

Embracing Localization Inaccuracy: A Case Study

Usman Raza, Amy L. Murphy, Gian Pietro Picco

November 2012

Technical Report # DISI-12-038

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible

Embracing Localization Inaccuracy: A Case Study

Usman Raza^{†‡}, Amy L. Murphy[†], Gian Pietro Picco[‡]

[†] *Center for Information Technology - IRST, Bruno Kessler Foundation, Trento, Italy*
{raza, murphy}@fbk.eu

[‡] *Department of Information Engineering and Computer Science, University of Trento, Trento, Italy*
gianpietro.picco@unitn.it

Abstract—In recent years, indoor localization has become a hot research topic with some sophisticated solutions reaching accuracy on the order of ten centimeters. While certain classes of applications can justify the corresponding costs that come with these solutions, a wealth of applications have requirements that can be met at much lower cost by accepting lower accuracy. This paper explores one specific application for monitoring patients in a nursing home, showing that sufficient accuracy can be achieved with a carefully designed deployment of low-cost wireless sensor network nodes in combination with a simple RSSI-based localization technique. Notably our solution uses a single radio sample per period, a number that is much lower than similar approaches. This greatly eases the power burden of the nodes, resulting in a significant lifetime increase. This paper evaluates a concrete deployment from summer 2012 composed of fixed anchor nodes throughout one floor of a nursing home and mobile units carried by patients. We show how two localization algorithms perform and demonstrate a clear improvement by following a set of simple guidelines to tune the anchor node placement. We show both quantitatively and qualitatively that the results meet the functional and non-functional system requirements.

I. INTRODUCTION

Indoor localization has been a much-researched topic with increasingly complex solutions such as those based on ultra wide band achieving accuracy on the order of 10 cm. Such solutions often come with a high cost both to establish the required infrastructure for and the maintenance of that infrastructure. Recently, several techniques have been developed to exploit low-power wireless sensor networks (WSNs), specifically using RSSI. However, due to the high, inherent variability of RSSI, these approaches tend to take many samples in a small time window, fusing the results in sophisticated ways, and even so, the results have large errors. Frequent sampling not only saturates the communication network, it also drains the batteries, increasing maintenance costs.

While some applications demand high accuracy, a wealth of others can accept lower accuracy and still meet the application needs. In this paper we explore one such application in a nursing home, detailed in Section II. In brief, our system offers both real-time information to caregivers about patient location as well as summary data for offline evaluation for changes in movement patterns. Neither of these objectives requires a high degree of accuracy, therefore our goal was to identify a solution with sufficient accuracy at low cost.

Our solution relies on standard WSN nodes, some anchored to known locations in the environment, and some carried by the patients. We use our own energy-efficient contact detection

protocol to identify when a mote is in range of one or more anchors. The actual location is calculated using two off-the-shelf RSSI-based techniques described in Section III. We then quantitatively and qualitatively evaluate all these solutions in Section IV, showing that their accuracy is sufficient for our application. The qualitative perspective is missing from most relevant literature, but it is critical to a complete analysis of the techniques that relates accuracy and user acceptability. We also show that by properly tuning the anchor node placement according to a simple set of guidelines, detailed in Section V, significant accuracy gains are possible at the same cost.

We end the paper with an overview of related work in Section VI and brief concluding remarks.

II. THE APPLICATION SCENARIO: ACUBE

The study presented in this paper is part of ACUBE (Ambient Aware Assistance), a four-year locally-funded project. The goal of ACUBE is to provide a technological infrastructure for improving the quality of life for elderly and disabled persons. The system combines a number of core sensing technologies such as video, audio, and WSN to offer support to caregivers.

While the system has been deployed in several test facilities, we focus here on the summer 2012 installation at a nursing home. The goal of this system is to support the staff in monitoring and evaluating the patients. For example, inside a patient's room, the video subsystem is used to detect repetitive movements as is common for Alzheimer's patients. Our focus here is on the WSN, which has two primary tasks. First, it is used to indicate the patient's approximate location on the floor. This is important as the common areas, shown in Figure 1, are quite large and patients are not always in view of a caregiver. A related functionality is an immediate warning if a patient leaves the monitored area, e.g., exiting into a stairwell. Second, the caregivers are provided an offline summary of patient movements, reporting the time a patient spent in each area.

A. ACUBE System Architecture

The overall ACUBE system is designed to decouple the sensing subsystems from the applications that exploit the produced data. This is accomplished with an Apache Active-MQ message queue middleware [1] to which applications register to receive data published by the various sensor subsystems. For example, a GUI application visualizes the patient locations generated by the WSN subsystem. A separate application collects the locations for offline analysis.

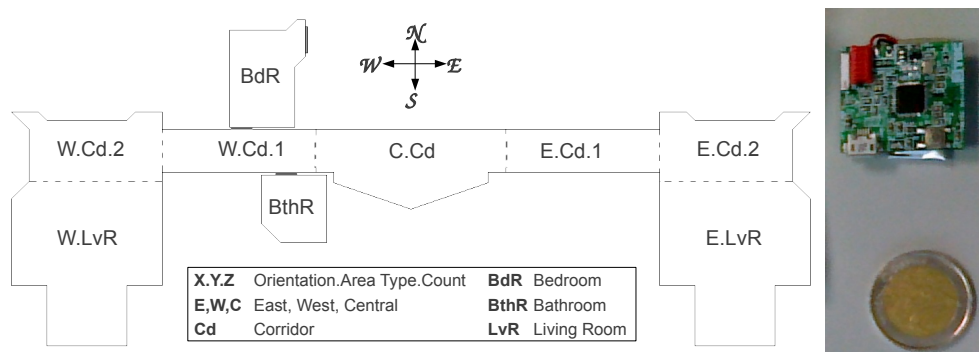


Fig. 1. Monitored areas of a single 45x20 m floor in a nursing home. The total monitored area is 307 m². Dashed lines indicate virtual boundaries between areas, while solid lines indicate walls or doors. The photo shows a 2 euro coin and the mobile mote carried by the patients.

B. WSN Subsystem

The WSN subsystem is composed of approximately 20 anchor nodes attached to the walls at a height of 2 m. Each patient is also given a mobile mote fastened to their clothing, e.g., pinned inside a shirt pocket. While the anchors are standard TelosB motes, the mobile motes are a custom modification that reduces the physical footprint to approximately 2 cm², as shown in Figure 1. Both anchors and mobile motes are battery powered and run TinyOS. The lifetime of the anchor motes running with two standard AA batteries is approximately 45 days while the mobile motes, with a rechargeable lithium ion LR2450 button battery, have a 5-day lifetime. The local processing and communication protocols running on all motes are described in Section III.

To collect data from the motes, we deployed two sink motes connected via USB to Gumstix embedded PCs. These PCs were connected via wired Ethernet to a PC running the so-called WSN adapter, which elaborated the raw data from the motes and then published events on the message queue. In this paper we focus on the localization events, reporting the (x, y) coordinates of the mobile motes.

III. CONTACT DETECTION AND LOCALIZATION

As noted previously, the primary goal of the WSN is to approximately track the patient, specifying the area where they are located. Most RSSI-based localization techniques use the signal strength between the mobile and multiple anchors to periodically offer the location of the mobile units. Further, common instantiations of these localization techniques collect signal strength many times in each period, ensuring that sufficient measurements are taken to balance the inherent inaccuracies of using RSSI to approximate location. Our instantiation separates the collection of RSSI measurements into a *contact detection* module that detects when the mobile is within range of one or more anchors, and reports the signal strength of the detected contact. Notably, our contact detection is based on a single message from the mobile to the anchor in each period. This greatly reduces the energy burden on the motes, however it simultaneously limits the accuracy of the RSSI-based localization technique.

This section outlines our contact detection mechanism, then briefly describes the two RSSI-based localization techniques we adopt. The next section returns to the accuracy analysis.

A. Contact Detection

Our localization approach takes input from the contact detection performed by the mobile and anchor elements of the WSN, thus it is important to consider the maintenance costs of this network, in particular the node lifetime. The anchor nodes are powered by two AA batteries, and changing these batteries is a labor intensive task. Instead, the mobile motes have rechargeable batteries, and it is straightforward to dedicate two motes per patient, one worn and one charging. Therefore, our contact detection protocol places a heavier power burden on the mobile units than on the anchor nodes.

As noted, the localization algorithm must report a new location periodically, with this period corresponding to the refresh frequency of the graphical interface presented to the caregivers. This period is a tunable parameter, set to 5 s in our experiments. We determined this value experimentally, as the nursing home patients are not fast-moving, and their location will not change dramatically in 5 s.

This period is important to the functionality of the contact detection algorithm, as we must guarantee that all anchors in range of the mobile unit detect the contact once in each interval. To do this in an energy-efficient fashion, the anchor nodes put their radios to sleep for half of the interval, then operate in a Low Power Listening mode (specifically BoX-MAC [2] with a sleep interval of 50 ms) for the remainder of the period. In contrast, the mobile motes transmit a broadcast beacon packet at -15 dBm once every half period, for a total of two transmissions per period. This combination ensures that the beacon will be received by all anchors in range, thus triggering contacts between the mobile and all nearby anchors.

The RSSI of all detected contacts are forwarded from the anchor nodes to the sinks and then to the WSN adapter where the location is calculated once per period according to one of the following two localization approaches.

B. Localization with Maximum RSSI

Our first localization approach was designed to be simple, and to adhere to the goal of merely approximating location.

Recall that the algorithm must work in a periodic manner, offering a location once per period. Therefore, our first technique collects the RSSI measurements of all contacts in each period, then identifies the anchor mote with the highest RSSI value. The assumption is that the mote is closest to this anchor.

C. Localization with REWL

Our second approach sought to offer a more refined location, but without modifying the functionality of the contact detection system. Therefore we adopted from the literature a free-range localization scheme called Relative Span Exponential Weighted Localization, REWL [3]. Notably, REWL has an extremely low computational overhead and our Java implementation requires only 6 lines of code. As with Max-RSSI above, REWL collects all RSSI values in a given period, then offers an (x, y) location by combining these values with the known locations of the anchor motes.

Intuitively, one can visualize the functionality of REWL by placing the mobile unit at the central point of a star graph where each edge is between the mobile and an anchor with which a contact was detected in the period. The RSSI values reported by each anchor specify edge weights that affect the length of the edge, like a spring. Stronger RSSI signals shorten the edge while weaker signals allow for longer edges. When the springs reach equilibrium, the coordinates of the mobile at the center of the weighted springs is reported as the location.

Notably, it is possible for the reported location to be outside of the monitored area. For example, if the patient is standing in the central corridor (C.Cd), weak contacts may be detected with anchors in the bathroom (BthR). The resulting localization, therefore, may place the patient in the “empty space” between the corridor and the bathroom. Although this occurred infrequently, when it did, we placed the mote on the edge of monitored area that contained the mote with the strongest RSSI contact in the given period.

IV. EVALUATION

To evaluate the combination of our power-efficient contact detection algorithm with the two localization schemes, we placed several anchor motes in the nursing home facility available to the ACUBE project. We then ran a number of test scenarios with researchers standing in for patients. Unfortunately due to legal privacy concerns, we were unable to run experiments with actual patients. Nevertheless, the experiments were executed in a controlled way, allowing for repeatability and complete coverage.

The test scenarios and their objectives are outlined next, followed by both quantitative and qualitative assessments.

A. Deployments

Over the course of our experiments in the summer of 2012, we explored two different physical deployments of the nodes; the second being a refinement of the first, designed to improve the accuracy without increasing system cost by adding new nodes or changing the energy consumption of the motes.

In both deployments, we applied TRIDENT [4], a connectivity assessment tool, to ensure that at least one anchor was visible from all possible patient positions. Thus if the patient remains in the monitored area, their location can be estimated.

Our first deployment spread the motes throughout the floor as shown in the left of Figure 2. All anchors are placed on walls or building support columns in the middle of the rooms.

Based on the collected results, described next, we chose to move the anchors for the second deployment to increase the quantitative accuracy. To achieve this, we applied a simple heuristic: maximizing the distance between anchors placed in adjacent areas. The result is shown on the right of Figure 2. Notably, we removed one node from an area whose coverage was too dense and moved to the area interiors most anchors that were formerly placed on the borders between two areas. This deployment evenly spreads anchors throughout the monitored area. This is especially important for REWL, as uneven deployment could incorrectly “pull” the location towards an area with more anchors, even if those anchors report weak signals. Additional details concerning guidelines for node placement appear in Section V.

B. Scenarios

Conversations with the caregivers revealed that certain areas are more critical than others. For example, as the patients spend most of their time in either the living rooms or their bedroom, it is more important that locations are accurate for these areas. Similarly, the central corridor (C.Cd) is also the location of the elevators and primary stairwell. It is important that if a person enters this corridor, they are detected so that if a patient is subsequently not detected, the caregivers will know from which area they exited.

To evaluate the behavior of the system, we ran three scenarios with a single mobile mote. In all scenarios, the ground truth was also collected.

- **Area experiments** are designed to evaluate the detection performance in individual areas/rooms. Specifically, the mobile user moves randomly in each single area for a duration of 5 minutes. As described earlier, low accuracies recorded for these experiments in our first deployment led us to move multiple anchors.
- **Walk experiments** are designed to assess system accuracy in tracking a mobile node as it moves around the monitored floor. The mobile user walks across the monitored area on the path shown in Figure 3, ensuring that the user traverses each area twice.
- **Patient emulation tests** focus on the prospective of patients and caregivers, the end users of the system, and are designed specifically for qualitative evaluation. Here, our test subject tries to realistically emulate patient behavior including sitting periods and walking pace. Simultaneously one of the caregivers qualitatively evaluates the conformance of the tracking visualization.

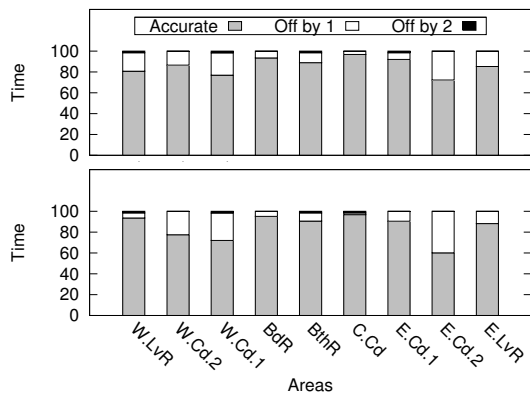


Fig. 4. Area accuracy for the second deployment using Max-RSSI (top) and REWL (bottom).

also an alternate application that sounds an alarm when the patient leaves the area. By knowing generally where the patient was last known to be, the number of exits to check is greatly reduced. Thus a measurement that is off-by-1 is sufficient.

Walk Experiments. To evaluate the system when the patient is moving, we conduct four repetitions of the walk scenario, collecting ground truth as before. The localization accuracy shown in Figure 5 indicates that REWL outperforms Max-RSSI. This is due to the fact that REWL bases the location on beacon receptions by *all* in-range anchors, making it less prone to movement dynamics and body shielding between the mobile node and anchors.

For space reasons we do not present the analogous plots for the first deployment, but remark that the average accuracy increased from 70% to 85% for REWL and from 67% to 82% for Max-RSSI.

In addition to the inherent uncertainty of RSSI measurements, our contact detection protocol also influences the accuracy during movement. Recall that the mobile mote sends two beacons in every period, and each anchor in range receive only one of these beacons as anchors sleep for half of the period. This means that, in any given period, on average half of the anchors receive the first beacon while the other half receive the second. Nevertheless, the RSSI values based on both beacons are integrated to offer the location for that period. Thus, if the patient moves a significant distance in 2.5 s, the location could be negatively affected by this integration. Nevertheless, given the average walking speed of the patients in a nursing home in combination with the expected use of the location information, this is unlikely to cause any problems.

D. Qualitative Evaluation

While the previous section presents quantitative accuracy measurements that are acceptable given our low-power, low-cost solution, numbers cannot assess whether the system provides a reasonable service to the caregivers. To evaluate this, we enlisted a caregiver to observe the ACUBE visualization during patient emulation tests. At the end of each experiment, the operator was asked to rate how well the visualization




Fig. 5. Accuracy of walk experiments for the second deployment using Max-RSSI (left) and REWL (right).

represented the actual movement of the patient and whether the information provided on the GUI would be useful for finding a patient. Based on these experiments, we extracted the following three observations.

Qualitative evaluation depends heavily on the visualization.

The visualization offered to the caregiver essentially shows a patient icon placed on a floor map similar to those shown in this paper. For Max-RSSI, we evaluated two different visualizations, namely placing the patient icon at the center of the room or co-locating the icon with the anchor that reported the maximum RSSI value.

Interestingly, despite the fact that the underlying data is identical, the caregiver gave a much more favorable rating to the less precise visualization that placed the node at the center of the area. Consider that the GUI is updated every 5 s. If the user stays within a single area, the icon remains at the center of the room, even if different motes in that area report the maximum RSSI. Instead, in the latter visualization, the icon frequently jumps from one anchor location to another every 5 s, producing a visually disruptive view. In the end, the caregivers preferred the visualization that contained less specific location information.

REWL is preferred over Max-RSSI. The third visualization offered to the caregiver placed the patient icon at the (x, y) location indicated by REWL. Although the icon changed position on the GUI after every 5 s period, this visualization was much preferred over the Max-RSSI visualization at the center of the area. Consider that when the patient moves from room to room, REWL produces subsequent (x, y) locations that are relatively close together. Instead, with the other approach, the icon moves from the center of one area to the center of another. This is typically a much larger distance that does not represent a realistic movement. Thus it is more visually disruptive.

Additionally, for locating a patient, if the REWL (x, y) location places the patient on the border between two areas, even if it is in the wrong area, it is much less of a problem than either of the other approaches which place the icon well inside the wrong area. Further, with REWL, the caregiver receives an indication of the direction of the patient movement, which is a further aide in finding the patient.

Low accuracy does not yield an unacceptable solution. Here we reported qualitative results only for the second deployment. Nevertheless, it is worth mentioning that preliminary experi-

ments with the first deployment, especially with the REWL visualization, also received positive qualitative assessments from our fellow researchers participating in the testing. This is despite the fact that the quantitative accuracy was, for some areas, very low. Instead, it was our own dissatisfaction with the accuracy of the first deployment that led us to move the nodes and fully evaluate a system with better quantitative accuracy.

V. LESSONS LEARNED

As seen in the previous section, we improved the accuracy by slightly modifying anchor node placement. Here we summarize the guidelines we followed to aide in future deployments: *i)* Maximize the distance between anchor nodes deployed in two adjacent monitored areas. *ii)* Place anchor nodes toward the center of the monitored areas. *iii)* Do not unevenly place more nodes in similarly sized areas. *iv)* Exploit physical barriers in the environment to provide radio shielding. To illustrate the final point, we note that anchor 19 was placed on a vertical column in the middle of the room. For the second deployment, we moved it to the opposite side of the column, allowing it to shield the anchor from motes in corridor E.Cd.2.

Looking back at our experience, we summarize the key lessons: *i)* The qualitative results are much better than the quantitative measurements would lead one to expect, implying that our application has a high tolerance for low accuracy. *ii)* Anchor placement can be refined to match the environment as well as the localization algorithm by following a few simple guidelines. *iii)* Adjacent areas without physical barriers such as walls decrease accuracy. Further, we note that accuracy could be further improved by processing the resulting stream of locations to remove transient locations that do not make logical sense, e.g., a mobile unit that moves too far in the period or transiently moves off the expected path. Additionally, improvements in the GUI could be made to smooth the motion of the patient icon, perhaps further increasing the positive evaluation by the caregivers.

VI. RELATED WORK

A plethora of localization techniques have been proposed that are broadly classified in two categories: range-based and range-free [5]. Range-based localization estimates the distance or angle between two nodes to calculate their relative location by employing Time of Arrival (TOA), Time Difference of Arrival (TDOA) and Angle of Arrival (AOA) or using RSSI to calculate the distance. TOA, TDOA and AOA require specialized hardware such as redundant transceivers, high precision clocks and ultra-wide band technology to achieve high localization accuracy. However, the accompanying increase in the cost and power consumption make them a less than ideal for large-scale WSNs of battery-powered nodes. It is also been observed that distance estimation with RSSI is very inaccurate in practice even for “ideal” outdoor environments [6].

For our target scenario, range-free techniques, such as centroid-based algorithms [7], [3], offer cost-effective solutions for low-power scalable localization system as they use standard radio technology. Although range-free localization

is simple, reportedly it does not provide good accuracy [8]. To mitigate this accuracy problem, additional mechanisms for redundancy and aggregation [9] have been researched, however they are computationally complex and require additional network traffic. These mechanisms must be evaluated in each application scenario to weigh the additional cost versus the increase in accuracy. Instead, most studies of these techniques have been conducted in simulations or on testbeds without any reference application. One quantitative study in [8] evaluated four different RSSI-based localization techniques in an indoor environment, with the conclusion that RSSI is unsuitable for accurate localization. Nevertheless, our work clearly shows that RSSI can produce acceptable results for end users, therefore the quality of RSSI-based localization is suitable for at least some real-world applications.

VII. CONCLUSION

Despite the widespread belief that RSSI-based localization techniques yield poor results, the qualitative evaluation presented here clearly shows that for a real scenario in a nursing home high accuracy is not required. Instead, our low cost, low maintenance system formed of standard WSN components and simple localization techniques is more than acceptable for this scenario and for many others.

ACKNOWLEDGMENTS

The authors wish to thank Davide Molteni, Silvia Demetri and Ramona Marfievici for their help during the deployment. This work was funded by ACUBE, a 2008-2012 Province of Trentino *Grandi Progetti*.

REFERENCES

- [1] (2012) Apache activemq. [Online]. Available: <http://activemq.apache.org/>
- [2] D. Moss and P. Levis, “BoX-MACs: Exploiting physical and link layer boundaries in low-power networking,” Stanford University, Tech. Rep. SING-08-00, 2008.
- [3] C. Laurendeau and M. Barbeau, “Relative span weighted localization of uncooperative nodes in wireless networks,” *Wireless Algorithms, Systems, and Applications*, pp. 358–367, 2009.
- [4] M. Chini, M. Ceriotti, R. Marfievici, A. Murphy, and G. Picco, “TRIDENT, untethered observation of physical communication made to share,” in *Proc. of the 9th ACM Conf. on Embedded Networked Sensor Systems*. ACM, 2011, pp. 409–410.
- [5] M. Bal, M. Liu, W. Shen, and H. Ghenniwa, “Localization in cooperative wireless sensor networks: A review,” in *13th Int. Conf. on Computer Supported Cooperative Work in Design (CSCWD)*. IEEE, 2009, pp. 438–443.
- [6] K. Whitehouse, C. Karlof, and D. Culler, “A practical evaluation of radio signal strength for ranging-based localization,” *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 11, no. 1, pp. 41–52, 2007.
- [7] N. Bulusu, J. Heidemann, and D. Estrin, “GPS-less low-cost outdoor localization for very small devices,” *Personal Communications, IEEE*, vol. 7, no. 5, pp. 28–34, 2000.
- [8] G. Zanca, F. Zorzi, A. Zanella, and M. Zorzi, “Experimental comparison of rssi-based localization algorithms for indoor wireless sensor networks,” in *Proc. of the Wkshp. on Real-world Wireless Sensor Networks*. ACM, 2008, pp. 1–5.
- [9] T. He, C. Huang, B. Blum, J. Stankovic, and T. Abdelzaher, “Range-free localization schemes for large scale sensor networks,” in *Proc. of the 9th Int. Conf. on Mobile computing and Networking*. ACM, 2003, pp. 81–95.