusters are matched to each of the sub-areas. KAILOS [7] breaks the floor plan into a grid (each cell is either $3\text{ m} \times 3\text{ m}$ or $2\text{ m} \times 2\text{ m}$). Each cell is then linked to its neighbours and itself to produce a Hidden Markov Model (HMM). Traces recorded in the environment are mapped onto this HMM and the associated WiFi fingerprints for each cell are used for localisation.

The key problem we solve is building the fingerprint database via crowdsourcing without relying on explicit floor plan information. In the next section a description of the effort needed to deploy a fingerprint-based RSSI system in a small test deployment is presented.

III. MANUAL CONSTRUCTION OF THE FINGERPRINT DATABASE

Deploying a fingerprinting system manually is a time consuming process. The fingerprint data is collected at pre-planned locations throughout the building. The test deployment was conducted in the Owheo Building, the floor plan of which is shown in Figure 1. It is a typical academic building with a mixture of shared office space (for postgraduate students), private office space (for academic staff) and corridors to access the rooms. The mezzanine ($7.1\text{ m} \times 10.2\text{ m}$) joins the two wings (the usable area of the mezzanine is reduced by a staircase occupying the south-west corner). The west corridor is comprised of two sub-corridors joined by a right angle (east-west $13.7\text{ m} \times 2.1\text{ m}$; the north-south $1.9\text{ m} \times 25.8\text{ m}$). The postgraduate suite is a shared office for students ($10.8\text{ m} \times 13.2\text{ m}$). The remaining rooms are private staff offices and are inaccessible for our experimentation.

Once the positions have been marked the collection of samples can begin. By default WiFi beacon frames are broadcast by the access points approximately every 0.1024 s . These frames contain the network’s Basic Service Set Identifier (BSSID), Service Set Identifier (SSID), and other parameters. As the beacon frames are received, the RSSI is measured, and

TABLE I: Estimated and reported data collection times for published crowdsourced fingerprint construction algorithms. The estimated time was calculated from (1) with $N_s = 25$, $N = 1$.

	A	R	T_e (hours)		Reported (hours)
			$T_p=1\text{ s}$	$T_p=5\text{ s}$	
LiFS [5]	1600	0.25	3.0	17.0	20.0
SubArea [6]	460	0.39	1.5	7.5	2.0
CRAFT	250	1.00	2.0	10.5	2.5

recorded along with the BSSID, SSID, and time. The dataset for the experiments is composed of 25 passively collected WiFi beacon frames per ground truth location. This collection took approximately 30 s per location to collect. Depending on the devices, sampling the WiFi may take between 1 s to 5 s per location.

$$T_e \propto \frac{ART_p}{N} \quad (1)$$

We estimate the amount of time, T_e , to perform the data collection according to (1). Where A is the area (m^2), R is the resolution measured as the number of sample points per m^2 , T_p is the estimate of the time per point, and N is the number of people performing the collection. This formulation does not take into account the administration effort, such as recording the position, moving from one position to the next, management of people, and so on. Our dataset was collected with a resolution of $1\text{ m} \times 1\text{ m}$, to match other published indoor localisation work, and resulted in a total of 97 sample locations throughout the accessible areas of the Owheo Building as shown by the dots in Figure 1. From the dimensions and time for WiFi scans the following parameters are obtained: $A = 250.7\text{ m}^2$, $R = 1$, $N_s = 25$, $T_p = 1\text{ s}$, $N = 1$, and thus, $T_e \propto 125.35\text{ min}$, or just over 2h. The collection for the dataset used in this work took 2.5 h to complete. Some examples of previous reported collection times and parameters (for (1)) are presented in Table I along with T_e . While crowdsourcing can allow for less manual effort, the location problem remains: how can the crowd know where they are, without any interaction on their part, in order to supply good quality WiFi fingerprint data?

IV. CROWDSOURCING THE FINGERPRINT DATABASE

We deploy a temporary localisation system in the environment, gather crowd sourced data, and construct a WiFi fingerprint database using the locations provided by the temporary system. The state-of-the-art RSSI trilateration system, Emender [21], is used with a deployment of Wireless Sensor Network (WSN) anchor nodes for the temporary localisation scheme. The specifics of this approach in context are discussed below. While we use Emender, any localisation system can be used provided that it is: quick to deploy, requires minimal training, and accurate for the purpose. The trilateration approach uses empirical data and a mathematical model with tuned parameters to estimate the distance from the measured RSSI. The WSN is used as we assume no knowledge of the WiFi access deployment.

Figure 2 shows the components of CRAFT. The three phases Setup, Construction, and Query, along with their respective tasks are discussed in the remainder of this section.




Fig. 2: The CRAFT algorithm, where the boxes indicate processes while ovals are for data. The purple tasks are related to trilateration, the green are the fingerprinting stages. The blue tasks on the right are to improve the accuracy of the system. The orange Crowdsourcing task provides data for the fingerprint database. The grey boxes indicate that the element can occur at either of those locations and is evaluated empirically.

By using the crowd to collect data the effort is reduced for a single individual. There are other advantages such as: adaptive resolution, where more frequented areas will have a higher density of samples; reusable anchor nodes, which can be used in other areas or buildings once collection is complete (thus the monetary cost is decreased); and finally, fitting in with the overall aims of this study, ubiquitous technology can be used for localisation with the resulting fingerprint database.

A. Setup

In the setup phase the temporary localisation system is deployed in the building and trained. The phase is broken into three tasks: deployment, zoning, and training. The deployment task is where the anchor nodes are positioned around the environment. CRAFT’s only requirement is that some of the anchor nodes are deployed around the perimeter of the building. The positions are recorded along with the corresponding unique identifier of each node. Note that this task is required by any trilateration system and is not specific to Emender. The next task, zoning, is where the anchor nodes in the building are grouped into ‘zones’ which can be allocated on a per-room basis, or in the case of a large area, to sub-areas as needed. This is discussed further the Construction phase below. The final, training, task is concerned with calibrating the temporary localisation system. Emender, like any trilateration scheme, uses a linear-least-squares method to fit the Log-distance path loss (LDPL) model to the WSN data for each anchor node and transmission power level independently [21]. The amount of data needed to train the system is discussed in Section V.

B. Construction

Collecting data from the crowd for fingerprinting is a straightforward task. Each member of the crowd acquires data from both the temporary localisation system and the WiFi networks detected by their smart phone, or the collection device. The collected WiFi data contains the RSSI, BSSID, SSID, frequency, and time. We use the timestamps to associate WiFi and WSN data.

The bounding boxes are constructed by selecting the maximum and minimum x and y components from the locations of the anchor nodes for each zone. An example of this is shown in Figure 3. Each node is allowed to belong to multiple zones so that the resultant bounding boxes can overlap. Constuction




Fig. 3: Example bounding box computation for the shared postgraduate suite. The nodes 38 and 33 provide the left and right edges, while 37 and 30 provide the top and bottom edges respectively.




Fig. 4: Computed bounding boxes overlaid on the floor plan.

of the bounding boxes in this manner is automatic and requires no explicit information about the floor plan.

To improve the accuracy, the location estimates are restricted using the bounding boxes. A check is performed using the location estimate against each of the bounding boxes. The location estimate is only valid if it is inside at least one bounding box, if the location estimates are not valid they are ignored. This task appears twice in Figure 2 because this check can be performed before construction of the fingerprint database or after the fingerprint database has a location estimate. The evaluation of where, or if, to perform the check appears in Section V.

The fingerprint database is constructed in the traditional manner, as described by RADAR, from the location estimates and fingerprint data collected by the crowd [22], we follow this approach as the emphasis of this paper is to reduce the effort needed to construct the fingerprint database. Fingerprints that share the same location are averaged across the BSSIDs. Once this task is complete, the temporary localisation system can be removed from the area. Queries are then made to find the closest RSSI vector in the fingerprint database. We compute the Euclidean distance between the query and stored fingerprint RSSI vectors, penalising missing values by a fixed amount (in this case 100 dB) larger than the sensitivity of the WiFi radios. This avoids problems with differing dimensionalities of the query and fingerprint RSSI vectors while remaining as authentic to RADAR as possible.

V. EXPERIMENTAL RESULTS

In this set of experiments each component of CRAFT is isolated and the performance is investigated. The first experiment is an assessment of the accuracy of the WiFi fingerprinting scheme with a fingerprint database constructed from the known locations. The subsequent experiments first test the accuracy of the trilateration step and finally, the accuracy of the trilateration constructed database. All these experiments test against the ground truth. The samples at each location, for each access point/anchor node, were collected from both technologies sequentially until the desired number of scans was reached. The WSN was able to respond with multiple transmission power levels so we record the anchor node ID, WSN scan number, transmission power, RSSI and the time of the sample. While for the WiFi we collected the WiFi scan number, RSSI, SSID, BSSID and time of the scan. Unless otherwise stated, the queries comprised the results of a single scan with no further processing performed before estimating the location.

A. Bounding Boxes

The bounding boxes computed from the positions of the anchor nodes quickly and suitably approximate the size and shape of the rooms in our deployment. The differences between the bounding boxes and the actual rooms are firstly, the areas outside the building at the top of the postgraduate suite—the bounding boxes consider them inside the building when they are outside; secondly, where the vertical corridor joins the horizontal—the angled piece is missing; thirdly, there is a minor misalignment in the horizontal corridor; and finally, the mezzanine’s box fails to account for the staircase where experimental data were not collected.

B. Accuracy of WiFi Fingerprinting

The fingerprint database was constructed using the ground truth and the WiFi fingerprints in a 80/20 split between training and test data. Each fingerprint and the resulting location estimate are treated independently, meaning that we do not perform any additional processing on the location estimates even if we know them to be from the same ground truth point. From this experiment, the mean error was 0.36 m with a standard deviation of 1.96 m and maximum error of 28.43 m. All the ground truth locations were contained within one or more bounding boxes so none of the location estimates were ignored.

C. Training of Trilateration

To determine the number of sample points needed for this approach we evaluated an increasing number of randomly chosen sample points from the deployment area. We started by taking ten locations (selected at random) and training the LDPL model, and storing the Least-squares model’s coefficient of determination (R^2). At each location we performed the same 80/20 split between testing and training datasets. This process was repeated 50 times increasing the number of selected locations from ten to the entire dataset (97). The R^2 values plateaued (to approximately 0.6) when 25 random locations were chosen. From this analysis we can see that the number of training locations has been reduced to approximately 25%.

D. Accuracy of Trilateration

The evaluation of the trilateration accuracy isolates the component and evaluates against the ground truth locations. When the fingerprint database is constructed, the WiFi fingerprints collected at the same point are merged. In this experiment, the average RSSI for each BSSID in each ground truth location are used, thus the maximum number of location estimates is 97. The results are presented in Table II. The error decreases from 2.40 m to 2.04 m and the number of valid location estimates decrease from 97 to 71 due to the bounds checks. This is an expected consequence of the trilateration process, and there is a trade-off to be made between higher error and complete coverage, or, lower error with patchy coverage.

E. Accuracy of CRAFT

From table III, we see that the error is *decreasing* between this experiment and the best performing previous trilateration results. There are two cases under evaluation in this experiment: unbounded (where none of the bounds checks are performed), and bounded (where at least one bounds check is performed: post-fingerprinting, pre-fingerprinting, or both). In both the unbounded and bounded cases the mean error increases between the experiments by 0.29 m and 0.70 m and 0.16 m respectively. The 0.29 m mean error increase is broadly consistent with the expected outcomes.

The trade off between lower error with patchy coverage or higher error with complete coverage explains the difference between the 0.70 m and 0.16 m errors. The largest error occurs when the trilateration location estimates are excluded before constructing the fingerprint database but are more accurate, whereas the best performance can be seen when the fingerprint database has full coverage and the bounds are imposed only after the fingerprint step resulting in a mean error of 2.20 m.

Consider the case where the trilateration phase estimates the location outside the bounds and suppose that fingerprint is the closest match in the database. If this point and associated data were to be excluded, the fingerprinting algorithm would search for the next closest match inside the bounded area, and, result in a deviation from the true location thus increasing the total error. If, on the other hand, the point and associated data were to remain in the database then the fingerprint algorithm selects it as the closest match and is then checked against the bounds and only then is eliminated. This does not increase the error but the number of ‘not found’. This is borne out in the Table III where there are 130 troublesome queries, by comparing the number of successful location queries between the second row and any of the others. These 130 queries correspond to the same points eliminated in the previous experiment.

Within this experiment, the error increases from 2.69 m unbounded to 2.74 m pre-fingerprint is a result of the bounds checks because they restrict the fingerprint database as discussed above. The last two rows of Table III are identical for a similar reason—all the location estimates outside the bounds have already been removed in the trilateration stage, thus the bounds check has no effect at the fingerprinting stage.

By constructing the fingerprint database from these estimated locations from trilateration on the temporary network introduces approximately 16 cm to 70 cm of error—broadly

TABLE II: Trilateration localisation results compared with the ground truth.

	Valid	Error (m)			
		Mean	St.Dev.	Median	Max
Post F'print	71	2.04	1.24	1.80	4.96
None	97	2.40	1.43	2.18	7.25

TABLE III: The accuracy of the CRAFT algorithm. The bounds checks are performed after the named stages.

Bounds	Check	Valid	Error (m)			
			Mean	St.Dev.	Median	Max
None		479	2.69	2.34	2.28	30.95
Post F'print		349	2.20	1.62	1.90	16.01
Pre F'print		479	2.74	2.32	2.18	20.62
Both		479	2.74	2.32	2.18	20.62

consistent with the error introduced by the fingerprinting algorithm as seen in Section V-B.

VI. CONCLUSION

This paper presented CRAFT, a solution to the problem of the effort involved with constructing a fingerprint database using crowdsourcing. The approach uses crowdsourcing and leverages the advantages of a temporary anchor node deployment and trilateration to produce location estimates for a fingerprint database. Since this temporary deployment is adaptable and the trilateration is flexible to each unique environment, it is suitable for general use. The thorough and comprehensive evaluation in a real building demonstrates the practicality of the approach. The mean location error is 2.20 m and while this error remains relatively high, any future improvements to RSSI trilateration or alternative easy to deploy systems like Emender are anticipated to improve the performance of this method. While the experiments were conducted with a WSN it is possible to use this approach with other digital radio hardware and thus is suitable for ubiquitous devices.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the support of the University of Otago's Doctoral Publishing Bursary.

REFERENCES

- [1] R. Harle, "A Survey of Indoor Inertial Positioning Systems for Pedestrians." *IEEE Commun. Surveys Tuts.*, vol. 15, no. 3, pp. 1281–1293, 2013.
- [2] S. Goswami, *Indoor Location Technologies*, 1st ed. Springer-Verlag New York, 2013.
- [3] S. Sorour, Y. Lostonlen, S. Valaee, and K. Majeed, "Joint Indoor Localization and Radio Map Construction with Limited Deployment Load," *IEEE Trans. Mobile Comput.*, vol. 14, no. 5, pp. 1031–1043, may 2015.
- [4] A. Rai and K. Chintalapudi, "Zee: zero-effort crowdsourcing for indoor localization," *MobiCom*, 2012.
- [5] C. Wu, Z. Yang, and Y. Liu, "Smartphones Based Crowdsourcing for Indoor Localization," *IEEE Trans. Mobile Comput.*, vol. 14, no. 2, pp. 444–457, feb 2015.
- [6] B. Wang, Q. Chen, L. T. Yang, and H.-C. Chao, "Indoor smartphone localization via fingerprint crowdsourcing: challenges and approaches," *IEEE Wireless Communications*, vol. 23, no. 3, pp. 82–89, jun 2016.
- [7] S.-h. Jung, S. Lee, and D. Han, "A crowdsourcing-based global indoor positioning and navigation system," *Pervasive and Mobile Computing*, 2016.
- [8] M. Werner, *Indoor Location-Based Services*. Cham: Springer International Publishing, 2014.
- [9] I. Bisio and M. Cerruti, "A trainingless wifi fingerprint positioning approach over mobile devices," *IEEE Antennas Wireless Propag. Lett.*, vol. 13, pp. 832–835, 2014.
- [10] D. Milioris, G. Tzagkarakis, A. Papakonstantinou, M. Papadopoulou, and P. Tsakalides, "Low-dimensional signal-strength fingerprint-based positioning in wireless LANs," *Ad Hoc Networks*, vol. 12, pp. 100–114, jan 2014.
- [11] B. Wang, S. Zhou, L. T. Yang, and Y. Mo, "Indoor positioning via subarea fingerprinting and surface fitting with received signal strength," *Pervasive and Mobile Computing*, jun 2015.
- [12] S. Salous, *Radio Propagation Measurement and Channel Sounding*. John Wiley & Sons, 2013.
- [13] M. Ficco, C. Esposito, and A. Napolitano, "Calibrating Indoor Positioning Systems with Low Efforts," *IEEE Trans. Mobile Comput.*, vol. 13, no. 4, pp. 737–751, apr 2014.
- [14] Q. Chang, Q. Li, Z. Shi, W. Chen, and W. Wang, "Scalable indoor localization via mobile crowdsourcing and Gaussian process," *Sensors*, vol. 16, no. 3, p. 381, mar 2016.
- [15] M. Redzic, C. Brennan, and N. O'Connor, "SEAMLOC: Seamless Indoor Localization Based on Reduced Number of Calibration Points," *IEEE Trans. Mobile Comput.*, vol. 13, no. 6, pp. 1326–1337, jun 2014.
- [16] D. Yuanfeng, Y. Dongkai, Y. Huilin, and X. Chundi, "Flexible indoor localization and tracking system based on mobile phone," *JNCA*, vol. 69, pp. 107–116, 2016.
- [17] H. Nurminen, A. Ristimaki, S. Ali-Loytty, and R. Piche, "Particle filter and smoother for indoor localization," in *International Conference on Indoor Positioning and Indoor Navigation*. IEEE, oct 2013, pp. 1–10.
- [18] H. Shin, Y. Chon, Y. Kim, and H. Cha, "A Participatory Service Platform for Indoor Location-Based Services," vol. 14, no. 1, pp. 62–69, jan 2015.
- [19] Y. Kim, H. Shin, Y. Chon, and H. Cha, "Crowdsensing-based Wi-Fi radio map management using a lightweight site survey," *Computer Communications*, vol. 60, pp. 86–96, apr 2015.
- [20] Z. Xiao, H. Wen, A. Markham, and N. Trigoni, "Lightweight map matching for indoor localisation using conditional random fields," in *International Symposium on Information Processing in Sensor Networks*. IEEE, apr 2014, pp. 131–142.
- [21] P. Crane, Z. Huang, and H. Zhang, "Emender : Signal Filter for Trilateration based Indoor Localisation," in *PIMRC*, 2016.
- [22] P. Bahl and V. Padmanabhan, "RADAR: An in-building RF-based user location and tracking system," in *Proc. IEEE INFOCOM'00*, vol. 2. IEEE, 2000, pp. 775 – 784.