Using Micro-Climate Sensing to Enhance RF Localization in Assisted Living Environments

Anthony Rowe, Zane Starr and Raj Rajkumar
Real-Time and Multimedia Systems Laboratory
Department of Electrical and Computer Engineering
Carnegie Mellon University, Pittsburgh, PA 15213
{agr, zcs, raj}ece.cmu.edu

Abstract

In this paper, we propose micro-climate sensing as an effective means of enhancing conventional RF-based localization. Our system targets people tracking applications in dynamic indoor environments, such as nursing homes, hospitals and office spaces that require simple deployment and where conventional RF tracking may suffer from time-varying signal attenuation or dropped packets. RF-based localization approaches suffer as the environment changes over time. To help mitagate these effects we use time synchronized windows of sensor samples to associate a mobile node with its nearby beacon nodes. In assisted living environments sensor networks likely already have basic sensors to collect contextual information about users and to monitor the environment. We propose using light, humidity, temperature and audio data samples over a short window of time to model the micro-climate of a beacon node. Using micro-climate matching in conjunction with RF signal strength decreases the worst-case localization error significantly by a factor of more than 3 (from 25m to 8m) while making the system more resilient to environment changes. Micro-climate data helps ensure at least room level location tracking even in buildings like hospitals with many rooms in close proximity.

I. Introduction

Newly emerging low-power wireless technologies such as IEEE 802.15.4 are making it increasingly more feasible to deploy large-scale RF-based tracking systems. In settings like hospitals, tracking of patients and inventory requires not only meter-level accuracy, but more importantly the ability to distinguish between rooms. Hospitals are a particularly difficult environment due to dynamic features such as movable walls, heavy metal machinery and people constantly moving throughout the environment. In these situations, conventional RF-based localization can suffer greatly from dropped packets or unusual signal attenuation due to new obstructions. A practical system must address scalability in terms of covering large areas and of tracking multiple mobile entities. In this paper, we present a technique for utilizing micro-climate sensing using wireless sensor networks in conjunction with a conventional RF-based scheme to enhance localization accuracy.

A micro-climate represents the physical attributes associated with a local environment, and can be characterized by features derived from sampling different sensors. For instance, a particular region of a room might be cooler than another or have a higher relative humidity than its surroundings. These features are not limited to simple intensity values. A time-synchronized window of audio samples may contain components such as voice or regular sounds generated by nearby machinery or HVAC ducts. Previous work [1] has shown that light analyzed in the frequency domain shows characterizable differences between florescent, incandescent, natural sun light and the glow of an LCD panel or monitor. Many sensor networking systems already provide the needed sensors making the addition of micro-climate matching to an existing RF-based system relatively simple.

In an assisted living environment ideally the system should know a user's location with a high degree of accuracy. Figure 1 shows an example scenario where sensor nodes are placed around a nursing home. The system uses the sensors to ascertain contextual information about the users so that it can monitor daily patterns and generate reports for care takers or family members. Given that the

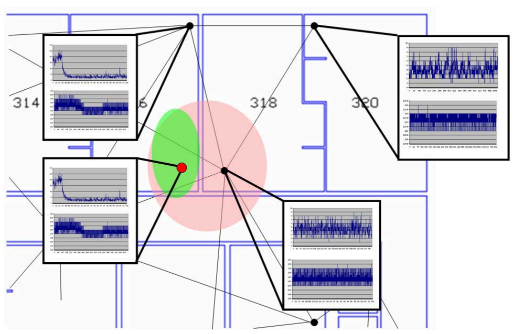

Fig. 1. This diagram shows the floor plan of a nursing home outfitted with sensor nodes. The boxes associated with each node show sample sensor data. In this scenario, the micro-climate data helps associate the mobile node with the correct room instead of only basing the position on potentially misleading RSSI values.

sensors are already present in the system, we propose utilizing this data to help improve location tracking. In Figure 1 we see an example where RF signal strength data would likely incorrectly locate a mobile node in room 318. In cases like this sensor data can correct such errors.

We also characterize how the performance of IEEE 802.15.4 body-worn nodes compares with the results of previous work using IEEE 802.11. We show that the higher node density that one expects from shorter range 802.15.4 nodes compared to typical 802.11 networks improves location accuracy to better than 1 meter. Unfortunately, using RF-only techniques we found that by the 95^{th} percentile, the error becomes greater than 15 meters. In other words, 5% of the locations detected have errors of 15 meters or greater. Introducing the micro-climate information rectifies many of these outlying points bounding even the *maximum* error to within 8 meters (previously 25 meters).

For the sake of convenience, we show the performance of two RF-only based localization techniques using a testbed setup in an active office environment. While not as ideal as testing in a nursing home or hospital, we believe our office environment is likely more homogenous with respect to sensor data and hence less advantagous for our approach. Using data from the testbed we explore the failure cases of these different schemes. We investigate the performance of different sensors and sensor features for micro-climate localization. Finally, we show how micro-climate matching can be used to improve existing localization techniques and help guarantee room-level location accuracy.

A. Related Work

Various RF-based localization systems have been studied in the past and range from indoor centimeter accurate tracking to outdoor GPS-style systems. Multiple systems have tried to augment RF systems with an alternative transmission medium like ultrasonic or IR. The Active Badge system [3] uses IR which suffers from limited range, reliability problems in intense ambient light as well as scaling issues with many nodes in a single IR collision domain. The MIT cricket [4] localization system uses the time-of-flight difference between RF and ultrasonic signals. The system is capable of extremely accurate range measurements, but requires line-of-sight communication, careful node positioning and high node density. It suffers worse scaling problems than IR since the ultrasound pulses propagate more slowly and travels father than IR.

Outdoor systems such as GPS and the wireless Enhanced 911 system use multiple RF-based approaches to localize positions. GPS uses the time difference of arrival (TDOA) of radio signals, while Enhanced 911 uses TDOA or multiple antennas to calculate the angle of arrival (AOA). Due to multi-path effects and high indoor path-loss exponents, these systems are only sutable outdoors.

Extensive work has been done on modeling RF propagation and deploying systems based on received signal strengths[5], [6], [7], [8]. These systems require mapping signal strengths to distances based on known radio characteristics. This type of approach is highly dependent on signal propagation in the environment as well as individual antenna and receiver-specific properties.

The RADAR [2] system was one of the first to use pre-recorded RF signatures to build a database of signal strength values at particular locations. A mobile node receives signal strengths from the stationary beacons which can then be compared with reference points to find the best match. The MoteTrack [9] system uses a similar approach with an emphasis on distributed operation in a sensor network. We use the site survey database approach to establish a baseline in which to compare our reference-free and micro-climate-based approaches. We also find that the database approach performance degrades over time with only subtle changes in the environment.

The research reported in [10] describes a range-free localization technique where each mobile node derives its position by calculating the center of the location of all beacons it hears. If beacons are placed well, localization errors can be decreased [11], but this is not possible in all environments. In contrast, we improve the centroid algorithm by adding micro-climate sensor information.

Utilizing information from the environment is perhaps the most basic and frequently overlooked source of localization information. Cameras and video have been used [12], [13] to localize objects in the environment. The eWatch [1] is a mobile computing platform that uses a sensor feature database for coarse-grained location identification. The work focused on identifying if a mobile user was in their apartment, office, riding in a bus, etc. Many of these environmental features such as background noises will slowly change over time, hence it becomes difficult to uniquely classify a large number of locations. In our system, we use instantaneous time-synchronized sensor features to correlate with nearby locations. Hence, we do not require a database and the approach does not suffer from slow changes in the environment over time.

B. Organization of The Paper

The rest of this paper is organized as follows. Section II provides a quick overview of the FireFly sensor network we use in our experiments. Section III describes the RF-only localization techniques we evaluate in our testbed. Section IV introduces our approach to sensing and fusing micro-climate information. In Section V, we present our detailed evaluation of the RF-only localization and micro-climate sensing approaches. Finally, in Section VI, we draw some conclusions and identify future work.

II. Wireless Sensor Network Used

We now briefly describe the wireless sensor network that we use for collecting both RF and micro-climate data while performing real-time localization.

Our wireless sensor network uses FireFly sensor nodes developed at Carnegie Mellon University. Each FireFly node consists of an 8MHz ATMEGA1281 8-bit micro-controller with 128KB of Flash and 8KB of RAM. The processor is coupled with Chipcon's cc2420 IEEE 802.15.4 radio. The node includes a light sensor, microphone, 14-bit accurate temperature / humidity sensor, and a passive infrared motion detector. Each FireFly node uses the

Nano-RK sensor operating system [14], and runs the RT-link communication protocol [15]. Nano-RK is a real-time multi-tasking priority-driven reservation-based power-aware OS specifically designed for sensor nodes. RT-Link is a globally time synchronized link layer protocol that communicates using scheduled TDMA slots. Mobile nodes communicate in a slotted aloha contention period or can be temporarily leased scheduled slots. Nano-RK with RT-link supports a `nrk_wait_until_tdma_slot()` syscall which suspends a task until a particular TDMA time slot. Since the link layer is globally time-synchronized, it offers a simple and powerful means for running distributed tasks and taking (near-)simultaneous snapshots across multiple nodes.

In our deployment, sensor samples are scheduled for the beginning of each TDMA cycle. The size of the TDMA cycle can be adjusted based on the expected maximum number of users and desired localization update rate. Both of these parameters directly affect overall system energy consumption. Since all nodes in the system record their sensor samples simultaneously, increasing the number of mobile nodes does not increase the number of required beacon sensor samples. It also allows for mobile nodes to compare micro-climates with each other which can be useful for applications such as collecting social networking data.

The architecture can be set up such that the processing takes place on the beacon nodes or the mobile node. We choose to allocate most of the processing on the infrastructure nodes since this approach tends to be more memory efficient. Each mobile node in the system broadcasts a request to be localized at most once per TDMA cycle. This request contains the mobile node's sensor information and is repeated across multiple packets to allow for a stable RSSI reading. Beacon nodes are nodes permanently placed in the environement at known locations. The beacon nodes perform the sensor correlation and sends the result back to the mobile node during the beacon node's scheduled transmit slots. The beacon nodes include their locations in the reply packet so that the mobile node has all information needed to perform the weighted-centroid algorithm. The mobile node can then relay its location information back to a gateway. Other variants are also possible.

III. RF-based Localization

In this section, we describe the localization techniques that we use and that only rely on RF information. Our other techniques to be discussed later will also use RF information but will augment it with micro-climate sensing information.

A. Signature Database

As described in Section I-A, database signal strength signature methods have been shown to achieve meter-level accuracy under good conditions [2], [9]. This approach requires that the environment be populated with fixed beacon nodes at known locations. A database of RF signatures is created by moving a transmitter to specific locations within the environment called reference points. The reference point transmitter broadcasts a stream of packets that are recorded by any beacons able to receive the data. The reference point signature is composed of the average received signal strength indication (RSSI) values received by beacons. Experiences with the RADAR system showed that the directionality of the transmitter's antenna can introduce error. As they suggest, we record RSSI values while the antenna is pointing in four different directions. Once the signature database of reference points has been constructed, a mobile node is localized using a nearest-neighbor matching algorithm. The list of beacons and associated RSSI values for a particular mobile node is stored as a vector which is matched against existing RSSI vectors in the database. Similar to the approach taken in [9], we employ a Manhattan distance metric by taking the sum of the mean differences of RSSI values across the vector. The reference point with the lowest score is selected as the nearest match. If more than reference point accuracy is desired, the centroid of the set of best reference points can be found weighing each coordinate based on the closeness of the matches. This refinement is similar to our weighted-centroid technique described next.

B. The Weighted-Centroid Scheme

We now describe a weighted-centroid approach that does not require an RF propagation model or site survey. Similar to the database approach, this method requires the placement of beacons in the environment at known coordinates. RSSI values received from each node are sorted and then the highest m RSSI values are used to find the centroid of the convex polygon that falls within the selected beacon node locations. We use the following notation. The location of a mobile node is denoted by $Mobile_{(x,y)}$, and comprises of its (x, y) coordinates represented by $Mobile_x$ and $Mobile_y$ respectively. $Beacon[i]$ is the i^{th} element of the list of beacons that are in range of the mobile node. $Beacon[i]_{rssi}$ is the RSSI value between the mobile node and beacon i. $w[i]$ is a weight factor that can be associated with each RSSI value. We may also choose to apply different weight factors $w[i]_x$ and $w[i]_y$ for the x and y coordinate dimensions respectively. When we adopt the RF-only localization technique, $w[i]$ has a value of 1, but is assigned other values in later sections.

$$Mobile_{(x,y)} = (Mobile_x, Mobile_y) \tag{1}$$

$$RSSI_{total} = \sum_{i=0}^{m} (Beacon[i]_{rssi} * w[i]) \tag{2}$$

$$Mobile_x = \sum_{i=0}^{m} \left(\frac{Beacon[i]_{rssi} * w[i]}{RSSI_{total}} * w[i]_x \right) \tag{3}$$

$$Mobile_y = \sum_{i=0}^{m} \left(\frac{Beacon[i]_{rssi} * w[i]}{RSSI_{total}} * w[i]_y \right) \tag{4}$$

Since the centroid will always be bounded within the beacon nodes, mobile nodes cannot be accurately tracked once they travel outside the perimeter of the beacons. This seems to be a reasonable trade-off given the computational efficiency as compared to other trilateration techniques. The weighted-centroid approach does not require RSSI values to be converted into a distance. In an environment where the path-loss component can drastically vary, this method helps mitigate errors by weighting distance relative to all nearby RSSI values. In Sections III we explore adjusting the number of beacons used to find the centroid.

IV. Micro-Climate Sensing

A micro-climate represents the physical attributes of a local environment that can be characterized by features obtained by using various sensors. In this section, we describe in more detail the features that we explored in order to associate mobile and beacon nodes. Our choices were influenced by the minimal memory and processing capabilities on sensor nodes as well as the maximum network packet size that can be used to transport data using IEEE 802.15.4 low-power radios. For rapidly fluctuating sensor values such as light and sound, we analyze their frequency components. For slow-changing sensor values like temperature and humidity, we simply use the amplitude of the value. We evaluated the following five features: light intensity, light frequency, audio frequency, temperature, and humidity.

We have previously shown that light and audio contain characteristic frequency components that can be used to identify a particular location [1]. In a time-synchronized sensor network, this matching process is greatly simplified since signals can be compared against each other directly in real-time rather than to a large database. A major drawback to sensor database approaches is that as the number of locations increases, the difference between signatures at locations is less prominent. In our system, since a mobile node only tries to correlate with beacon nodes that are within RF range, there are very few locations to compare against, greatly decreasing the chance of ambiguity.

A. Frequency Component Features

When analyzing frequency components, each node must simultaneously sample the sensors at a consistent frequency. Once the data has been loaded into a buffer, we perform an FFT, shown in Equation (5), translating the signal into the frequency domain. Next, we look at the normalized spectral energy density of each signal as shown in Equation (7). Since the signal is time-synchronized and normalized, the correlation comparison can be done using subtraction. The correlation error is then simply the sum of the errors shown in Equation (8).

$$X_k = \sum_{n=0}^{N-1} x_n e^{-\frac{2\pi i}{N}nk} \tag{5}$$

$$k = 0, ..., N-1 \tag{6}$$

$$P = norm\left(\frac{Xconj(X)}{N}\right) \tag{7}$$

$$Error = \sum_{n=0}^{N/2} \left(P_n - P'_n\right) \tag{8}$$

The most computationally involved step in this process is the FFT which requires $\mathbf{O}(n\ log\ n)$ operations. Performing a 128-point FFT on our 8MHz micro-controller takes approximately $15ms$. A full 1024-point FFT takes about $170ms$.

B. Micro-Climate Sensor Data Fusion

In this section, we discuss approaches used to fuse together micro-climate sensor data in order to pick the overall best match between a mobile node and its closest beacon(s). Frequency features return a ranking of nodes based on their correlation error, while steady-state features simply return the beacons ranked by the nearest value to the mobile node. Each feature contains a ranking of beacons which most closely correlate to the mobile node. The next logical step is to merge this information together so as to identify the overall closest or closest set of beacons. Additionally, a confidence metric should be given with each beacon that can be used to set the weight factor $w[i]$ from Equation (2). Possible approaches include various voting schemes, weighting schemes or even environment-specific training approaches. For simplicity, we evaluate a weighting scheme where each feature has a statically defined weight. For instance, if audio and light frequencies match well between a mobile node and its closest beacon(s), but humidity or temperature are farther apart, we still want that association to be considered strong. We empirically evaluated different weighting factors and use the best combination based on results from our initial tests. Subsequent testing in other environments indicates that our initial weight factors work reasonably well under other conditions.

The centroid-weighting scheme lends itself well to additional sources of information because of the ease at which extra weight values can be added. Survey-based approaches can also use micro-climate information by applying a centroid-weighting scheme among the closest n reference points. We compute the confidence of our sensor-based correlations by looking at the difference between the best and average correlation score of the data set. If all of the beacons appear to be similarly correlated, then the sensors do not heavily influence the RSSI-based localization. If one particular point is highly correlated, but has a weak RSSI value, then this value is effectively increased pulling the mobile node toward that beacon. We use this confidence value for the best sensor-based correlated feature to set the $w[i]$ parameter from Equation (2).

V. Experimental Results

In this section, we describe our experimental setup as well as data collection and processing methodology. We then evaluate the performance of the various schemes described in Sections III and IV. Our motivations for these tests are to (a) understand how IEEE 802.15.4 systems compared with previous 802.11 base systems, (b) investigate the performance difference between database and non-database approaches, (c) better understand the failure cases of these current approaches, and (d) evaluate the performance of using micro-climate sensing.

A. Experimental Methodology

As mentioned in Section II, we use FireFly sensor nodes and the Nano-RK operating system [14] running the RT-

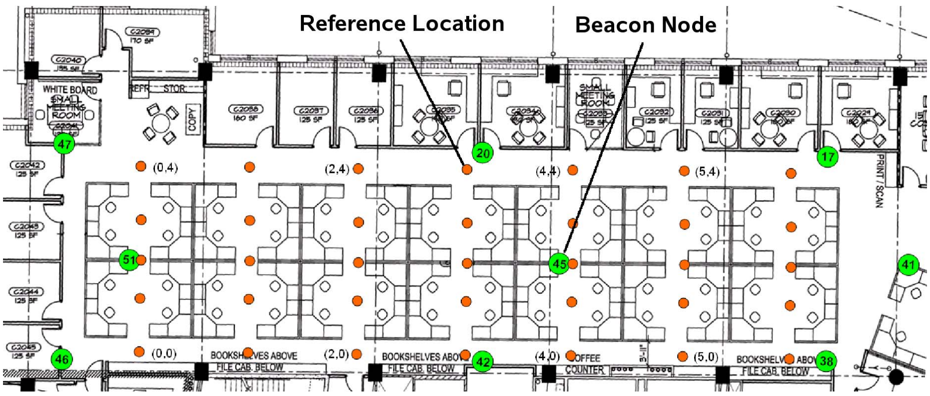

Fig. 2. Our experiments were conducted on the second floor the CIC building at Carnegie Mellon University. The 9 large circles with numbers inside them represent beacons, while the 35 smaller circles represent reference point locations.

Link [15] communication protocol. The FireFly sensor networking platform provides a collision-free time-division multiplexed access (TDMA) protocol in conjuction with a real-time operating system capable of time synchronizing tasks. We placed 9 beacon nodes as shown in Figure 2 (larger grey circles with numbers inside them) across a 40 x 15 meter area filled with cubicles in the CIC office building at Carnegie Mellon. Nodes were placed 2 meters above the ground such that each node had line-of-sight communications between each other. A desktop computer connected to a gateway node logged data from all beacon nodes in the environment. The beacons could relay sensor information as well as RSSI values across multiple hops from mobile nodes to the gateway. Our mobile node was mounted 1 meter above the ground on a tripod that could be rotated pointing the node in specific directions at each reference location.

B. Data Collection and Processing

We first performed a detailed site survey as described in Section III-A. A mobile transmiter sent 300 packets over a 30-second interval from each of the 35 reference location in the room shown in Figure 2. Each beacon that was able to receive a packet forwards it to the central gateway. We repeated the experiment at each reference point in 4 different directions recording 42,000 packets per survey. This survey was conducted late at night when nobody was using the office space. One month later, we repeated the survey during working hours so that the presence of people would induce realistic errors. Since the human body acts as a significant attenuator when present between the node and a beacon, we also conducted a second detailed survey with a person carrying the node to emulate the mobile sensor device being worn as a badge. We also recorded reference points for half of the room at three different transmit power levels.

After collecting data, we tested our various schemes using a test-set cross-validation approach. We randomly put aside 10% of our data from each reference point data set to be used as labeled input data later. This allows for simple evaluation of many data points that were realistically sampled from the environment. Since our underlying sensor network uses TDMA and is hence deterministic, any missing packets are logged as dropped ones, allowing for accurate reproduction of channel loss.

C. Results of RF-Only Localization

We first tested the site survey nearest-neighbor approach along four directions to give us a baseline set of error measurements. As expected, the site survey database approach performed well in general, giving a best-case 50^{th} percentile error of better than half a meter. Without the directional information, however, this tracking resolution is increased to 1.5 meters. It must be noted that this test should perform extremely well since all data points in the test-set lie exactly on reference points. Thus if the nearest-neighbor match is successful, then the error will be 0. We then cross-validate with different reference point configurations, by removing certain reference points from the training data while keeping them in the test-set.

We saw that as the density of reference points decreases, the error nearly linearly increased. At a density of 1 reference point for every 17, 28, 42 and 50 sq meters, we saw a 50^{th} percentile error of 3.1, 3.75, 4.25, 4.6, and 5 meters respectively. We also saw that in all reference point configurations, there is increased error when few packets are sent. This is due to the combination of packets being dropped and spurious RSSI readings. We see that the location error stabalizes when the RSSI of 50 or more packets are averaged.

As the density of beacon nodes decreases, the error also increases. With beacon node densities of 1 beacon every 66, 100 and 150 sq meters, we saw a 50^{th} percentile error of 1.5, 4.2 and 6.3 meters respectively. This confirms the intuitions that increasing node density will reduce error since there is (a) more RSSI information and (b) less distance between the mobile node and the beacons.

Using multiple power levels in our environment was problematic since at levels below full transmit strength, many locations were not able to receive packets from the mobile node. In cases when the signal was weak, the increased packet loss increased tracking error. Also, given our current radio hardware, there is little difference between transmitting at low power and high power with respect to energy consumption.

Figure 3 shows the error distribution of test-points using

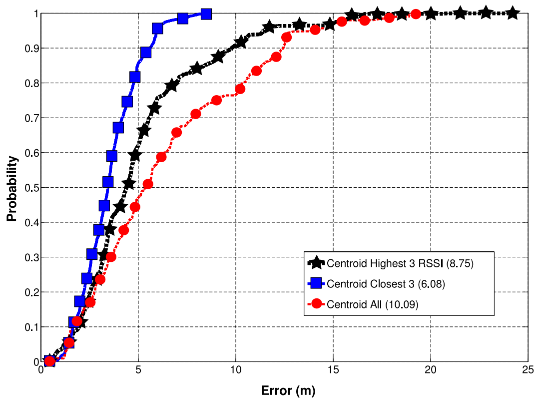

Fig. 3. Cumulative distribution of error using the weighted centroid approach with various beacon selection schemes. The three highest RSSI line shows the error when only the highest three RSSI valued beacons are used for the weighted centroid. The closest three line shows what would happen if an oracle informed the system of the closest three beacons to the mobile node. This shows that extra information about beacon selection can improve worst case error by more than a factor of 2. Total error shown in parentheses (x10^4)

variants of the weighted-centroid approach. The "centroid all" (right-most) line corresponds to the performance when all receiving beacons are used in the centroid algorithm to localize the mobile node. The "centroid highest three" (middle) line shows the performance of the centroid algorithm if the highest three RSSI-valued beacons are used for the localization. We see a 30% increase in performance at the 80^{th} percentile, but both ends of the distribution tend to follow the "centroid all" line. The left-most line shows what the performance would be if an oracle told the mobile node the physically closest three beacons that should be used for localization. We see that if the mobile node knows the correct beacons to use for the centroid calculation, then the results are nearly twice as good as the worst case. In Section V-E we show that this information in part can be supplied using micro-climate sensor data correlation.

D. Micro-climate Results

Figure 4 shows a map of temperature, humidity and light samples taken across our test-bed. Though they are not time-synchronized (since this would require an enormous density of nodes), these images do show that there is significant variation across a seemingly homogeneous environment. Many of the lighter regions found in the images correspond to artifacts in the environment such as air-conditioning vents or large glass windows.

Next, we evaluated the performance of our frequency features. Figure 5 shows an example of an audio signal being compared between a mobile node and three different beacon locations using the method described in Section IV-A. The top row of graphs shows the normalized spectral power density of the mobile node's audio sample. The middle row shows the normalized spectral density of audio samples at three different locations sampled in time within $10\mu s$ of one another. The bottom row of the graph shows the difference between the upper two signals. The sum

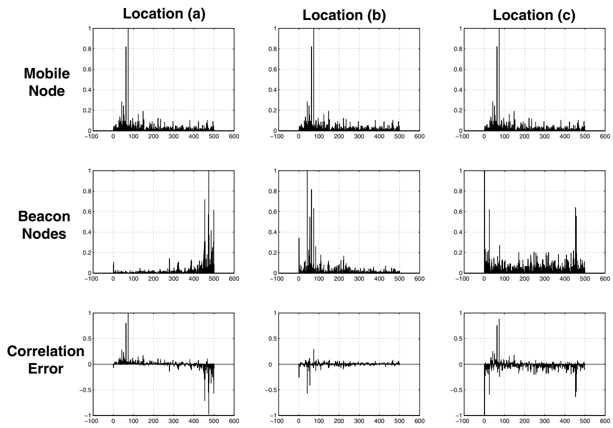

Fig. 5. This figure shows a sample set of the spectral energy densities of 4 different time synchronized audio recordings. Each column shows the mobile node's data on the top as compared to the beacon at a particular physical location in the second row. The bottom row shows the difference between the top and middle signals. We present the signals seen by three beacons at three different locations. In this example, the mobile node was closest to location b, hence the least error.

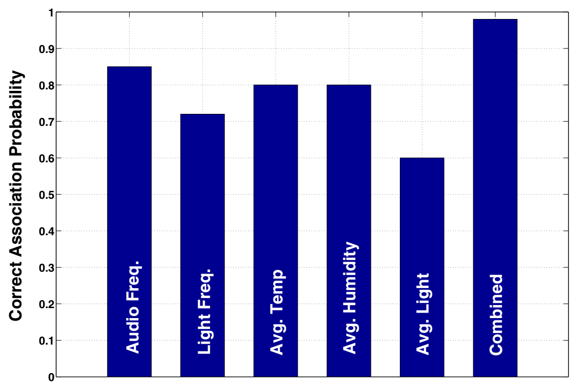

Fig. 6. Performance of various features in determining the closest beacon to the mobile node. The performance of the fused sensor data is shown in the last bar.

of all of the bars in the bottom corresponds to the total correlation error. In this example, the mobile node was nearest to location (b) which is supported by its low correlation error. Though simplisitic, this approach is both effective and computationally inexpensive.

We then performed a test where we moved a mobile node throughout the environment and recorded synchronized features for each known location. We then compared how well the correlation matched the closest beacon to the mobile node's location. Figure 6 shows the overall performance of the different features. We see that the audio performs best and hence should be weighted more aggressively in biasing RSSI values. We also note that when frequency components have a strong correlation, they are nearly always correct. The "combined" category is the performance of using all of the features combined with

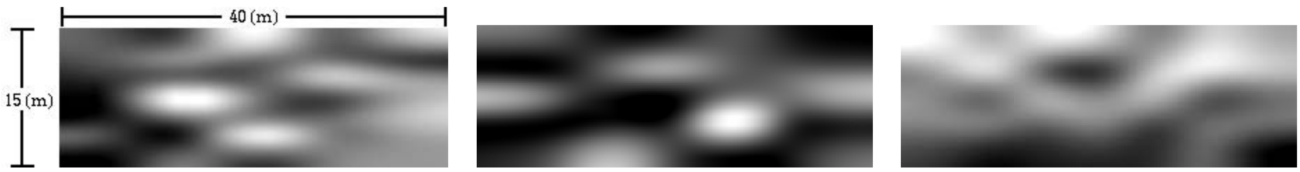

(a) Temperature (b) Humidity (c) Light Intensity

Fig. 4. These plots show temperature, humidity and light normalized across our testbed. We see that in seemingly homogenous environments, there still exists significant variation. Many of the intense spots in the image are positioned around air vents found on the floor of the building. Lighter colors represent higher temperature, humidity and light levels.

equal weighting. If all of the sensors were completely independent of each other, this would perform at 99.98% accuracy. We observed an actual performance of 98% accuracy which we attribute to slight mutual dependencies among sensor values measuring micro-climates.

E. Localization with Micro-climate Sensing

We now discuss the performance of our centroid-weighting scheme in conjunction with the micro-climate data. Figure 7 shows the overall comparison of the performance of the signature database localization, the centroid localization and the centroid augmented with micro-climate data approaches. The signature-based approach performs extremely well up to the 30^{th} percentile. The standard centroid approach has the worst overall performance. The centroid with micro-climate data is nearly twice as good as the standard centroid approach and outperforms the signature database approach significantly. In specific, the worst-case error using micro-climate sensing is bounded by 8 meters, while the errors from the reference signature and standard centroid approaches go as high as 20 meters. This ability to reduce the worst-case error can be used to very accurately provide room-level resolution.

F. Localization Performance Over Time

One of the primary problems with any radio communications environment is the changing nature of the wireless channel. In particular, the performance of the site survey approach is very susceptible to such changes. Figure 8 shows the effect of environmental change over time on the localization schemes we studied. The signature data taken from our original site survey were cross-correlated with data from the second site survey taken 1 month later. The second site survey was done during working hours, but when correlated with itself was still able to perform well. However, when we emulate changes and have the correlation done with the second site survey based on training data from the first survey, there can be an 100% increase in error. This is not surprising since the signature databases captures idiosyncrasies in the environment, and the environment likely suffered at least some minor changes. By comparison, our weighted-centroid approach using micro-climate sensing performed within 2% of the original total cumulative error between the first and second surveys. In other words, the micro-climate sensing approach is significantly more impervious to changes, since it utilizes correlations carried out at actual localization times.

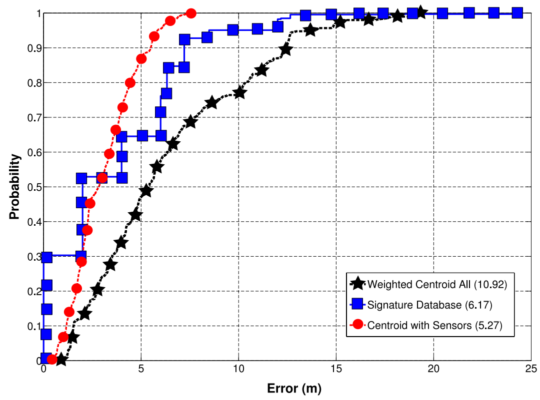

Fig. 7. Error vs Probability of 3 different localization schemes. Total cumulative error shown in parentheses (x10^4).

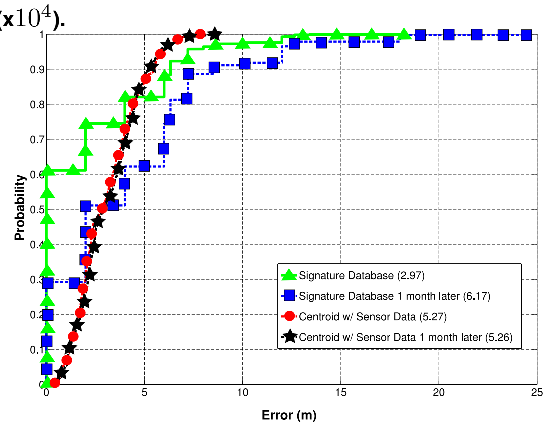

Fig. 8. Cumulative distribution error of the original signature database and then the original signature database against test points collected 1 month later. Total cumulative error shown in parentheses (x10^4). The centroid weighting scheme with sensor data shows almost no change in performance over time.

VI. Concluding Remarks

In this paper, we introduced sensor-based micro-climates to assist conventional RF-based location tracking. Our system is intended for people-tracking applications in dynamic environments where conventional RF localization suffers. We see that conventional signature site survey-based approaches deteriorate over time if the environment changes. We also show that a survey-free weighted-centroid approach performs better given additional information about which beacons are closest to the mobile node. The centroid scheme is easy to deploy, but suffers a worst-case error in our tests of more than 20 meters. We introduce the notion of micro-climate sensing in conjunction with the weighted-centroid scheme. We found that the worst-case error reduced to less than 8 meters. In our office environment tests, with subtle changes over time, the centroid with micro-climate data approach outperforms the signature database approach significantly in the worst case. Given that many buildings and structures may already have sensors for environmental monitoring, our approach is very attractive to increase localization accuracy. Our future plans include the study of adaptive weighting techniques under a variety of configurations.

References

[1] Maurer, U., Rowe, A., Smailagic, A., Siewiorek, D. eWatch: A Wearable Sensor and Notification Platform. *IEEE Workshop on Wearable and Implantable Body Sensor Networks*, 2006.

[2] Paramvir Bahl and Venkata N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *INFOCOM (2)*, pages 775–784, 2000.

[3] Want, A., Jones, A., Hopper, A. A New Location Technique for the Active Office. *IEEE Personal Comm.*, 1997.

[4] Priyantha, N.B., Chakraborty, A., Balakrishnan, H. The Cricket Location-Support System. *MOBICOM*, 2001.

[5] Krumm, J., Williams, L., Smith G. SmartMoveX on a Graph - An Inexpensive Active Badge Tracker. *UbiComp*, 2002.

[6] Christ, T., Godwin, P. A Prison Guard Duress Alarm Location System. *IEEE ICCST*, 1993.

[7] Smailagic, A., Small, J., Siewiorek, D. P. Determining User Lcation For Context Aware Computing Through the Use of a Wireless LAN infrastructure. 2000.

[8] Youssef, M., Agrawala, A., Shankar, A.U. WLAN Location determination via clustering and probability distributions. *IEEE PerCom*, 2003.

[9] K. Lorincz ,M. Welsh. A robust, decentralized approach to rf-based location tracking. Technical Report TR-04-04, Harvard University, 2004.

[10] Bulusu, N., Heidemann, J., Estrin, D. GPS-less Low Cost Outdoor Localization for Very Small Devices. *IEEE PerCom*, 2000.

[11] Bulusu, N., Heidemann, J., Estrin, D. Density Adaptive Algorithms for Beacon Placement in Wireless Sensor Networks. *IEEE ICDCS*, 2001.

[12] Kanade, T., Okutomi, M., Nakahara, T. A multiple-basline stereo method. *DARPA Image Understanding Workshop*, 1992.

[13] Rungsarityotin, W., Starner, T. Finding location using omni-directional video on a wearable computing platform. *IEEE ISWC*, 2000.

[14] Anand Eswaran and Anthony Rowe and Raj Rajkumar. Nano-RK: an Energy-aware Resource-centric RTOS for Sensor Networks. *IEEE Real-Time Systems Symposium*, 2005.

[15] Anthony Rowe and Rahul Mangharam and Raj Rajkumar. RT-Link: A Time-Synchronized Link Protocol for Energy-Constrained Multi-hop Wireless Networks. *SECON*, 2006.