

Robust Indoor Localization on a Commercial Smart-Phone

Nisarg Kothari Balajee Kannan M. Bernardine Dias

CMU-RI-TR-11-27

August, 2011

The Robotics Institute
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

Copyright © 2010 Carnegie-Mellon University

Abstract—This technical report outlines a system for robust indoor localization on a commercial smart-phone. Towards achieving effective localization, in the absence of GPS, we combine complementary localization algorithms of dead reckoning and Wifi signal strength fingerprinting. Dead reckoning is performed using the on-board accelerometer, magnetometer, and gyroscope sensors to detect motion and estimate orientation. At the same time, Wifi signal strength fingerprinting is employed to provide an independent position estimate. These measurements along with a pre-built map of the environment are combined using a particle filter towards more accurate pose estimation. We discuss a procedure for collecting Wifi calibration data which uses a robot to reduce the amount of time needed to train the system for a given environment. The system was tested using multiple participants in two different indoor environments. It was found to achieve localization accuracies of within 10 meters, sufficient for a variety of navigation and context-aware applications.

I. INTRODUCTION

In recent years, smart-phones have re-defined the notion of mobile computing platforms. Ever improving features of affordability, ubiquity, and portability, increased sensory and computational power along with low power consumption fueled by readily available batteries, have opened up a number of interesting applications. One such application is software that is *location-aware*. Location aware systems in indoor environments have the promise of significantly improving the quality of life, especially for the visually impaired. By tracking a user's location and providing directions audibly in terms of easily recognized features like doorways and intersections, a location aware smart-phone application could improve the independence and effective mobility of the visually impaired.

Most current location-aware applications use GPS tracking technology to provide directions and locate places of interest. However, these systems suffer from their dependency on GPS and can therefore only function accurately when there is a clear view of the requisite number of satellites. In the absence of GPS, achieving sufficient localization accuracy on a consumer device is extremely challenging. By using a variety of onboard sensors and exploiting complementarity and redundancy, GPS-free localization accuracy in the range of ~10 meters can be achieved even in conditions where one or more sensors are disrupted. We propose one such application for robust indoor GPS-free localization using commercial mobile phones. The core research idea involves combining complementary localization approaches for dead-reckoning and WiFi signal strength strategies towards robust GPS-free localization. The complementary and redundant characteristics of the two approaches allow the system to robustly operate even in environments where one or more individual sensors maybe disrupted.

The suitability of a localization method depends on a number of factors, including, the size and cost of the hardware, the time and money needed to deploy the system in a new environment, and the accuracy that can be achieved. The system introduced here operates on a consumer smart-phone of the sort people already own. By exploiting Wifi access point infrastructure that in many cases is already installed in a building, this system minimizes installation time and cost. The system utilizes the measurements of wireless radio signal strength indication (RSSI) for position estimation. A fingerprinting technique is used to identify RSSI position dependency. The system does require a calibration step for Wifi signal strengths. Using a robot, mapping a floor of a roughly 10,000 square foot building takes about an hour. Further, once a calibration database of the environment is generated, it can then be used across different phones without the need for re-calibration. Accuracy of about 10m is generally achievable, sufficient for determining the room a user is in or their approximate location along a hallway.

Rest of the paper is organized as follow. Related work is discussed in Section II, followed by a discussion of the localization system in Section III. The performance of the system in different environments is evaluated in Section IV, followed by conclusions and future work.

II. RELATED WORK

There is an extensive literature in pedestrian localization using custom hardware and rigid mounts. Specialty shoes [1], helmets [2], and wearable modules [3] fitted with sensors have all been shown to be effective at the localization task. These solutions have potential applications in specialized areas such as disaster response or the military, but are unlikely to have the same success in consumer applications due to the expense imposed by having to acquire specialty hardware and the inconvenience of setting up the solutions. A localization solution implementable on a phone is much more likely to obtain broad acceptance. A cell phone does not incur a social stigma to use in public, and many people already own one. Distribution is simple since the localization program can be downloaded as an 'app'. A modern smart-phone such as the Google/Samsung Nexus series has a vast array of available sensors, most of which are applicable to the localization problem. These include GSM and Wifi radios, Bluetooth, Near Field Communication (an extension of RFID), a camera, accelerometer, magnetometer, and gyroscope sensors.

GSM signal strengths have been used to determine which floor of a building the user is in [4]. However, accurate tracking within a floor of a building has not been achieved with GSM. Similarly, the short ranges of Bluetooth and Near Field

Communication (NFC) have been used to constrain the estimate of the user's location [7][8][11]. Bluetooth and NFC have the drawback of requiring installation of markers in each environment where the solution is used. We prefer an approach that does not require the installation of additional infrastructure. Localization using visual data is a rich field of research. A major drawback to its use in a cell phone is that it requires the user to actively take pictures of the environment to determine their location. Another drawback is the need to either enhance the environment with artificial landmarks [9], or to create large databases of training information [10]. We chose to focus on techniques that do not require active user intervention to perform continuous localization. There has been extensive work in Wifi approaches [5][6] from which a baseline implementation using a particle filter and a k-nearest neighbor approach has emerged. However, RSSI positioning alone does not have the responsiveness achievable with dead reckoning. A dead-reckoning system comprising of accelerometer, magnetometer, and gyroscope sensors can provide fast and accurate estimation of local pose. Steinhoff and Schiele [12] have demonstrated an experimental dead reckoning module with inertial sensors that operate from the user's trouser front pocket. This approach is able to exploit the swinging of the user's thighs as the user walks to extract strong movement signals. For this work, we focused on the problem of localizing the phone when it is held in the user's hands and consequently less information is available in the accelerometer signal. Although reasonably accurate over the short-term, this system is not a complete solution because the localization error significantly increases over time, predominantly due to sensory drift. We overcome this by combining multiple complimentary localization systems including dead-reckoning and RSSI fingerprinting for robust localization. The selected sensor mix is attractive because it can provide a bounded-error solution that does not require any additional infrastructure and is robust to signal dropouts and noise.

Not only are there many sources of localization data, there are also many ways of integrating disparate information sources together to derive a combined location estimate. Bayesian methods, which use a probabilistic framework to perform reasoning, are a theoretically sound way to combine multiple observations of different types. Extended Kalman Filters and Particle Filters are commonly used in the robotics domain for localization [13]. Kalman filtering methods break down when the state domain is restricted by the presence of obstacles and when multiple hypotheses need to be tracked. Both of these situations are common in indoor localization problems, so Kalman filtering approaches are not suitable to the task. Particle filtering methods, in contrast, are well suited for the indoor localization problem. They can readily incorporate nonlinearities introduced by Wifi signal strength filtering, restrictions imposed by the environment map, and other sources. Furthermore, they are able to track multiple hypotheses in a natural way, by dividing particles among hypotheses in proportion to their likelihood.

Evennou and Marx [14] have shown a combined dead reckoning and Wifi-based system which uses a particle filter to merge observations. They used a custom rigidly-mounted belt module with a Wifi radio, 2 axis accelerometer, gyroscope, and pressure sensor. Their inertial movement sensing was accomplished by a combination of step counting and gyroscope integration. However, they did not clearly explain their step counting algorithm or how they obtained a zero-reference for their gyroscope. Indeed, their sensor mix is not sufficient for establishing global heading in the dead reckoning component without requiring the user to be facing the same direction every time the system is started. Our work addresses some of the shortcomings in that of Evennou and Marx. We introduce the use of a magnetometer sensor, which allows the user's orientation to be derived in a global frame. We discuss methods of detecting distance travelled with an accelerometer in more detail, including a new variance-based algorithm that we have found to have better performance than step counting techniques. Most current method for building a RSSI to position database are tedious and require a large number of samples. Further, changes to environment require extensive re-calibration. We present an automated system for collecting signal strength calibration data which can significantly reduce the burden of creating the fingerprint database.

III. THE INDOOR LOCALIZATION SYSTEM

After consideration of the benefits and drawbacks to each of these sensors, we selected the Wifi radio combined with dead reckoning using accelerometer, magnetometer, and gyroscope sensors for use in our system. This sensor mix is attractive because it can provide a bounded-error solution that does not require any additional infrastructure investment, assuming that the environment already has an adequate number of Wifi access points. Our approach combines the multiple complimentary localization systems including dead-reckoning, Wifi, and GSM using a particle filter for robust localization over multiple-floors of an indoor building as shown in figure 1. A gait-based motion model combined with a heading estimator provides a pre-filtered dead-reckoning sensor estimate to the particle filter (PF). Simultaneously, a pose estimate based on fingerprinting between observed wifi signal strength readings and pre-collected database of RSSI estimates. The combined sensor data is fused and filtered using a particle filter (PF) which results in a smooth and continuous pose estimation state.

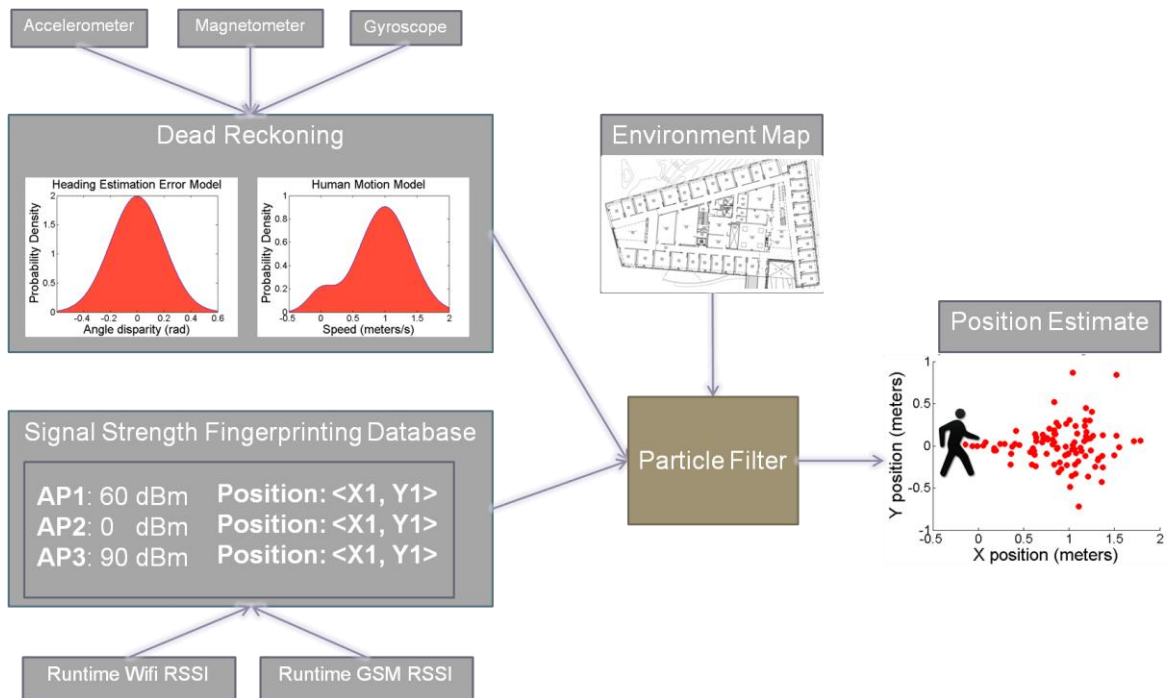


Figure 1: System Architecture

The dead reckoning component can be decomposed into heading estimation, discussed in Section A, and movement detection, discussed in Section B. The Wifi signal strength fingerprinting approach requires a calibration step, discussed in section C. At runtime Wifi signal strength fingerprinting is used to initialize the system and provide a rough global location estimate. The user's movement is also tracked at high frequency using dead reckoning. We introduce a particle filter in Section D to combine these disparate sources of information. To minimize the computational requirements of the solution, it is desirable to keep the dimensionality of the particle filter as low as possible. For this reason, we do not track heading within the filter, but estimate only the linear position. The dead reckoning is performed in a pre-processing step, and all the particles in the filter are periodically updated based on a model of the variance of the dead reckoning estimate.

A. Determining Heading

The system should provide accurate heading estimates regardless of how the phone is being held. For example, the phone may be stored vertically in the user's shirt pocket or horizontally in their hands as they interact with it. This means that the orientation of the phone must be tracked over all three spatial axes.

We track the orientation using two complimentary methods. The first method is to employ an accelerometer and magnetometer. The accelerometer gives a reference direction for gravity, and the component of the magnetic field perpendicular to gravity gives a reference for magnetic north. The benefit of this method is that each measurement is externally referenced, so orientation errors do not accumulate over time. The drawback is that magnetic anomalies which interfere with the operation of the compass are common in indoor environments due to electrical wiring and use of metal furniture or building materials. Additional error will be imparted from imperfect separation of the gravity signal from the linear accelerations imposed on the phone. The second orientation tracking method is to use a gyroscope. Gyroscope sensors are much less noisy than accelerometers or magnetometers, and are not susceptible to external interference. However, because they measure angular velocity rather than angular position, only the relative movement of the phone can be derived from the gyroscope readings.

Furthermore, the error in the orientation tracking will accumulate over time without a bound, since the gyroscope readings are open-loop. To robustly determine heading, these two methods can be merged. The combined orientation filter continuously accounts for drift in the gyroscope and error conditions in the magnetometer by using one sensor to compensate for the failings of the other. The inputs to the filter are raw accelerometer, magnetometer, and gyroscope readings as they are received. The output is an estimate of the azimuth, pitch, and roll of the phone in a global frame (figure 2 a). As acceleration and magnetic field strength readings are received, the direction of gravity and of north may be estimated.

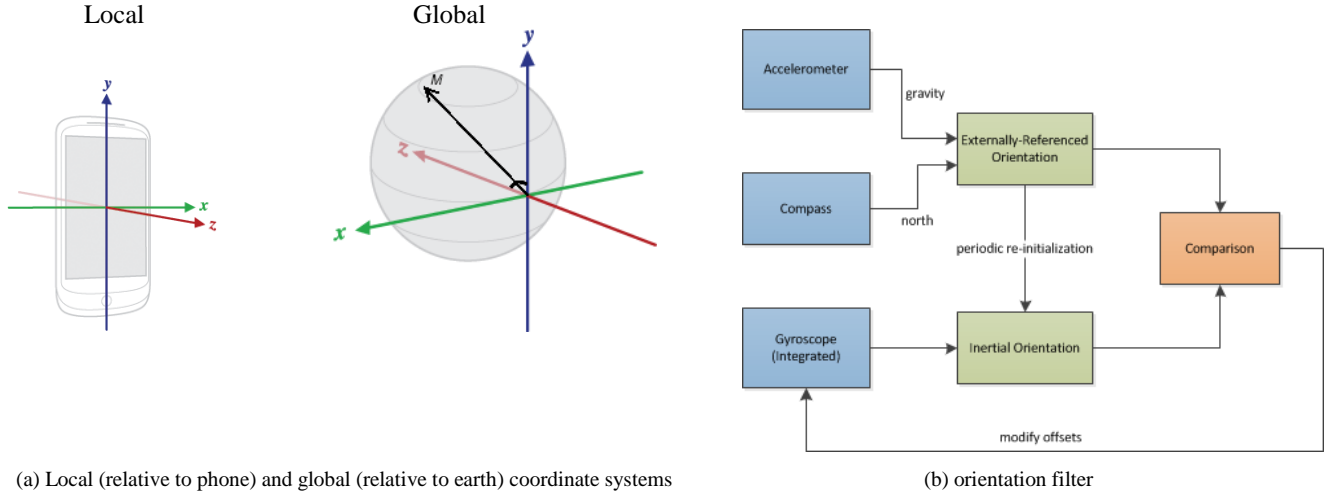


Figure 2: Heading estimation using accelerometer, magnetometer, and gyroscope

Figure 2b outlines our low-pass orientation filter which is applied to the sensor readings to remove transient effects, based on the assumption that the phone itself is not moving rapidly. Next, a preliminary analysis is done to assess the validity of the gravity and compass estimates. If the acceleration signal deviates significantly from 1g, or the magnetic field strength deviates significantly from the strength expected for the user's approximate global location (~55uT in the Northeastern United States [15]), the datum is marked as unreliable and is not used. The gyroscope data must be processed differently from the other sensors. Because it measures the derivative of the quantity of interest (angular velocity rather than angular position), the data is inherently less susceptible to noise and need not be filtered in the same manner. However, the tradeoff is the introduction of integration drift. For the purpose of input to the main body of the filter, the gyroscope data is integrated using the trapezoidal method to convert from angular velocity to change in angular position in the local coordinate system (3). At this point, the change in angular position is converted into the global orientation frame. If an estimate of the transformation from the gyroscope's local frame to the global orientation frame is already available, this can be done by applying the transformation matrix R to $\Delta\theta_{local}$:

$$\begin{aligned}
 accel_{filtered} &= \text{lowpass}(accel_{raw}) \quad (1); \quad magnetic_{filtered} = \text{lowpass}(magnetic_{raw}) \quad (2) \\
 \Delta\theta_{local} &= \frac{1}{2}(gyro_{clean}(t) + gyro_{clean}(t-1)) * \Delta t \quad (3); \\
 \Delta\theta_{global} &= R \Delta\theta_{local} \quad (4); \quad \theta_{global} \leftarrow \theta_{global} + \Delta\theta_{global} \quad (5)
 \end{aligned}$$

If the transformation is *not* available, it must be derived from $accel_{filtered}$ and $magnetic_{filtered}$. If one of those inputs was deemed unreliable because it was out of range, then there is insufficient information to initialize the orientation estimate and the system falls back on the particle filter until reliable orientation data are available.

If $accel_{filtered}$ and $magnetic_{filtered}$ are within nominal range, the computation of R and θ_{global} can proceed. $accel_{filtered}$ is assumed to point directly down, and the component of $magnetic_{filtered}$ perpendicular to $accel_{filtered}$ is assumed to point to magnetic north. In (6), a unit vector in the 'down' direction is created by normalizing $accel_{filtered}$ to have a length of one. In (7), a unit vector in the direction perpendicular to both 'down' and 'north' is found by taking the cross product and normalizing. This is the west direction in the coordinate system shown in Figure 2a. Finally, a unit vector in the 'north' direction can be generated by taking the cross product of the two other unit vectors.

$$\begin{aligned}
 Down &= \text{normalize}(accel_{filtered}) \quad (6) \\
 West &= \text{normalize}(magnetic_{filtered} \times Down) \quad (7) \\
 North &= Down \times West \quad (8)
 \end{aligned}$$

The transformation matrix between local and global coordinate systems can be expressed by creating a matrix of the global axis vectors expressed in the local coordinate system. While the matrix alone is sufficient to express the orientation of the phone, but for convenience the orientation can also be expressed in Euler angles:

$$Azimuth = \text{atan2}(West_y, North_y) \quad (9)$$

$$Pitch = \text{asin}(-Down_y) \quad (10)$$

$$Roll = \text{atan2}(-Down_x, Down_z) \quad (11)$$

$$\theta_{\text{global}} = \langle Azimuth, Pitch, Roll \rangle \quad (12)$$

At this point, we perform an additional check to verify the quality of the magnetic data. In the above calculations, the component of the magnetic field pointing in the down direction was discarded in the cross-product operations. However, the angle of that component, called the magnetic inclination, provides important information about the magnetic field. If the inclination angle varies significantly from the expected value for the user's coarse global location, this is another indicator of a magnetic anomaly. The inclination 'I' is calculated as follows:

$$I = \text{atan2}(North \cdot magnetic, Down \cdot magnetic) \quad (13)$$

Occasionally, θ_{global} must be re-initialized from the accelerometer and magnetometer data to limit the accumulation of error from gyroscope integration.

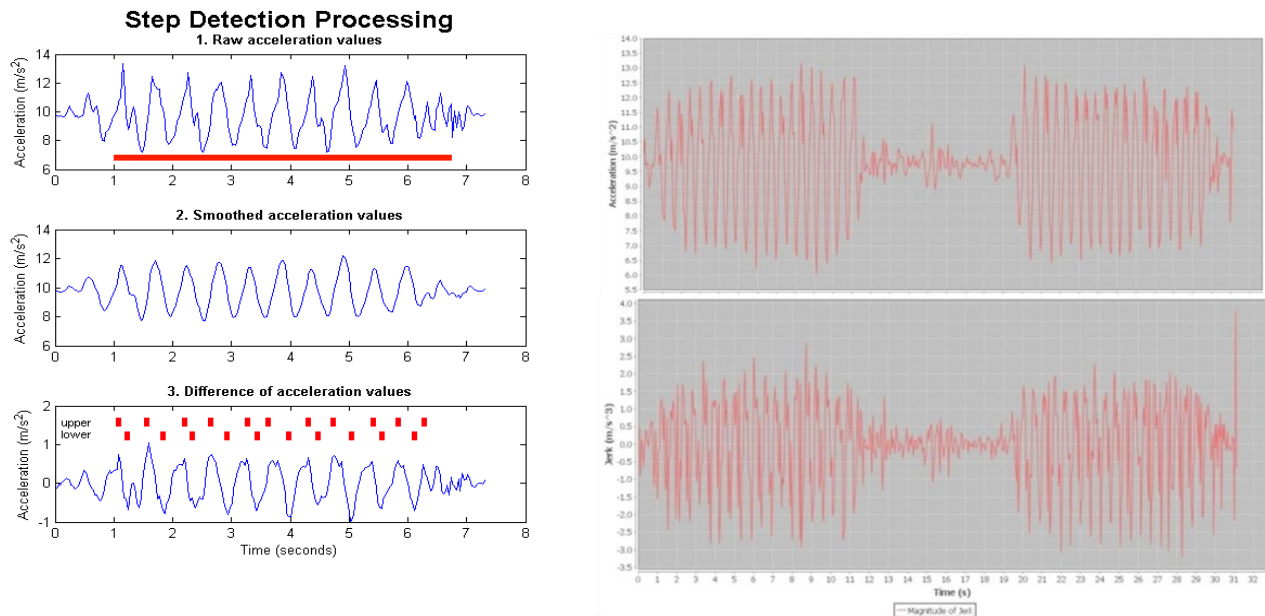


Figure 3: Accelerations during a short walk.

B. Movement Detection

Detecting when the user moves is relatively easy and reliable with only an acceleration sensor. Two different approaches have been explored, both of which exploit the periodic nature of the walk cycle. The first method uses a peak detection filter with alternating high and low thresholds to detect individual steps. The second method, which has the advantage of detecting movement continuously in time rather than at the step level, simply looks at the variability of the acceleration readings, with the assumption that higher variability corresponds to movement.

1) Step Detection

There are two major effects that must be compensated for to detect steps accurately using a threshold filter. The first is momentary accelerations and noise which could create spurious peaks. The second is the effect of gravity. Depending on the orientation of the phone, accelerations from walking may alternately cancel and add to the gravitational signal, or add to it in both parts of the walk cycle. Furthermore, MEMS accelerometers often have poor calibrations, meaning that the base gravity signal cannot be reliably assumed to be equal to the correct value of 9.81m/s^2 . These effects are handled by first smoothing the accelerations using a running average filter, and then removing offset errors and drifts by taking the first difference of the smoothed result. To further suppress spurious peaks, the peak detector has a short 'cool down' period whereby it doesn't report additional peaks found 400 milliseconds after an initial one. Figure 3 shows a sample of the registered variations in the acceleration reading during a short walk. The line under graph 1, on the left, indicates the times the standard deviation test would indicate 'movement'. The points in graph 3 indicate when the upper and lower thresholds of the peak detection filter would fire.

2) Continuous Movement Detection

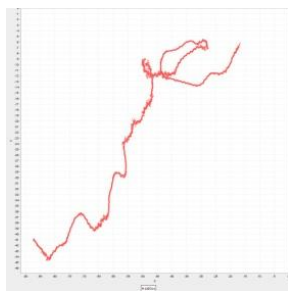
The precision of the step detection was found to be wanting in scenarios where the user was turning frequently. If the state of ‘movement’ was detected instead of individual steps, the user’s estimated position could be updated much more frequently, which is especially useful during rapid changes of heading. A simple way to detect movement is to take the standard deviation of the acceleration values over a fixed time span. The size of the time span and the threshold that must be met to report movement both contribute to a tradeoff between the delays that is incurred before the filter is able to detect movement, and the possibility of false positives. Figure 4 shows the comparative performance between the two different dead-reckoning algorithms when compared against a GPS in an outdoor environment. While both system show fairly good performance, the continuous movement model outperforms the step-detection method by a factor of 2. Furthermore, there were less observed false positives with the continuous method when compared against the step-detection method.



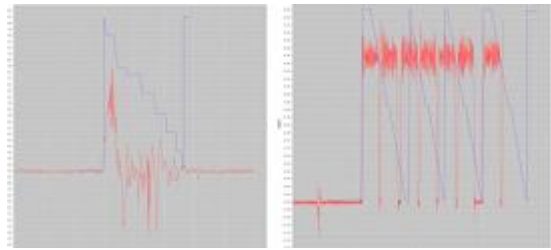
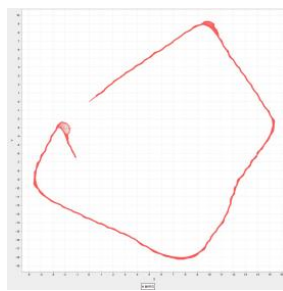
Localization Method	Error (m)	Percentage Error (%)
GPS	1.0	0.5
Step Detection	17.5	8
Continuous Movement	9.0	4

Figure 4: Comparison between implemented dead-reckoning algorithms over a 220m run outdoors at Carnegie Mellon University campus.

Despite the fast estimation time and promising results, there are a few drawbacks to using dead-reckoning only estimation. Since the dead reckoning algorithm is a local estimation technique, it has to be seeded with an initial position for valid estimation. This may not be always feasible. Additionally, the on-board gyroscopic sensor, while effective for estimating pose over short distances, accumulates error and drifts un-bounded over time (see figure 5b). Furthermore, the magnetometer can be easily distorted by magnetic anomalies, which are relatively common in cluttered indoor environments (figure 5a).



(a) Effect of magnetic anomaly on pose estimation



(b) Gyroscopic drift over time (red- phone, blue ground truth), x-axis represents time while y-axis is the measured angular rate

Figure 5: Sensitivity of on-board sensors

C. Signal Strength Calibration

To compensate for the drawbacks of the dead reckoning system, we use a Wifi-based signal strength fingerprinting approach. To perform the signal strength fingerprinting, it is necessary to create a database of signal strength information from the environment correlated to a free space map of the environment. This pose-RSSI database forms the basis for comparison and subsequent estimation. Collecting this information by hand would be laborious and prone to error, so we developed an automated solution to this problem which uses a robot equipped with a SICK LMS200 laser rangefinder. A phone is placed on the robot and the robot is tele-operated through the environment. The phone collects signal strength information continuously in time, and then the signal strength readings are correlated with the position of the robot at the time they were taken. At the same time, the robot builds a map of the environment using the laser rangefinder. The result is an accurate, high density sample of signal strength information in a short amount of time. Figure 6 shows the robot setup and the RSSI database for a sample environment.

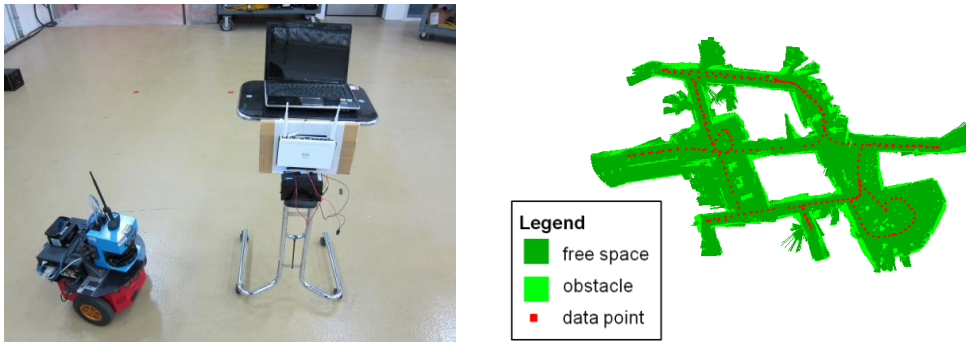


Figure 6: RSSI database creation using pioneer robot

D. Particle Filter

The use of a particle filter is critical for minimizing the accumulation of localization error over time [16]. Particle filters are able to estimate the parameters of a system (such as the user's location) using Bayesian principles. They use sampling methods to track many possible hypotheses, updating them every time new information becomes available. Unlike a Kalman filter, particle filters are easily adapted to handle the presence of obstacles like walls, non-linearities and non-Gaussian noise models, multiple hypotheses, etc. without special extensions to the filter. This is particularly important when such a wide variety of sensing types are being combined. A particle filter will usually be more computationally expensive than a Kalman filter, but they can still work with reasonable speed on a cell phone processor. Unlike that of a Kalman filter, the performance of a particle filter can be scaled with available computation power by varying the number of particles that are tracked.

A particle filter has 3 major components: a *motion model* that updates the positions of particles, observation model that sets particle weights, and a *re-sampling algorithm* for modifying distribution to reduce variance. For our system, the motion model is derived from the dead reckoning information given by the accelerometer, magnetometer, and gyroscope. They are treated similar to how wheel odometry is used in particle filters for localizing robots. The heading and movement speed given by the dead reckoning methods described above are extended into a spatially directed Gaussian distribution (see Figure 7). The variance in the forward component of the distribution is related to the walking speed of the user, and the variance of the angular component is related primarily to the measurement error introduced by the gyroscope during turns. During the movement model update, each particle is moved by an amount that is sampled from that distribution.

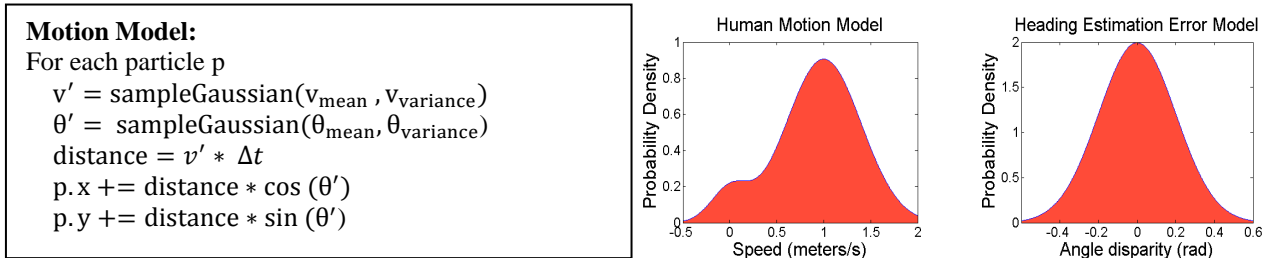


Figure 7: Motion model for the particle filter

The external measurements used by the filter to update particle likelihoods come from two different sources. The first is a map of the floor of the building. Whenever a particle enters a space that is indicated to be impassible in the map, its likelihood is reduced to zero, so it is thrown away during the resampling step. This ensures that the particle filter does not give location estimates that are impossible. If the particle falls outside of the map by leaving it through an open area, it is not removed. This allows the localization to continue to function even if the user has left the mapped area. The second source of external measurements is Wifi signal strengths. The signal strengths to several Wifi access points (APs) measured by the phone at localization time are compared with a signal strength map generated earlier. For each particle, the difference between the expected reading for that position in the map and the actual signal strength reading is used to adjust the weight of the particle. We used a Euclidean distance metric in signal space between the APs common to both readings. Effective distance is calculated as a weighted average of the nearby calibration points to reduce noise.

$$\text{Distance} = \sqrt{\sum (\text{measured signal} - \text{reference signal})^2}$$

In addition, a penalty was imposed for APs that were disjoint between the two readings in proportion to the signal strengths of those APs. The intuition is that a faint signal would be expected to drop in and out with changing conditions, but a strong signal should always be visible in a fixed location. We use the importance sampling algorithm to resample particles [16]. This method reduces the variance of the particle distribution by throwing away unlikely particles and duplicating particles

that are more likely. The pre-built environment map structure is used for weighing the sampling algorithms. Particles that lie outside the constraints of the environment are weighed lower than others. Specifically, the new particle distribution is formed from the old one by selecting particles (with replacement) with probability proportional to their weight. This focuses the particle distribution on the area of maximal interest. Figure 8 shows how the filter varies over time for a sample environment.



Figure.8: Particle filter estimate over time for environment 2. The standard deviation of the PF bubble is $\sim 1\text{m}$ after convergence over the duration of the run

IV. EXPERIMENTS AND RESULTS

We developed and implemented our localization solution for the android platform using a Nexus S smart-phone. Towards analyzing the feasibility of the implemented solution, we tested the system in two different indoor environments. A closed loop test was performed in which the participants walked a path through hallways delineated with cones. The participants were instructed to hold the cell phone device in their hands pointing forward and to walk at their normal pace. Four individuals participated in the test for environment 1 (Figure 9a) and three for environment 2 (Figure 8). Each set of experiment was repeated multiple times for consistency. The total length of the traversed path in environment 1 was about 120 meters, and the path in environment 2 was about 72 meters. There were about 30 Wifi access points (figure 8b) in both environments. The dead-reckoning information was obtained at a faster rate (30Hz) than the wifi signal measurement (1 HZ). For each dataset the performance of the system using only dead reckoning, only Wifi, and both dead reckoning and Wifi were evaluated to give a clear idea of the individual and combined performance of these subsystems (figure 8). Further two additional sets of metrics were analyzed, online performance and offline performance. In the case of the offline system, sensor data collected from the above runs were run through the particle filter and post-processed offline on a laptop. For the online experiment, an online version of the filter was run on the device and the resultant output tabulated. This is implemented as a standalone Android platform activity that is used for continuous pose estimation. The limited processing capabilities of the Android handset required major modification to the particle filter algorithm from the off-line version. Specifically, the updater (dead reckoning) portion of the particle filter has been separated so that the orientation accuracy can still be maintained while the particle filter processes the WiFi data. Further, the online version runs on a much more sparse data so as to not overwhelm the system set whereas the offline system has the advantage of running the full data set.

Over the relatively short duration of the experiments, the dead reckoning method by itself was able to track the participants with relatively low error (figure 8c) and outperformed the wifi only estimation (figure 8d). For environment 1, the off-line dead reckoning had a mean path error of around 5m, whereas the wifi was on the order of 10 m. While promising, this is somewhat misleading, because of two factors. First, in both instance the particle filter was manually initialized to the correct position for the dead reckoning only results, whereas it was given a uniformly random initialization in the trials that used Wifi. Having a confident estimation of the starting location is not a given for all situations. Ideally it would not be necessary to manually initialize the filter, as this places an additional burden on the user of the system. Second, dead reckoning is known to drift from the correct position over time; it would be unlikely to recover from that error. Finally, as mentioned before, a dead-reckoning based method that uses magnetometer is susceptible to magnetic anomalies prevalent in the environment. We specifically chose to test in an environment that had lower magnetic anomalies, so as not to bias the system.

Interestingly, the online system shows a lower mean error variation than the offline system (Table1) and both had similar closed loop error measurements ($\sim 6\text{ m}$) averaged over the different runs. The higher computation of a laptop allows the offline system to process the data at a much faster rate than the online system. From figure 8g, we can see that the offline PF takes about 0.5 sec per processed reading whereas the online system can only output a reading about once every 1.5 sec. However, the larger data set introduces an increase in sensor variation, especially from wifi readings, leading to a reduced performance output. The slower processing rate of the online system results in the error growing at a much slower rate leading to an improved estimation. As part of our future work, we will look to identify the ideal balance for the number of sensor readings to use for optimal performance. This indicates that indoor localization is viable even on devices which do not have the necessary sensors for dead reckoning.

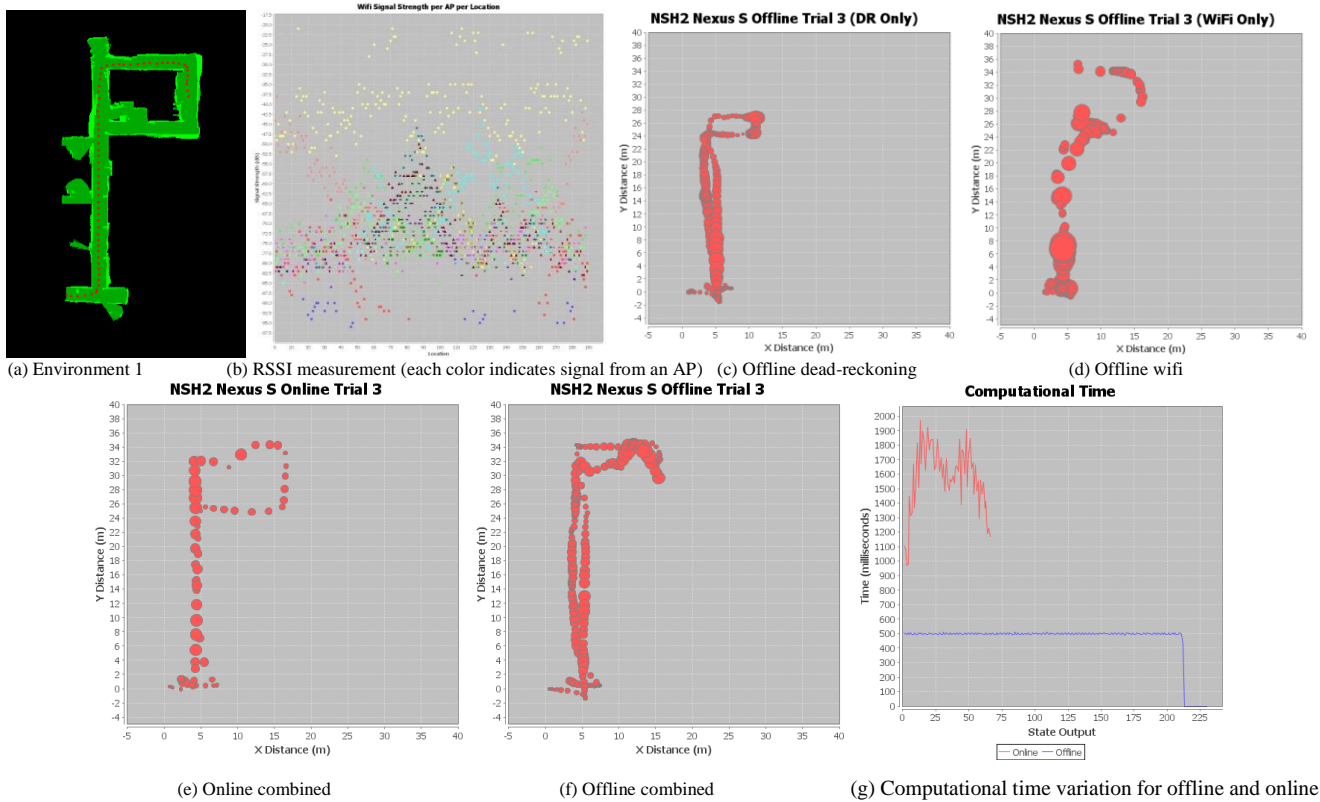


Figure 9: Experimental results for environment 1

However, the performance of Wifi is restricted by the extent of the collected RSSI database. This shortcoming could be potentially be mitigated using better signal strength interpolation methods. The combined dead reckoning and Wifi method is able to achieve the precision of dead reckoning and the favorable initialization and error recovery properties of Wifi. It is also able to degrade gracefully if one of the components is unavailable, by simply not feeding that measurement to the particle filter. If errors are intermittent, the filter will continue to function without any significant disruption.

Table 1. Experimental Results				
	Mean Error of Most Probable Location		Additional Features	
	Environment 1	Environment 2	Automatically Initializes Position	Error Bounded over Time
Dead Reckoning Only	5 ± 3 meters	6 ± 2 meters	No	No
Wifi Only	15 ± 10 meters	10 ± 4 meters	Yes(<5 seconds)	Yes
Dead Reckoning and Wifi (offline)	5 ± 3 meters	7 ± 2 meters	Yes(<5 seconds)	Yes
Dead Reckoning and wifi (online)	3 ± 3 meters	<i>we currently do not have comparative results for environment 2</i>	Yes(<5 seconds)	Yes

V. CONCLUSION AND FUTURE WORK

In this paper, we have outlined a smart-phone based indoor localization system using dead-reckoning and wifi RSSI fingerprinting that is precise enough for giving way-finding directions and for use in context-aware applications. We have also outlined a robot-based technique for reducing the time needed to acquire calibration data and building the RSSI database. As a result, the time and effort needed to set the system up are minimized, maximizing its commercial relevance.

There are several interesting directions to explore in future work. The weakness of wifi signal estimation can be improved by incorporating other sensors into the localization framework, such as GSM, Near Field Communications (NFC), or even a camera. GSM is useful for extending the system to localize over multiple floors, whereas NFC and a camera could impose additional constraints within a floor to increase accuracy. Further, the quality of localization can be refined by exploring the extent to which the system can improve through learning. The history of users' paths through an environment provides information that could be exploited to increase the system's performance.

VI. ACKNOWLEDGEMENTS

The authors would like to acknowledge Evan Glasgow, Victor Marmol and other members of the rCommerce Lab at Carnegie Mellon University for their important contributions towards software development and testing. This work was sponsored by the Boeing Company under contract CMU-BA-GTA-1. The content of this publication does not necessarily reflect the position or policy of the sponsors and no official endorsement should be inferred.

VII. REFERENCES

- [1] R. Stirling, J. Collin, K. Fyfe, and G. Lachapelle. "An Innovative Shoe-Mounted Pedestrian Navigation System" *Proc. European Navigation Conf. GNSS*, 2003, pp. 22-25.
- [2] S. Beauregard, "A helmet-mounted pedestrian dead reckoning system," *Proceedings of the 3rd International Forum on Applied Wearable Computing (IFAWC 2006)*, Citeseer, 2006, p. 15-16.
- [3] C. Randell, C. Djiallis, and H. Muller. Personal position measurement using dead reckoning. *Seventh IEEE International Symposium on Wearable Computers, 2003. Proceedings.*, 166-173. Ieee. doi: 10.1109/ISWC.2003.1241408.
- [4] V. Otsason, A. Varshavsky, A. LaMarca, and E. De Lara, "Accurate gsm indoor localization," *Ubiquitous Computing*, 2005, p. 141-158.
- [5] I.K. Adusei, K. Kyamakya, and F. Erbas, "Location-based services: advances and challenges," *Electrical and Computer Engineering, 2004. Canadian Conference on*, IEEE, 2004, p. 1-7.
- [6] V. Honkavirta and T. Perala, "A comparative survey of WLAN location fingerprinting methods," *WPNC 2009* pp. 243-251
- [7] S. Feldmann, K. Kyamakya, A. Zapater, and Z. Lue, "An indoor Bluetooth-based positioning system: concept, implementation and experimental evaluation," *International Conference on Wireless Networks*, 2003, p. 109-113.
- [8] M. Kourogi, N. Sakata, T. Okuma, and T. Kurata, "Indoor/outdoor pedestrian navigation with an embedded gps/rfid/self-contained sensor system," *Advances in Artificial Reality and Tele-Existence*, 2006, p. 1310-1321.
- [9] K.J. Yoon and I.S. Kweon, "Artificial landmark tracking based on the color histogram," *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, IEEE, 2001, p. 1918-1923.
- [10] S. Se, D. Lowe, and J. Little, "Vision-based mobile robot localization and mapping using scale-invariant features," *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*, pp. 2051-2058.
- [11] Using Inquiry-based Bluetooth RSSI Probability Distributions for Indoor Positioning, Ling Pei, Ruizhi Chen, Jingbin Liu, Heidi Kuusniemi, Tomi Tenhunen, Yuwei Chen, *Journal of Global Positioning Systems* (2010), Vol.9, No.2 :122-130 DOI: 10.5081/jgps.9.2.122
- [12] U. Steinhoff and B. Schiele. Dead reckoning from the pocket - An experimental study. *2010 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 162-170. Ieee. doi: 10.1109/PERCOM.2010.5466978.
- [13] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. Cambridge, Massachusetts: The MIT Press, 2005.
- [14] F. Evennou and F. Marx. "Advanced Integration of WiFi and Inertial Navigation Systems for Indoor Mobile Positioning" *EURASIP Journal on Advances in Signal Processing*, vol. 2006, pp. 1-11
- [15] NOAA. "US/UK World Magnetic Model -- Epoch 2010.0" Available: <http://www.ngdc.noaa.gov/geomag/WMM/DoDWMM.shtml>
- [16] A. Howard, S. Siddiqi, and A. Sukhatme. An experimental study of localization using wireless ethernet. *Field and Service Robotics* (p. 145-153). Springer. Retrieved March 5, 2011, from <http://www.springerlink.com/index/mu1333820763207h.pdf>.
- [17] Google Inc, "Android API Reference," Available: <http://developer.android.com/reference/packages.html>