

A Measurement Study of Zigbee-based Indoor Localization Systems Under RF Interference

Seng-Yong Lau^a, Tsung-Han Lin^b, Te-Yuan Huang^c, I-Hei Ng^a, and Polly Huang^a

^aDepartment of Electrical Engineering, National Taiwan University, Taiwan

^bSchool of Engineering and Applied Sciences, Harvard University, US

^cComputer Science Department, Stanford University, US

ABSTRACT

With an expected market value of \$2.71 billion in 2016, supporting daily use of real-time location systems in households and commercial buildings is an increasingly important subject of study. A growing problem in providing robust indoor location estimations in real time is the use of wireless transmissions in RF frequencies in daily environments. Having implemented a simple RSSI-signature-based location system on a 24-node IEEE 802.15.4-based sensor network testbed, we are able to analyze the effect of background IEEE 802.11 traffic on localization error. The measurement results demonstrate that the 80th-percentile of the localization error may increase by 141% when the background 802.11 traffic is high. Such performance degradation results from RSSI reading loss as the beacon messages collide with background traffic.

Categories and Subject Descriptors

C2.4 [Computer-Communication Networks]: Distributed Systems; C3 [Special-Purpose and Application-Based Systems]: Real-time and Embedded Systems

General Terms

Algorithms, Design, Experimentation, Measurement

Keywords

Indoor Localization, Frequency Hopping, Interference, Co-existence

1. INTRODUCTION

The market for real-time location systems for assets and personnel tracking is expected to reach \$1.26 billion by 2011 [3], and \$2.71 billion in 2016 [1]. For widespread adoption and everyday use of real-time location systems in households and commercial buildings, the systems must be able to provide accurate and stable location estimations with little delay.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

Most indoor localization systems employ an RSSI-signature-based approach which exploits temporal stability in the received signal strength indication (RSSI) from a set of pre-deployed beacons at identified locations, which is referred to as the RSSI signature. When a target carrying a receiving tag enters the space, the received RSSI values are compared to the RSSI signatures. The location of the target is identified by the corresponding location of the closest RSSI signature. To tackle the variation of RSSI signatures, methods of ensuring robust mapping between the measured RSSI values and the pre-recorded RSSI signature have been studied intensively in recent years [8][11][23]. Such methods are needed to minimize localization errors induced by unstable RSSI values.

An often overlooked problem is the increasing use of wireless transmission of RF frequencies in everyday environments. Bluetooth (IEEE 802.15.3), WiFi (IEEE 802.11) and Zigbee (IEEE 802.15.4) all operate in the 2.4x GHz frequency band. The stability and availability of RSSI information for WiFi- or Zigbee-based localization systems may vary depending on interference from other WiFi, Bluetooth, and Zigbee sources.

After conducting a systematic set of experiments on a Zigbee-based sensor network testbed, we find that the 80th-percentile error of a simple RSSI-signature-based localization system may increase from 1.6 to 3.9 meters when the amount of background WiFi traffic increases from 68 to 2835 kbps. Having measured also the amount of WiFi traffic in our department building, we observe that there is a significant amount of time that the localization accuracy may suffer from the bursts of background noise. In a detailed analysis, we discover that the degradation in localization accuracy is mainly contributed by loss of beacon messages, rather than the variance of RSSI values. This agrees with previous studies that discovered variance in RSSI values is mainly due to the multi-path effect [2][26]. Background traffic does not add to the multi-path effect, rather causing the beacon messages to drop.

This study makes the following two contributions:

- The unique architecture of the proposed sensor network testbed enables low cost co-collection of data traces at the beacon nodes and the receiving tags.
- The systematic measurement study provides an understanding on the effect of background traffic to indoor RSSI-signature-based location systems.

The rest of the paper is organized as follows. First, the

RSSI-signature-based localization system and our testbed are detailed. Sections 4 and 5 describe the measurement methodology and the analysis of the effect of background WiFi traffic on the localization system. We give the implication of our analysis towards a robust localization system in Section 6.

2. LOCALIZATION SYSTEM

This study implements an RSSI-signature-based localization system. The underlying concept of this solution is to exploit the mapping between a tag’s location and RSSI values of packets received from pre-deployed beacons. The RSSI set is referred to as the RSSI signature or vector. These systems typically operate in two phases, training and tracking phases. In the training phase, the area is surveyed to construct the reference RSSI signature per sampled location. The collective set of RSSI signatures obtained at various locations is referred to as the radio map.

Using the radio map, the system compares the collected RSSI vector to the reference RSSI signatures in the tracking phase to identify the closest possible location. The system employs the k-nearest-neighbor (KNN) method for location inference. The k sample locations with RSSI signatures closest to the collected RSSI vector are selected. The KNN estimator then outputs a location by averaging the coordinates of the top k locations weighted by the distances between the RSSI vector and the signature.

2.1 Beacon

The beacons periodically transmit short packets containing the beacon ID. The packet sending interval is set to 200ms. The radio transmission power is set to -7dBm. Thus, the tag can detect nine or ten beacons at every location. Because the beacon packets are the basis for the RSSI readings, successful delivery of the packets is critical to the performance of the localization system. To avoid packet collisions among the beacons, the DESYNC[6] protocol is implemented. The protocol ensures that neighboring beacons have different sending time to avoid collisions.

2.2 Training Phase

In the training phase, the RSSI signature map is constructed. The survey area is divided into grids, approximately 30cm apart, which is about the distance of one step. During the survey phase, a receiving tag is connected to a portable PC held by the user, who then walks along the corridor. The user must wait at each grid for 8 seconds until the beacon packets are received. After collecting forty RSSI vectors, the received RSSI vectors are averaged to generate a single RSSI signature vector for each sampled location.

2.3 Tracking Phase and the KNN Estimator

In the tracking phase, the receiving tag collects the beacon packets for 220ms and sends the RSSI vector back to the localization system. The system then compares the received RSSI vector with the signature map to find the closest possible location. The KNN method is then used to find the k locations with the closest signature and computed the weighted average of the k location. In this study, the value of k was set to 3.

The signature distance employed is the *normalized Euclidean distance*. Restated, the Euclidean distance between two RSSI vectors is further divided by the number of bea-

Figure 1: Testbed Layout of 24 Beacon Motes. The beacon motes are all wired via USB to central PCs to simplify debugging. The signature maps are built along the training area denoted in blue.

cons with significant values in the RSSI vector. This step is necessary because beacon packet loss may be due to geographic distance as well as signal instability or collisions. The semantics of packet loss are ambiguous. Simply using the lowest RSSI value for a lost beacon packet introduces bias and estimation error. To avoid the ambiguity and bias, missing values are simply disregarded by normalizing the Euclidean distance.

3. TESTBED

The testbed served as the platform for measurement study. Twenty-four beacon nodes were deployed on the 6th floor of a department building at this university. The beacon nodes are telos-like modules [20] equipped with TI MSP430 microcontrollers and CC2420 802.15.4 radio. The software is implemented on TinyOS, and the default MAC, a CSMA/CA-like mechanism, is on for all beacon packet transmissions.

Figure 1 shows the floor plan. The smaller rooms, numbered 611 to 629, are faculty offices and the remaining are graduate assistant laboratories. The twenty-four Telos-like beacon nodes are small boxes distributed evenly along the corridor. To simplify testbed debugging, every beacon node was connected via USB to one of the two testbed PCs. The PCs were installed in room 621 and room 613, and each was connected to twelve beacon nodes. The PCs served as gateways to allow easy code upgrades and data logging via USB. MoteLab [25] shares the same wired setting, but this testbed uses lower cost off-the-shelf components for long distance USB connectivity. The testbed was designed to use AC power battery replacement and to enable long-term measurements.

3.1 USB Connectivity

The effective transmission distance of the standard USB interface was about 5 meters. For nodes more than 5 meters apart, a USB extender [14] was used. The USB extender is an off-the-shelf product that extends the effective transmission distance by up to 45 meters. A local unit on the extender modulates the USB input to signals that transmit on any standard CAT5e network cable. At the other end of the CAT5e network cable is the remote unit which demodulates the signal back to the USB format.

To avoid deploying numerous lengthy wires throughout

Figure 2: Testbed Wiring. (a) Chaining scheme to connect beacon nodes. USB extenders and USB hubs are used to avoid lengthy wires. (b) The power connection (VCC) between USB extender and USB hub is cut to ensure the USB hub only draws power from the AC power adapter.

the building, the chaining scheme in Figure 2(a) was used to connect nearby beacon nodes. At the beginning of the chain, a 4-port USB hub was connected to the gateway via a USB extender. Three beacon nodes were directly connected to this USB hub, and the next USB hub in the chain was connected via another USB extender. This chain was continued until it reached the maximum range of the USB extender. That is, the distance between the last USB hub and the PC did not exceed forty-five meters. This limitation is also why two PCs were required to cover the entire deployment area in the testbed.

3.2 Power Supply

Powering the beacon nodes and the USB devices was a challenging problem. The peak current consumption of the beacon node, the USB hub and the USB extender was approximately 60mA, 5mA, and 20mA, respectively. Therefore, a single chain would consume more than 800mA current. Sourcing power from the PC USB port was not feasible due to the high current requirement.

Instead, the USB hubs for the system were externally powered. As Figure 2(a) shows, each USB hub on the chain was connected to an AC power adapter providing maximum of 3 amperes of currents. Due to the excessive length of the beacon chains deployed and the large current consumption, the voltage at the last USB hub would have dropped by as much as 1 volt. Hence, although the standard supply voltage for USB is 5 volts, a 6 volt power adapter was used for the external power source to compensate for the voltage drop. Every device in the testbed was tested to ensure it could sustain the 6 volt power supply. The final step is to cut the power connection pin (VCC) between the USB hub and the USB extender remote unit (Fig. 2(b)), to ensure that the USB hub was drawing power solely from the external power source.

4. MEASUREMENT METHODOLOGY

The RSSI-based localization system is vulnerable to environmental noises. In a typical office or campus environment, background noise could be from WiFi, Bluetooth, a 2.4 GHz cordless telephone, a microwave-oven or other RF devices operating on a 2.4 GHz ISM band. Among these, WiFi traffic produces a significant amount of interference,

especially in an office environment. To determine the effect of WiFi noise on localization accuracy and the efficiency of the proposed frequency hopping mechanism, WiFi traffic was generated at different levels. As the WiFi traffic was transmitted at different levels, the following data were collected: (1) background WiFi traffic, (2) beacon messages received at the receiving tag and (3) beacon messages received at other beacons.

4.1 Location of Measurement

Because the generated WiFi traffic would be traveling in the space between the source and the access point (AP), the wireless LAN for generating traffic and the location of taking the beacon measurements were carefully selected. A survey of the 6th floor of the department building revealed more than ten APs. Six APs located on the ceiling of a corridor had been installed by the university to provide general wireless Internet access for staff and students. Others were deployed by individual laboratories and had restricted access. One of the generally accessible AP on the ceiling of a corridor was selected for testing. The localization testbed was set to operate on the channel that overlaps with the WiFi channel used by the AP. The receiving tag was positioned close to the AP.

4.2 Data Logging

To generate WiFi traffic at different levels, a laptop PC was connected to the internet via the selected AP and a large file was downloaded from an FTP server using FlashFXP, an FTP client that allows the user to set the upload/download speed limit.

Another laptop PC near the WiFi source was used as a sniffer to log all WiFi traffic in the channel and to ensure no unexpected extra traffic occurred. Dumpcap [7], a Linux packet header capturing tool built on the pcap library, was used to log the packets. The WiFi log, referred to as 'WiFi' data, was used to measure the background noise. The experiments were conducted in midnight during weekends. Only a very small amount of traffic other than the generated one was observed by the sniffer. This ensured that the interference patterns observed were coming from the same AP. In the five sets of experiments, the average background WiFi traffic rates were 68, 264, 1308, 1705, and 2835 kbps. The traffic rate reported here is not exactly the same as the traffic in the channel because the sniffer cannot capture corrupted packets. Also, no kernel loss is reported by the packet capturing tool.

The receiving tag was connected to another laptop PC through the USB interface. Each rate-limited file transfer session was slightly longer than 10 minutes. During that period, the receiving tag transferred all beacon messages received through the USB interface to the laptop PC. Similarly, all beacon nodes in the testbed passed the beacon messages received through the USB interface to the gateway PCs. Table 1 provides detailed information about the data sets.

The RSSI values in the beacon messages collected at the receiving tag were used to infer the localization errors. From the beacon messages collected within a beacon cycle (0.2 seconds), the location of the receiving tag was estimated using the mechanism described in Section 2. Each run produced 3000 location estimations. The beacon messages collected at the neighboring beacons were used to calculate the beacon

Table 1: Details of the Data Sets for 5 Different WiFi Data Rates.

Test Case (WiFi Data Rate)	68 kbps	264 kbps	1308 kbps	1705 kbps	2835 kbps
WiFi Log (pkt/MB)	48789/4.96	74257/19.34	211008/95.80	247684/124.88	450176/207.62
Beacon (pkt/MB)	530483/11.13	519557/10.90	455629/9.56	445589/9.35	368411/7.73
Beacon Length Error (pkt/Bytes)	21/462	16/352	29/638	28/616	25/550

packet reception rate (PRR). The beacon packet reception rate was calculated by zooming into the beacon link near the AP where the generated background traffic emerged. Taking a sliding window of 50, the percentage of beacon packets received in the past 50 cycles was calculated.

The WiFi and RSSI data sets were later used to study the effect of WiFi background traffic to the localization error. The RSSI and PRR data sets were used to observe the correlation between the beacon message reception rate at the neighboring beacons and the localization error to determine the optimal design of the frequency hopping mechanism.

4.3 Loss in Beacon Message Logging

Software running on the sensor nodes as well as the PCs collecting data through the serial port was prone to errors. During the course of the experiments, several errors in the data sets were identified. Most were software bugs and were quickly corrected. The remaining errors were caused by hardware and communication problems. To ensure that traces were not contaminated by software bugs and to accurately assess the quality of the traces, three error checks were implemented to identify hardware and communication problems.

(1) Message length check. Every beacon message generated had a fixed and identical length. For unknown reasons, the CC2420 radio stack sometimes received a valid packet with an altered packet length in high contention. This problem was also noted in the TinyOS CC2420 radio module file, but no fix is currently available. The incorrect length field could be longer or shorter than normal. If the length field reported a longer value, the packet payload would still be correct but with garbage bytes appended for the extra length. However, the RSSI reading would be invalid, and would usually show 0xFF. This problem, although not critical to other uses of sensor packets, is problematic for RSSI-based localization systems. Simply recording the RSSI value without checking produces erroneous location estimations. To correct this problem, the number of bytes in the packet were verified and packets with incorrect packet length were dropped.

(2) Serial error check. To capture bit errors during serial transmission, a 16-bit CRC checksum value was appended to each packet logged. Packets failing the checksum were discarded by the serial listener. In addition to checksum, a serial sequence number was also appended to each logged packet to check for possible serial loss. The serial listener determined the amount of packet loss from the sequence number. Throughout the experiments, no checksum failure or packet loss was reported. This also showed that wiring the sensor nodes to a central PC was effective for the measurement study.

4.4 Beacon Message Synchronization

The beacon sequence number in the beacon messages was used to synchronize the beacon traces. Upon receiving the beacon messages, the other beacon nodes and the receiving

Figure 3: Distribution of Localization Errors Under Different WiFi Data Rate. The localization error is larger when the WiFi data rate is higher.

tag time stamped the messages using their local clocks. Assuming that the time required for the beacon messages to travel one hop to the receiving tag was the same as that required to travel to the neighboring beacon the traces were synchronized by simplified Jigsaw approach [5].

More specifically, the local clock of the receiving tag was used as the global clock. Let t_m represent the timestamp of the reference packet received at the mobile tag and let t_k denote the timestamp of the reference packet received at the k_{th} beacon nodes. The local clock of the k_{th} beacon node would then adjusted by adding the time offset $t_m - t_k$. Since the testbed was a multi-hop network, no reference packet could be received by any beacons in the network. A queue of reference packets was implemented to transitively synchronize other beacon nodes not receiving the previous reference packets. The first packet received by the mobile tag was chosen as the first reference packet. Once a beacon was synchronized, the next packet it received/sent was added to the queue. Elements in the queue were be popped out sequentially until all beacons were synchronized.

5. TRACE ANALYSIS

We analyze the five sets of traces collected from the localization testbed with different levels of WiFi traffic in the background.

5.1 Localization Errors

Figure 3 depicts the cumulative distribution function (CDF) of the localization errors. The localization accuracy is pretty good with 50th percentile error 53cm. We believe such good accuracy comes from the following reasons. First, DESYNC is applied on beacons. Collisions are thus reduced, and the receiving tag can receive sufficient RSSI readings to give accurate location estimation. Second, the survey and the test conditions, e.g. antenna orientation and the way the receiving tag is wore, are held the same throughout the ex-

(a) Correlation between background traffic rates, number of beacon packets received and localization errors (b) Effect of beacon message loss (c) Distribution of the number of beacon messages

Figure 4: Impact of Beacon Packet Loss. Higher background traffic rate results in a larger localization error and a smaller amount of beacon packets received. Insufficient amount of beacon packets is the main cause of larger localization.

periments. The beacon density of the system is also high, with a beacon placed every five meters.

The test results showed that the localization errors were influenced by the background WiFi traffic. In the 50th percentile, the errors increased from 53cm to 81cm (53% increase) as the background WiFi traffic increased. The increase in the 80th percentile error from 160cm to 385cm (141% increase) was particularly large. This indicated that, as background traffic increases, the localization error and variance also increase. In cases of heavy background traffic, all beacon messages may be corrupted in some cycles. The localization error was set to a pre-defined maximum for analysis of these cases. In practice, the system can predict or simply report the location obtained in the previous cycle.

5.2 Beacon Message Losses

To understand how the background traffic impacts localization accuracy, we first look into the loss of beacon messages. In Figure 4(a), the left x-axis shows the average number of beacon message received during each 220-ms interval, and the right x-axis plots the 80-percentile localization error. It can be seen that the average number of beacon message received goes down in high background traffic rate and shows a strong correlation with the increasing localization error.

To further clarify the impact of beacon message loss, receiving cycles were classified by the number of beacon messages received for each trace in Fig. 4(b). The corresponding average localization error and the variance shown in the figure indicate that fewer received beacon messages increase localization error and variation. A location estimation based on only one or two beacon RSSI readings would be very imprecise. Average errors would be as high as 9 meters since several sample signatures share similar RSSI readings for a single beacon. The insufficient information received due to the fewer beacon messages would cause larger localization errors.

Figure 4(c) shows the probability of the tag receiving a specific number of beacon messages for different background traffic rates. The distribution of the number of received beacon messages indicates that high background traffic increase the number of cycles in which the receiving tag observes only a small number of beacon messages. In fact, the background

traffic interferes with the delivery of beacon messages and corrupts them. Thus, the overall number as well as variance in localization errors worsens when background traffic is high.

5.3 Beacon Impact

Next, we investigate when the number of beacon message loss is the same, does it matter to have a specific beacon message received? In other words, do all beacons influence the localization accuracy equally? To understand the effect of individual beacon message losses, we measure the importance of each beacon by defining the *beacon impact*. We take measurements where there is a fixed number of beacon message loss. For every beacon, we separate the data into two sets. In one set the message from the beacon is observed, and in the other set the message from the beacon is missing. The *beacon impact* of the particular beacon is defined as the difference of the average localization error calculated between the *missing* set and the *observed* set.

Intuitively, if the beacon is important, missing its RSSI value would result in a larger localization error. The beacon would have a positive beacon impact value, and we called that the beacon is desired. On the other hand, if the beacon impact is negative, including the beacon RSSI may increase the overall localization error and we said the beacon is undesired.

In our data we found that the beacons that are closer to the receiving tag are more desired and the beacons that are farther away are undesired. Figure 5 show the beacon impact values of six different beacons. Beacon 23, 24, and 25 that have higher RSSI values are closer beacons, while beacon 13, 14, and 29 are farther. We show the beacon impact value when the number of beacon messages loss is 3, 5, 7, and 9.

It can be seen that it is desired to have closer beacons. Having the closer beacons shows positive impact on the localization accuracy. If the message from a desired beacon is received, the localization error can be reduced by up to 5 meters, whereas the error is increased by 6 meters if an undesired beacon message is received. In addition, a larger amount of beacon messages shows more prominent positive or negative impact on localization accuracy, since when the

Figure 5: Impact of Different Beacons. Beacons closer to the receiving tag (with larger RSSI) are more important to the localization system. The beacon impact values for closer beacons are positive indicates that observing their RSSI values shows smaller localization error than those missing the RSSI.

information is insufficient, the quality of information becomes more important. The results suggest that weaker RSSI values are more ambiguous for determining locations while larger RSSI values are important for more accurate localization.

5.4 RSSI Values

Note that in Fig. 4(b), if the receiving tag manages to receive sufficient beacon messages, the localization errors are all similar for different background traffic rates. This suggests that produces less distortion of RSSI readings. Figure 6 shows the average RSSI readings and the standard deviation for each trace. The RSSI values from different beacons are slid slightly for clarity and ease of comparison. The deviation of RSSI readings again reveals no a clear trend as background traffic increases. Generally, RSSI variance cause some localization error. However, background WiFi traffic apparently has no significant effect on RSSI variation. As Fig. 3 shows, the overall increase in the magnitude of errors is mainly due to the beacon message losses caused by background traffic.

6. TOWARDS ROBUST INDOOR LOCALIZATION

Our experimental results in Section 5 indicate that the localization system is susceptible to WiFi interference. A larger amount of background WiFi traffic may cause bigger localization errors. We found that the localization errors results from the loss of beacon messages, and the loss increases as the WiFi traffic becomes more intensive. This is because the WiFi traffic might collide with beacon packets in the air, and therefore the localization system cannot obtain enough information for quality estimation. Note that our localization system already uses the normalized KNN(NKNN) for signature matching, which tries to mitigate the effect of in-

Figure 6: Effect of Background Traffic to RSSI Readings. Varing the background traffic rates does not have significant effects on beacon RSSI readings. The variation of RSSI is thus not the main cause of lareger localization error in higher background traf- fic.

Figure 7: Packet Reception Rate(PRR) vs Localization Errors. This figure shows the PRR on a specific link between two beacons. It can be seen that there is correlation between link PRR and localization error. Link PRR can potentially be a indicator on the background interference level.

Figure 8: WiFi Traffic Load. The WiFi traffic shown is measured during a busy working hour. The data rate exceeds 1 Mbps for substantial time periods, for which the localization system shows performance degradation.

sufficient information. However, our results show that the interference from WiFi traffic is so severe that the system still has a poor performance even with NKNN.

Figure 7 shows the relationship between the beacon packet reception rate under the influence of five different WiFi traffic rates and their resulting localization errors. From the figure, we can see that when the level of background traffic decreases, the packet reception rate rises and the corresponding localization error is then reduced. Since the impact of WiFi traffic is so significant, a robust location system should have the ability to evaluate the quality of each link. To do so, we can utilize the fact that the number of beacon message loss from a certain beacon to a receiving tag is also observable from the neighboring beacons. Thus, each beacon node could track the packet reception rates from neighboring beacons. Then, from the packet reception rate observed at the neighboring nodes, the beacon node could be aware of the existence of background wireless traffic.

We also observed sufficient amount of WiFi traffics in our department building as those used in Figure 7. Figure 8 shows WiFi data rate during a busy working hour at this department building. For more than one third of the observation period, the data rate exceeds 1 Mbps. Such a bursty pattern of traffic has also been reported in similar works elsewhere [21][12]. The 80th-percentile error during these periods may be as high as 2 or even 4 meters, while the error is usually lower than 1 meter. Worse, the localization system may sometimes temporarily blackout when WiFi traffic is high. Even under an ordinary amount of background WiFi traffics, the localization system might suddenly perform poorly and give unacceptable location estimation. This again shows that for a robust localization system, handling the coexistence of wireless signals is crucial. This suggests that for a RSSI-based localization system to perform stably with accuracy, it might be necessary for it to operate at a different frequency channel to avoid heavy WiFi interference. Dynamic control of frequency spectrum can be important for a robust indoor localization system that provides consistent and accurate location estimations.

Furthermore, our analysis in Section 5.3 shows that the closer beacons have greater influence on localization accuracy, especially when the information is incomplete. Thus, when deploying a robust location system, one should bear in mind that the density of beacon nodes would have great impact on location estimation. Also, when calculating locations, the system should give higher priority for the information provided by the closer beacons.

7. RELATED WORK

Co-channel interference is not a new problem, and channel hopping is a popular solution. The problem exists whether or not radio standards are similar. Srinivasan *et al.* [24] and Zhou *et al.* [27] have observed the impact of interference between 802.11 and Zigbee signals. Recently Gummadi *et al.* [9] analyzed the interference of Zigbee and cordless phone using 802.11 networks for data communication.

Indoor localization techniques can be divided into two categories: range-free and range-based. The former does not localize targets based on range estimation [10]. The latter uses absolute angle estimates [17] or point-to-point distance estimates by measuring signal propagation time [19][22], signal interference patterns [15] or received signal strength indicator (RSSI).

In particular, the range-based system that exploits the measured RSSIs infers the location either based on the decay model of distance to signal strength [18][16][13] or the RSSI fingerprints (also known as the RSSI signatures) to relate an observed set of signal strengths to ones at known locations [4]. The RSSI-signature-based techniques are broadly categorized into deterministic and probabilistic techniques. For deterministic method, RSSI readings are first collected at known locations and the location is the one which RSSI signature is closest to the measured RSSIs. Probabilistic methods, on the other hand, construct a probability distribution over the target's location. The extended Kalman filter [?] or particle filter [11] [23] based approach, though can reduce estimation error, requires predictive mobility model which is not obtainable without aid of other sensors.

8. CONCLUSION AND FUTURE WORKS

In this study, we build an indoor sensor network testbed and implement a RSSI-signature-based localization system on top of it. The design of the testbed to have both wired and wireless interface on every sensor node facilitates the measurements. In addition, the testbed is built with off-the-shelf hardware except the sensor nodes. Our measurement study shows that the performance of the localization system suffers from WiFi interference that the interference corrupts beacon packets. Without sufficient information, the system thus fails to provide quality estimation in a timely fashion. Our analysis also indicates that such interference can potentially be detected by monitoring the link quality between neighboring beacons. Switching the operating frequency of the system to a less interfered channel may be a viable solution. Future work includes to understand the behavior of other types of interference such as microwave oven, and design a frequency hopping mechanism for robust indoor localization.

9. REFERENCES

- [1] Real time locating systems 2006-2016 (RTLS). *GIExpress Market Research Report*, Jul 01, 2007. <http://www.giexpress.com/products/ix37643/>.
- [2] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *Proceedings of SIGCOMM '04*, pages 121–132, New York, NY, USA, 2004. ACM.
- [3] B. Bacheldor. RTLS market to grow 30 percent annually. *RFID Journal*, May 10, 2006. <http://www.rfidjournal.com/article/articleview/2325/1/1/>.
- [4] P. Bahl and V. Padmanabhan. An in building RF-based user location and tracking system. *IEEE INFOCOM 2001*, April 2001.
- [5] Y.-C. Cheng, J. Bellardo, P. Benkö, A. C. Snoeren, G. M. Voelker, and S. Savage. Jigsaw: solving the puzzle of enterprise 802.11 analysis. In *Proceedings of SIGCOMM '06*, pages 39–50, New York, NY, USA, 2006. ACM.
- [6] J. Degeys, I. Rose, A. Patel, and R. Nagpal. Desync: self-organizing desynchronization and TDMA on wireless sensor networks. In *Proceedings of IPSN '07*, pages 11–20, New York, NY, USA, 2007. ACM.
- [7] dumpcap man page. <http://www.wireshark.org/docs/man-pages/dumpcap.html>.

- [8] E. Elnahrawy, X. Li, and R. Martin. The limits of localization using signal strength: a comparative study. *Sensor and Ad Hoc Communications and Networks, 2004. IEEE SECON 2004. 2004 First Annual IEEE Communications Society Conference on*, pages 406–414, 4-7 Oct. 2004.
- [9] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan. Understanding and mitigating the impact of rf interference on 802.11 networks. In *Proceedings of SIGCOMM '07*, pages 385–396, New York, NY, USA, 2007. ACM.
- [10] T. He, C. Huang, B. M. Blum, J. A. Stankovic, and T. Abdelzaher. Range-free localization schemes for large scale sensor networks. In *Proceedings of MobiCom '03*, pages 81–95, New York, NY, USA, 2003. ACM.
- [11] J. Hightower and G. Borriello. Particle filters for location estimation in ubiquitous computing: A case study. In *Proceedings of Ubicomp'04*, pages 88–106, 2004.
- [12] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding link-layer behavior in highly congested ieee 802.11b wireless networks. In *Proceedings of E-WIND '05*, pages 11–16, New York, NY, USA, 2005. ACM.
- [13] K. Lorincz and M. Welsh. Motetrack: a robust, decentralized approach to RF-based location tracking. *Personal Ubiquitous Comput.*, 11(6):489–503, 2007.
- [14] C. F. I. LTD. USB extender. <http://www.cfi.com.tw>.
- [15] M. Maróti, P. Völgyesi, S. Dóra, B. Kusý, A. Nádas, Ákos Lédeczi, G. Balogh, and K. Molnár. Radio interferometric geolocation. In *Proceedings of SenSys '05*, pages 1–12, New York, NY, USA, 2005. ACM.
- [16] D. Niculescu. Positioning in ad hoc sensor networks. *IEEE Networks*, 18(4):24–29, July-August 2004.
- [17] D. Niculescu and B. Nath. Ad hoc positioning system (aps) using aoa. *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies. IEEE*, 3:1734–1743 vol.3, 30 March-3 April 2003.
- [18] N. Patwari. Relative location estimation in wireless sensor networks. *IEEE Transactions on Signal processing*, 51(8):2137–2148, August 2003.
- [19] N. Patwari, I. Hero, A.O., M. Perkins, N. Correal, and R. O’Dea. Relative location estimation in wireless sensor networks. *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, 51(8):2137–2148, Aug. 2003.
- [20] J. Polastre, R. Szewczyk, and D. Culler. Telos: enabling ultra-low power wireless research. In *Proceedings of IPSN '05*, page 48, Piscataway, NJ, USA, 2005. IEEE Press.
- [21] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan. Measurement-based characterization of 802.11 in a hotspot setting. In *Proceedings of E-WIND '05*, pages 5–10, New York, NY, USA, 2005. ACM.
- [22] A. Savvides, C.-C. Han, and M. B. Strivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of MobiCom '01*, pages 166–179, New York, NY, USA, 2001. ACM.
- [23] V. Seshadri, G. V. Zaruba, and M. Huber. A bayesian sampling approach to in-door localization of wireless devices using received signal strength indication. In *Proceedings of PERCOM '05*, pages 75–84, Washington, DC, USA, 2005. IEEE Computer Society.
- [24] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. Some implications of low-power wireless to ip routing. In *Proceedings of the Fifth Workshop on Hot Topics in Networks (HotNets V)*, Irvine, CA, Nov 2006.
- [25] G. Werner-Allen, P. Swieskowski, and M. Welsh. Motelab: a wireless sensor network testbed. In *Proceedings of IPSN '05*, page 68, Piscataway, NJ, USA, 2005. IEEE Press.
- [26] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proceedings of SenSys '03*, pages 1–13, New York, NY, USA, 2003. ACM.
- [27] G. Zhou, J. Stankovic, , and S. Son. Crowded spectrum in wireless sensor networks. In *Proceedings of the Third Workshop on Embedded Networked Sensors (EmNets 2006)*, Cambridge, MA, May 2006.