A scheme for indoor localization through RF profiling

Israat Tanzeena Haque, Ioannis Nikolaidis, and Pawel Gburzynski
Department of Computing Science
University of Alberta
Edmonton, Alberta, CANADA T6G 2E8
Email: [israat,yannis,pawel]@cs.ualberta.ca

***Abstract*—We present a certain approach to indoor location tracking based on sensing the strength of received RF signal. Our scheme employs a database of signal strength readings from known locations collected during a prerequisite *profiling* stage. Subsequently, the problem of estimating the location of a node emitting an RF signal from an unknown place boils down to data mining in the database—to select a best matched set of profile points and then average the coordinates of those points into an approximate location of the tracked sender. We show preliminary experimental results which confirm our hope that this approach will result in a better accuracy than, e.g., triangulation based schemes.**

I. Introduction

Tracking the locations of objects and people has been the subject of extensive research. While the cases of navigation and (to a large extent) monitoring in open areas have been well addressed by GPS [9], location tracking inside buildings poses specific problems. First of all, GPS doesn't work under a roof. Second, many interesting problems related to indoor location tracking demand a better accuracy than what is available to commercial GPS modules. This accuracy is typically qualified by the characteristics of the monitored area. For example, one would like to be able to authoritatively assess whether the tracked object (or person) is on one side of the wall or the other. The postulated accuracy may be higher in some places than in others. For example, if the distance of the tracked object from the wall (or some other element of the environment considered critical) is reasonably large, the quality of position estimation may be relaxed.

Using RF signals is not the only option for indoor location tracking. Schemes based on infrared signals, acoustic signals, various pressure sensors embedded into the floor, have been proposed and implemented. RF-based technology has the advantage of simplicity and low cost, especially with the proliferation of cheap wireless networks. Our proposed scheme utilizes an ad-hoc network where the cost of a single node (assuming mass production) is a few dollars. The trackable *tags* (being essentially the same nodes as those involved in their tracking) are in principle disposable, which compares favorably with the cost of a (raw) GPS module (which cannot be used indoor, anyway).

RF techniques employed for localization can be categorized with respect to the attribute of the radio signal being measured and correlated with the tag's estimated position. Thus, we have schemes based on measuring 1) Received Signal Strength (RSS), 2) Time Of Arrival (TOA) or Time Difference Of Arrival (TDOA), and 3) Angle Of Arrival (AOA). Among those, the RSS-based techniques [2] pose the minimal requirements on the complexity and cost of the RF hardware. Their commonly acknowledged disadvantage is the poor and unpredictable correlation of RSS with distance resulting from the multi-path effect, which is typically much more serious inside a building than in an open area. Although the other two techniques enjoy a slightly better reputation, especially in not-so-crowded environments, their performance under multiple paths tends to deteriorate quite drastically as well [2]. Thus, in combination with the complexity and cost of the requisite hardware [8], their advantages over RSS-based schemes in indoor environments are questionable.

From the viewpoint of functionality, location tracking techniques can be classified as 1) range-based (i.e., estimating the distance or angle between the transmitter and the receiver [2]), 2) range-free (driven solely by the perceived connectivity of the tracked tag with its neighbors [7]), and 3) profiled (whereby the perceived attributes of the tag's signal are compared against a pre-collected set of samples from known locations [2]). With the range-free approach, the localization problem is easy to solve, but the estimates tend to be crude. By resorting to profiling, one can hope to compensate for the intrinsic properties of the environment which render direct transformations of the perceived characteristics of RF signals into distances or angles highly unreliable. This hope underlies our work.

We present an RF-based localization scheme, dubbed LEMON,[1] which employs a wireless ad-hoc network consisting of cheap nodes deployed at fixed locations in the monitored area. Two features distinguish our approach from previous attempts at profiling-based solutions. First, our objective is not to minimize the number of "base stations," i.e., the receivers monitoring the strength of the RF signal emitted by the tracked object (e.g., as in [2] or [3]). By resorting to a cheap (and massive) ad-hoc network, we can afford to have a dense coverage of the monitored area. Second, we forego the tempting (but practically futile) idea of directly correlating the RSS with distance (e.g., based on a channel model).

[1]Location Estimation by Mining Oversampled Neighborhoods.

978-1-4244-3437-4/09/$25.00 ©2009 IEEE

While some correlation of RSS with location is implicit in our scheme, it is derived solely in an empirical fashion (i.e., determined by the environment). In this context, LEMON can be viewed as a combination of a range-free approach with "traditional" profiling: the scheme is driven by a (somewhat fuzzy) concept of neighborhood, while its objective is to produce an "educated" estimation of the actual location form the coordinates of selected profile readings.

II. Related Work

The most notable, from the viewpoint of LEMON, RF-based indoor location tracking system was proposed in [2]. The idea behind that scheme, dubbed RADAR, is to have a (small) number of *base stations*, each of them covering the entire area to be monitored. During the profiling (off-line) stage, a set of readings from known locations is collected, with each base station recording the received strength of the signal. During location tracking, time-stamped readings from the tracked tag are correlated by the base stations using the *nearest neighbor* search in the RSS space, with the objective of minimizing the Euclidean distance between the measured and recorded signal strength. As that straightforward approach resulted in unacceptable errors, the authors considered other metrics/criteria, including increasing the number of points to average, with mixed results. They also considered a channel model incorporating the impact of walls separating the tag from a base station, which were the source of serious discrepancies in the original estimation algorithm.

From the perspective of our experience, the primary reason for the limited success of RADAR was the small number of base stations which was constrained by the technology (heavy-duty expensive equipment). Within such a framework, the entire monitored area becomes a single large neighborhood, which makes it difficult to isolate and properly profile tricky sub-areas, e.g., in the proximity of walls. The inclusion of tag orientation in the collected profile data demonstrates that the authors tried hard to squeeze as much information as possible into the inherently uncertain profiled sample.

In [10], the authors carried out a comprehensive study of RSS characteristics, with the intention of applying them to location tracking. They focus on three generic attributes: the noise, the attenuation rate, and the effective range. The study considers both indoor and outdoor environments and factors such as elevation, transmission power, packaging, and the impact of obstacles. It concludes that elevation and transmission power are the most influencing factors for RSS which is also highly susceptible to environmental changes. The localization technique proposed in [10] is based on [7], whereby the stationary nodes (called anchors) flood their neighborhoods, while the tracked tags try to assess their shortest path distance from the perceptible anchors. The authors confess that their range-based approach did not work very well in indoor environments and recommend profiling as a more promising idea.

Yet another scheme, presented in [1], employs multiple channels trying to mitigate the multi-path effect. The collected values are subsequently transformed into distance for triangulation. The authors claim that their solution derives a close approximation of the actual distance from the theoretical propagation model. However, we feel that those claims have not been convincingly substantiated be experiments involving elaborate configurations of practical scenarios including obstacles, walls, etc.

In [11], the authors advocate using active RFID (Radio Frequency Identification) tags for indoor localization. The idea is to equip the tracked object with an active RFID tag. The environment is outfitted with a set of n RFID readers and m fixed RFID tags. Two sets of signal levels are collected: r_1 between the tracked object and the readers, and r_2 between the fixed tags and the readers. A similar approach was taken in LANDMARC [6]. The difference consists in the fact that while LANDMARC considers all readers and tags, the scheme described in [11] selects subsets of q readers and p tags that are reachable by all q readers. The location of the tracked object is averaged from the location of those fixed tags whose RSS indications at the readers appear to be close to the indications of the tracked tag.

III. LEMON

Our location tracking system is built around a wireless network of static nodes, called *Pegs*, whose exact locations need not be known. A monitored (mobile) device, which is essentially a node of the same type as a Peg, is called a *Tag*. During the profiling stage, the network collects and stores in a database (maintained on a central server) *samples* acquired from Tags located at known points within the monitored area. The logistics of collecting such samples may involve a person moving around the area with a special variant of the Tag node, e.g., equipped with a clickable map.

A single sample stored in the database can be viewed as a triplet $< C, \Omega, \tau >$, where C stands for the known coordinates of the sampled point, Ω is the so-called association set, and τ, called the sample's *class*, identifies the (settable) RF parameters of the transmitter (typically transmission power, bit rate, and channel number). Its role is to discern samples collected under different "options" of the Tag's transmitter, such that they will only be matched to (future) readings acquired under the same options. For now, we only consider a planar version of the problem, i.e., $C = (x, y)$. The association set Ω consists of pairs $< p, r >$, where p identifies a Peg, and r is the RSS value perceived by that Peg.

A tracked Tag periodically emits numbered RF packets. A Peg receiving such a packet will forward to the central server a report consisting of its own identifier, the Tag identifier, the packet number and class. Having received a collection of such reports referring to the same packet number, the server will built an association set Φ representing the combined momentary perception of the Tag's RSS by all the Pegs that can hear the Tag.

Note that the only difference between the profiling stage and actual tracking is that, in the former case, the association set built by the server upon the receipt of profiling reports is tagged with one more item of information: the known coordinates of the sampled point. To keep as much of the general environment identical in both cases, the Tag uses the same packet format in the two stages. The fields used to convey the coordinates in the profiling stage are filled with random bits during the actual tracking.

The initial step of the localization algorithm is to select from the database a subset of samples representing the best match to the association list Φ. First of all, only the samples of the same class as the received reports are subject to selection. To further narrow down the search, the server finds in Φ the pair $< p_m, r_m >$, such that r_m is the highest among all pairs. Then, it only considers those samples form the database whose association lists include p_m as one of the Pegs. This pre-selection boils down to the postulate that the Peg p_m appearing to be very close to the tracked Tag be a member of all samples that will be used for estimating the Tag's location.

Let $\Omega = \{\omega_1, \ldots, \omega_k\}$ and $\Psi = \{\psi_1, \ldots, \psi_m\}$ be two associations sets. By the distance between these sets, we understand:

$$D(\Omega, \Psi) = \sqrt{\sum_{j=1}^{N}(R_\Omega(j) - R_\Psi(j))^2}$$

where N is the total number of Pegs in the network and $R_\Omega(j)$ is defined as r_j, if the pair $< p_j, r_j >$ occurs in Ω, and 0 otherwise.

In the second step, the server evaluates the distance of each pre-selected sample (its association list) from the current association list Φ. Then, its selects K samples with the smallest distance.

In the last step, the coordinates of the selected samples are averaged to produce the estimated coordinates of the Tag. The averaging formula biases the samples in such a way that the ones with a smaller distance from Φ contribute proportionally more. Let D_{max} be the maximum distance from Φ among the best K selected samples and $S_D = \sum_i^K D_i$ be the sum of all those distances. The tracked coordinates are estimated as:

$$x_{est} = \frac{\sum_{i=1}^{K} x_i \times (D_{max} - D_i)}{K \times D_{max} - S_d}$$
$$y_{est} = \frac{\sum_{i=1}^{K} y_i \times (D_{max} - D_i)}{K \times D_{max} - S_d}$$

where (x_i, y_i) are the coordinates associated with sample i.

Note that the above approach does not attempt to assign any particular interpretation to RSS, except for using it as a numerical attribute of a sample whose value should be close to the observed value. The averaging formula does factor in the magnitude of discrepancy between RSS values (in terms of distance between points in Euclidean space), but this is a purely numerical interpolation (clearly, not an application of any RF propagation model). Our objective is not to minimize the number of Pegs (base stations) or samples (reference points), but instead impose no limit on that number hoping that the matching rules will tend to locate those samples that best apply to a particular "episode" of tracking. As the number of samples is allowed to be large, the role of the last-step interpolation is secondary: we do not purport to know that the RSS values encode some information about the distance. In particular, it may make sense to oversample the area, e.g., collecting multiple samples from the same point. For example, those multiple samples may correspond to the different orientations of the Tag, as in [2], except that in our case there is no need to mark them with any specific "orientation" attribute. If that attribute has any impact at all on the perception of RSS by the neighboring Pegs, that impact will be "discovered" by the system at the matching stage when the closest one of the multiple samples is selected.

IV. Experiments

A prototype LEMON system was implemented and tested in three indoor areas on our university campus. This section describes our experiments and discusses the results.

A. The hardware

Both Tags and Pegs in our LEMON network are based on the same wireless device, EMSPCC11 from Olsonet Communications,[2] which is a low-cost low-power mote for wireless sensor networking, programmable in PicOS [5]. The node employs the CC1100 RF module from Texas Instruments operating within the 916 MHz band. Owing to the flexibility of PicOS and the robustness of TARP as the ad-hoc wireless routing scheme [4], EMSPCC11 is a viable candidate for a production tracking system based on LEMON.

B. The logistics

As the implementation was intended for research rather than production, all data processing was done off-line, i.e., the system wasn't used for real-time location tracking (although there are no fundamental problems barring its true application). The interesting feature of LEMON is that the collected data, including the profiled samples as well as Tag readings from tracked locations, can be re-interpreted many times for solving the same problem. For example, one can try to ignore various subsets of the profiled samples, ignore some Pegs, use different metrics for selecting the best-matched set of samples, change the values of K (the number of best-matched samples), use different averaging formulas, apply various re-scaling factors to RSS readings, introduce thresholds or clustering (discretizing the RSS values), and so on. Consequently, the separation of data acquisition from the actual estimation (localization) is the natural methodology in this type of work.

A typical experiment would start by deploying a number of nodes within the monitored area. Generally, the interpretation of those nodes (Tags versus Pegs) was left until the off-line analysis stage. The networked program run in the nodes (the *praxis* according to PicOS terminology [5]) allowed us to obtain RSS readings between all pairs of directly reachable nodes and for any selected setting of the transmitter (output power,

[2] See http://www.olsonet.com/Documents/emspcc11.pdf.

bit rate, channel number). Thus, the deployment of nodes was usually followed by data collection: the nodes would exchange a massive number of packets, conveying to the central node (called the *master*) their parameters (sender/receiver ID, serial number, transmitter parameters, RSS). The master node was connected (via a USB dongle)[3] to a laptop, where all the data collected by the network were deposited.

During the off-line analysis, some of the collected readings would find their way into LEMON's profile database, some others would be interpreted as tracking data. As the locations of all nodes were known exactly, their roles as Pegs or Tags were exchangeable.

Three rooms were used in our tests: one was a relatively small 3×5m graduate office outfitted with three wooden tables, four chairs, several wooden shelves mounted on the wall, two PC computers, and a steel file cabinet. The other two were of comparable size, over 7×7m, with sparse furniture (mostly chairs). In the third room, the Tag node was moved around by a person: our objective was to determine whether the node's proximity to a body and, e.g., its orientation, have a significant impact on the accuracy.

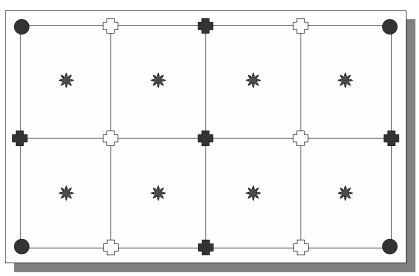

Fig. 1. A sample distribution of nodes in room 1.

Figure 1 shows a sample distribution of nodes for an experiment carried out in room 1 (3×5m). All marks represent nodes (their total number is 23). In one interpretation, the 4 solid circles (in the corners) were Pegs, while the 11 crosses (both transparent and solid) provided profile samples. The asterisks acted as Tags whose locations were to be determined.

The RF module of EMSPCC11 offers several settings regarding the transmitted power level, packet bit rate, and the channel number. The transmission power can vary from -30 dBm to 10 dBm (in 8 discrete steps), the bit rate options are 5 kbps, 10 kbps, 38 kpbs, and 200 kbps, and there are 256 different channels (numbered 0 to 255) with 200 kHz spacing. All combinations are possible and, in principle, sensible. The experiments reported on in the rest of this paper were carried out at the lowest power setting with 5 kps transmission rate using channel 0.

C. The results

Consider the configuration of nodes shown in Figure 1. The four corner nodes are used as Pegs, while the nodes marked with asterisks are interpreted as Tags whose locations are to be estimated. In the first crude experiment, we only use five profiled samples provided by the nodes marked with solid crosses. With $K = 4$ (i.e., four best matching samples out of the five), the average error in estimating the location of the eight Tags was 0.73 m. With 11 reference points (all crosses) and $K = 5$, the average error was 0.6 m. Other values of K, i.e., $K = 4$ or $K > 5$ gave worse results, as illustrated in Figure 2.

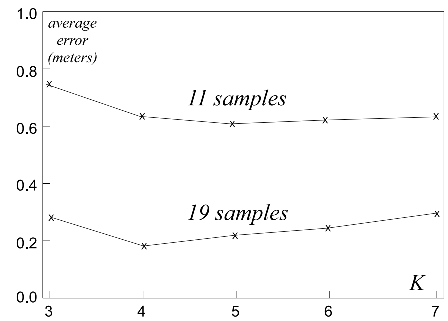

Fig. 2. Average error distance, room 1.

The set of samples can be augmented by including some readings from the nodes representing Tags, which will bring the total number of samples to 19. With those readings included in the database, the quality of location estimation improves rather drastically (see Figure 2). This time the best number of samples turned out to be four. Owing to the fact that the database contains samples collected from points situated very close (distance zero) from the estimated locations, using too many points in the averaging formula will have the tendency to pollute the contribution of the true best match, even if that match isn't 100% perfect, because of statistical fluctuations in RSS. Thus, this simple exercise suggests that dense profiling is likely to improve accuracy and require a lower value of K than a sparse set of samples.

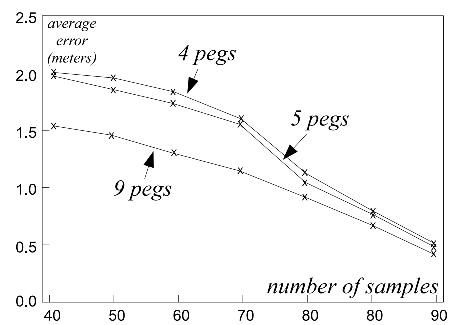

Fig. 3. Experiment in room 2.

In room 2, we set up an 8×8 grid of nodes spaced 1 m apart. Initially, we assumed only 4 pegs located in the corners. The remaining 60 nodes of the grid provided signals for profiled samples. Then we put some nodes in the centers of the grid

[3]EMSPCC11 is equipped with a raw UART interface which can be easily converted to USB.

squares to collect data for estimating their locations. The average error turned out to be around 1.6, with a not-so-well pronounced minimum ($1.56\,\mathrm{m}$) for $K = 6$.

There are two obvious ways to try to reduce that error. The fact that it is significantly higher than in the previous case clearly results from the much larger area being covered by the same (small) number of Pegs. Consequently, increasing the number of pegs is one obvious idea, while another approach is to increase the number of profiled samples. Figure 3 shows the impact of these two factors on the observed error. Suppose that the grid points are denoted $< u, v >$, where $u, v = 0, \ldots, 7$. Thus, the original four Pegs were placed at four corners of the grid and the case with 5 Pegs includes one extra Peg at location $< 3, 3 >$. Finally, the additional four pegs for the 9-Peg scenario are $< 4, 0 >$, $< 7, 4 >$, $< 3, 7 >$, $< 0, 3 >$. The extra samples were collected in such a way as to make their distribution as uniform as possible within the confines of the discrete grid.

The experiment conducted in room 3 involved the same grid as in room 2, except that the tracked Tag was carried by a person (both for profiling as well as the actual tracking). Initially, the results were perceptibly worse than those obtained under the more sterile conditions of the previous two environments. The analysis of data revealed that most of the problems resulted from assigning too much relevance to low RSS values, i.e., corresponding to weak reception, which would exhibit large statistical fluctuations. The numerical RSS readings presented by the RF module of EMSPCC11 are positive numbers, roughly between 80 and 150, representing a shifted dB signal level of the received packet. Assuming that those readings are always between *MIN* and *MAX*, we transformed them by applying this formula:

$$R_s = \left(\frac{r - \mathit{MIN}}{\mathit{MAX} - \mathit{MIN}} \right)^{\alpha} \quad ,$$

where $\alpha > 1$. Note that for $\alpha = 1$ we effectively obtain the original (not re-scaled) case, as a linear transformation of all RSS readings does not change the outcome of our algorithm. The best results have been observed for $\alpha = 3$. In 90% cases we were able to estimate the Tag location with an error less than $1\,\mathrm{m}$. The average error went down to about $0.5\,\mathrm{m}$ with 76 profile samples.

V. Conclusions and further work

We believe that our system presented in this paper can be effectively applied to tracking equipment as well as people in indoor environments with the average error well below $1\,\mathrm{m}$. Even in its present form (far from complete), the scheme practically guarantees this kind of accuracy, and we haven't even scratched the surface of all the options that remain to be explored. LEMON appears simpler and more accurate than other approaches, especially that the limits of its accuracy have not been explored yet. For example, the average error reported in [6] was about $2\,\mathrm{m}$ using a finer grid than in our case. Finally, the uniformity and low cost of the equipment (even our underfunded lab could afford a few hundred of Olsonet nodes) makes LEMON a highly viable and very practical solution.

One of our important goals for immediate future is to investigate how reliably LEMON can decide on which side of a wall the tracked Tag is located. In most scenarios where indoor location tracking is of practical relevance, decisions of this kind play a significant (often critical) role. Also, it seems that the proper way of generalizing the scheme to 3-D is to tackle the "wall problem" (or the "ceiling problem") first. This way, once the floor has been determined correctly, the remaining part of tracking can be safely carried out in the 2-D domain.

We haven't yet used the *class* attribute of a sample (see Section III), i.e., all samples used in our experiments so far were of the same class. The observation that lower values of RSS tend to be confusing to the sample selection algorithm as well as the averaging formula hints at the possibility of using two or more transmission power levels at the Tag. This is an easy thing to implement, as the Pegs need not re-tune themselves: the information about the power level is included in the packet. A strong signal will confuse nearby Pegs (their readings will not be reliable), but some further Pegs may then offer assistance in interpreting the reading. This effectively calls for re-scaling the RSS readings at both ends of the scale and using multiple matches (within different classes of samples) to arrive at a better approximation of the Tag's location.

References

[1] S. Ali and P. Nobles. A novel indoor location sensing mechanism for IEEE 802.11 b/g wireless LAN. In *4th Workshop on Positioning, Navigation and Communication, WPNC '07*, pages 9–15, March 2007.

[2] P. Bahl and V.N. Padmanabhan. RADAR: an in-building RF-based user location and tracking system. In *Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. (INFOCOM 2000)*, Israel, March 2000.

[3] T.W. Christ, P.A. Godwin, and R.E. Lavigne. A prison guard duress alarm location system. In *International Carnahan Conference on Security Technology*, pages 106–116, Ottawa, Canada, October 1993.

[4] P. Gburzyński, B. Kaminska, and W. Olesinski. A tiny and efficient wireless ad-hoc protocol for low-cost sensor networks. In *Proceedings of Design Automation and Test in Europe (DATE'07)*, pages 1562–1567, Nice, France, April 2007.

[5] P. Gburzynski and W. Olesinski. On a practical approach to low-cost ad hoc wireless networking. *Journal of Telecommunications and Information Technology*, (1):29–42, January 2008.

[6] L.M. Ni, Y. Liu, Y.C. Lau, and A.P. Patil. LANDMARC: indoor location sensing using active RFID. *Wirel. Netw.*, 10(6):701–710, 2004.

[7] D. Niculescu and B. Nath. Ad hoc positioning system (APS). In *International Conference on Global Communication (GLOBECOM)*, pages 2926–2931, Ottawa, Canada, 2001.

[8] D. Niculescu and B. Nath. VOR base stations for indoor 802.11 positioning. In *MobiCom '04: Proceedings of the 10th Annual International Conference on Mobile computing and networking*, pages 58–69, New York, NY, USA, 2004. ACM.

[9] United States Coast Guard Navigation Center. Global positioning system standard positioning service specification, June 1995.

[10] K Whitehouse, C Karlof, and D Culler. A practical evaluation of radio signal strength for ranging-based localization. *SIGMOBILE Mob. Comput. Commun. Rev.*, 11(1):41–52, 2007.

[11] Guang yao Jin, Xiao yi Lu, and Myong-Soon Park. An indoor localization mechanism using active RFID tag. *Sensor Networks, Ubiquitous, and Trustworthy Computing, International Conference on*, 1:40–43, 2006.