

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, Polly Huang^a

^aDepartment of Electrical Engineering, National Taiwan University, Taiwan

^bSchool of Engineering and Applied Sciences, Harvard University, MA, USA

^cComputer Science Department, Stanford University, CA, USA

sylau@ntu.edu.tw, thlin@eecs.harvard.edu, huangty@stanford.edu,

b91901152@ntu.edu.tw, phuang@cc.ee.ntu.edu.tw

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. 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% at worst in an office building with active use of IEEE 802.11 for data. 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, Interference, Coexistence

1. INTRODUCTION

The market for real-time location systems for assets and personnel tracking is expected to reach \$1.26 billion by 2011 [4], and \$2.71 billion in 2016 [2]. For widespread adoption and everyday use of real-time location systems in households and commercial buildings, it is required that the systems can

provide accurate and stable location estimations with little delay.

Most indoor localization systems employ an RSSI-signature-based approach which exploits temporal stability in the received signal strength indication (RSSI) of wireless signals. In that, at every known location, the RSSIs collected from a set of pre-deployed beacons form an RSSI signature for the corresponding location. When a target carrying a receiving tag enters the space, the RSSI values collected on the tag are compared to the RSSI signatures. The location of the target is identified by the corresponding location with the closest RSSI signature. To tackle the temporal variation of RSSI signatures, methods of ensuring robust mapping between the measured RSSI values and the pre-recorded RSSI signatures have been studied intensively in recent years [8, 12, 24].

An often overlooked problem is the increasing use of wireless transmission in RF frequencies in the urban environment. Bluetooth (IEEE 802.15.3), WiFi (IEEE 802.11) and Zigbee (IEEE 802.15.4) all operate in the 2.4x GHz frequency bands. The stability and availability of RSSI information for WiFi- or Zigbee-based localization systems may vary depending on the level of interference generated from other WiFi, Bluetooth, and Zigbee sources. Aiming at a better understanding on the effect of RF interference, we build a Zigbee-based sensor network testbed for RSSI-signature-based indoor localization.

After conducting a systematic set of experiments on the testbed, we find that the RSSI-signature-based localization system is susceptible to WiFi interference. The 80th-percentile error of the system may increase from 1.6 to 3.9 meters when the average 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 the 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 [3, 27]. Background traffic does not add to the multi-path effect, rather causing the beacon messages to drop. Ultimately, the study reveals the need of interference-resilient indoor localization mechanisms and ways towards effective detection and mitigation of RF interference.

This study makes the following two contributions:

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 republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WiNTECH'09, September 21, 2009, Beijing, China.

Copyright 2009 ACM 978-1-60558-740-0/09/09 ...\$10.00.

- 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.

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 the RSSI values of packets sent from pre-deployed beacons. The set of RSSIs collected at a single location is referred to as an RSSI signature or vector. Typically, these systems operate in two phases, the training and the tracking phases. In the training phase, the area is surveyed to construct the reference RSSI signature at every sampled location. The collective set of RSSI signatures obtained at various locations is referred to as the radio map.

Using the radio map, in the tracking phase, the system compares the collected RSSI vector to the reference RSSI signatures in order 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 the weighted centroid of the top k locations with the weights determined by the closeness between the RSSI vector and a given signature.

2.1 Beacon

The localization system has a set of pre-deployed beacons, periodically sending short packets containing their ID. We set the packet sending interval to 200ms, and the transmission power of the radio to -7dBm. In general, we observe that the receiving tag can hear 9 out of 24 deployed beacons at every location in our setup (see Section 3). Since the beacon packets are the basis for RSSI readings, successful delivery of the beacon packets is critical to the performance of the localization system. To avoid packet collisions among the beacons, i.e., to minimize the effect of self-interference, we implement the DESYNC[7] protocol. The protocol ensures that neighboring beacons send packets at different time, and thus, avoids packet collisions.

2.2 Training Phase

In the training phase, we measured the RSSI signature at every location. To define a location, the training area is divided into grids approximately 30cm apart, which is about the distance of one step. The measurement is done with a receiving tag connected to a portable PC held by a person, who then walks along the corridor. The person waits at each grid until 40 beacon packets are received. The received RSSI vectors are averaged to generate a single RSSI signature for each sampled location.

2.3 Tracking Phase and the KNN Estimator

In the tracking phase, the receiving tag collects the beacon packets and the RSSI readings along with them. Every 200ms, the receiving tag sends the RSSI vector collected from the past 220ms back to the localization system to determine its location. Every beacon is supposed to broadcast




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.

only once within the interval. We use such a short time interval in order to reduce the delay of the localization system. Note that this interval bounds the location update rate of the system. A person in average can move 26cm during the interval with velocity 1.3m/s, which is roughly the same as the location resolution of our system.

The system compares the received RSSI vector with the radio map to find the closest possible location. In that, we use the KNN estimator that extracts the top k locations with the closest RSSI signature. The location of the receiving tag is approximated as the weighted average of the k selected locations. In this study, the value of k is set to 3. The distance metric between two RSSI vectors employed in the KNN estimator is the *normalized Euclidean distance*. Restated, the Euclidean distance between an RSSI vector and an RSSI signature is further divided by the number of beacon messages received in the vector. This is to handle the potential missing values in the RSSI vector since the uncertainty of wireless medium may lead to unexpected packet drops. Normalization is to exclude the bias from the missing values.

3. TESTBED

In this section, we will describe the testbed environment that allows us to build a localization system and conduct the measurement study. We deploy 24 beacon nodes on the 6th floor of a department building at National Taiwan University, with the beacons mounted on top of the ceiling. The beacon nodes are Telos-like modules [21] equipped with TI MSP430 microcontrollers and CC2420 802.15.4 radio. The software is implemented on TinyOS, and the default media access control (MAC), a CSMA/CA-like mechanism, is turned 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 student laboratories. The 24 Telos-like beacon nodes are small boxes distributed evenly along the corridor. To simplify testbed debugging, every beacon node is connected via USB to one of the 2 testbed PCs. The PCs are installed in room 621 and room 613, and each is connected to 12 beacon nodes. The PCs serve as gateways for code upgrades and data logging via USB. We share the same wired setting with MoteLab [26], but our testbed uses lower-cost off-the-




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.

shelf components for long distance USB connectivity. The testbed is designed to use AC power and to enable long-term measurements.

3.1 USB Connectivity

The effective transmission distance of the standard USB interface is about 5 meters. However, our testbed spans a 12mx50m space. Therefore, we use USB extenders [16] to connect nodes more than 5 meters apart. The USB extenders are off-the-shelf products that extend the effective transmission distance up to 45 meters. Each USB extender consists of a local unit and a remote unit. The local unit modulates the signals from the USB format to the format over CAT5e network cable; while the remote unit demodulates the signal back to the USB format.

To avoid deploying numerous lengthy wires throughout the building, the chaining scheme in Figure 2(a) is used to connect beacon nodes. A chain of 4-port USB hubs are connected by USB extenders and CAT5e cables. The beginning of the chain connects to the PC gateway. Each USB hub is further connected with three beacon nodes. This chain is continued until it reaches the maximum transmission range of USB extenders. That is, the distance between the last USB hub and the PC cannot exceed 45 meters. This limitation is also the reason that 2 chains and PCs are required to cover the entire deployment area of our testbed.

3.2 Power Supply

Powering the beacon nodes and the USB devices is a challenging problem. We measured the peak current consumption of the beacon node, the USB hub and the USB extender separately, and the current draw for each component is approximately 60mA, 5mA and 20mA, respectively. Thus, with 12 beacon nodes and 4 sets of USB hubs and extenders, a single chain will consume more than 800mA current. The high current requirement makes it infeasible to source power directly from a PC's USB port.

Instead, in the testbed, the USB hubs are powered externally. As Figure 2(a) shows, the USB hubs are connected to an external AC power adapter that is capable to provide a maximum of 3A current. However, due to long length and the large current consumption of the beacon chain, a voltage drop as much as 1V can be observed at the end of the chain. Hence, although the standard supply voltage for

a USB interface is 5V, a 6V power adapter is used for the external power source in order to compensate for the additional voltage drop. We tested all the devices in the testbed beforehand to ensure they can sustain the 6V power supply. Note that the power connection pin (VCC) between the USB hub and the USB extender remote unit is cut as shown in Figure 2(b). This is to ensure that the USB hub draws power solely from the external power source and not from the PC. The excessive amount of current draw from the PC's USB interface can cause the interface to function incorrectly.

4. MEASUREMENT METHODOLOGY

The RSSI-based localization system is susceptible to environmental noises. In a typical office or campus environment, background noise could be sourced from WiFi, Bluetooth, a 2.4 GHz cordless telephone, a microwave-oven or other RF devices operating on the 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, we generate WiFi traffic at different levels, and collect the following data: (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

A survey of the 6th floor of the department building reveals more than 10 APs operating in the area. Among them, six APs, which are located on the ceiling of the corridor, are installed by the university to provide general wireless Internet access for staff and students. The Others are deployed by individual laboratories and have restricted access. In order to understand the effect of background WiFi traffic on the localization system, we select one of the generally accessible APs on the ceiling and establish WiFi connection with the AP. In the meantime, the localization system is set to operate on the channel that is closest to the center frequency of the AP's operating channel. The receiving tag is placed on the corridor right below the AP to capture the effect of WiFi interference to the IEEE 802.15.4 interface at the worst case.

4.2 Data Logging

In our experiment, three additional laptops are used. One serves as WiFi traffic generator, the second one measures the amount of WiFi traffic in the channel, and the third one is a data logger that records all the packets received by the receiving tag. To generate WiFi traffic at different levels, the traffic generating laptop is connected to the Internet via the selected AP and a large file is downloaded from an FTP server using FlashFXP [9], an FTP client that allows the user to set the upload/download speed limit. Each rate-limited file transfer session is slightly longer than 10 minutes.

The second sniffer laptop is placed near the receiving tag to sniff all WiFi traffic in the channel. In this way, we can measure the actual WiFi traffic rate on the air and also ensure there is no unexpected extra traffic occurred during the measurement. Dumpcap [1], a Linux packet header capturing tool built on the pcap library, is used to log all the packets the laptop hears. We conduct the experiments in midnight during weekends; as a result, only a very small amount of traffic other than the generated one is observed

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 Traffic (pkt/MB)	48789/4.96	74257/19.34	211008/95.80	247684/124.88	450176/207.62
Beacon Messages (pkt/ fMB)	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

by the sniffer. This ensures that the WiFi interference to the localization system is mostly generated by us and coming from the same AP. In the 5 sets of experiments, the averages of observed background WiFi traffic rates are 68, 264, 1308, 1705, and 2835 kbps. No kernel loss is reported by Dumpcap during the experiments. However, the traffic rate reported here should be lower than the actual amount of the traffic in the channel, since the sniffer cannot capture corrupted packets.

The receiving tag is connected to the third data logger laptop through the USB interface. During each file transfer session, the receiving tag sends all the received beacon messages to the logger laptop. Similarly, we have all the beacon nodes in the testbed also pass the beacon messages received from each other to the gateway PCs through the USB interface. Table 1 details the basic statistics of the data sets.

The localization error can be calculated using the beacon messages collected on the data logger laptop. An RSSI vector can be assembled by picking up the RSSI values from the beacon packets received within 200ms. Given an RSSI vector, the location can be estimated as described in Section 2. 3000 location estimations are generated during an FTP download session. In addition, we also collected the beacon messages overheard by neighboring beacons. The data is used to compute the packet reception rate (PRR) of beacon-beacon links. PRR is a natural metric to measure link quality, and potentially may vary with different amount of background noise. To validate this point, we examined the PRR of beacon-beacon links near the AP where the generated background traffic emerges. PRR is calculated as the percentage of beacon packets received within a sliding window of 50 beacon cycles.

In summary, the WiFi and RSSI data sets are used to study the effect of WiFi background traffic to the localization error. The RSSI and PRR data sets are used to observe the correlation between the beacon message reception rate at the neighboring beacons and the localization error to facilitate discussions towards robust indoor localization.

4.3 Loss in Beacon Message Logging

Software running on the beacon nodes as well as the PCs to collect data through a serial port is prone to errors. During the course of the experiments, several errors in the data sets have been identified. Most of them are software bugs and are quickly corrected. The remaining errors are caused by hardware and communication problems. To ensure that traces are not contaminated by software bugs and to accurately assess the quality of the traces, 2 error checks are implemented to identify hardware and communication problems.

Message length check. Every beacon message had a fixed length. For unknown reasons, the CC2420 radio stack sometimes received a valid packet with an altered packet length under high contention. This problem was also noted

in the TinyOS CC2420 radio module documentation, but no fix is currently available. The incorrect length field could be longer or shorter than the correct one. A shorter length value would result in incomplete packet payload and should always be discarded. In contrast, 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, in this case, the RSSI reading from the radio is usually 0xFF, which means invalid. This problem, although not critical to other uses of sensor data packets, is problematic for RSSI-based localization systems. Simply recording the RSSI value without checking would lead to erroneous location estimations. To correct this problem, the number of bytes in the packet are verified and packets with incorrect packet length are discarded.

Serial error check. To capture bit errors during serial transmissions, a 16-bit CRC checksum value is appended to each packet before sent to a PC or laptop via a serial port. Packets failing the checksum check are discarded by the serial listener. In addition to checksum, a serial sequence number is also appended to each packet to check for unexpected serial losses. The serial listener determines the amount of packet losses from the sequence number. Throughout the experiments, no checksum failure or packet loss is reported. This also shows that wiring the sensor nodes to a central PC is effective for the measurement study.

4.4 Beacon Message Synchronization

To find out the relationship between the localization errors and the PRR of beacon-beacon links, synchronization among the tag and beacons is required. Instead of adding an additional online clock synchronization protocol on the beacons, we performed an offline trace synchronization similar to Jigsaw [6] to reduce extra network traffic. The trace synchronization method utilizes the beacon messages in the trace itself to determine clock offsets between neighboring beacons. In that, when a beacon message is received, it is timestamped by the receiver’s local clock. Assuming two beacons receive the same beacon message at exactly the same time, the clock offset can be determined by comparing the two timestamps. Once the all-pair clock offsets are determined, the traces can be synchronized to a single reference clock. Note that the beacons on the testbed form a multi-hop network, and the clock offset is not directly available for every pair of beacons. However, the all-pair clock offsets can be calculated transitively.

Consider the example in Figure 3, with three logs from beacon 1, 2, and 3. Each trace contains a list of beacon messages the node overhears during the experiment. The clock offset between beacon 1 and 2 can be calculated since they both received the same message coming from beacon 4. The offset is $72.2 - 30.4 = 41.8$. Similarly, the offset between beacon 2 and 3 is $13.1 - 74.4 = -61.3$. The clock offset between beacon 1 and 3 can then be calculated transitively from the previous two values as $41.8 + (-61.3) = 19.5$. In

Beacon 1			Beacon 2			Beacon 3		
ID	Seq	Time	ID	Seq	Time	ID	Seq	Time
2	20	30.1	4	31	72.2	8	20	12.8
4	31	30.4	1	15	72.4	5	26	13.1
7	25	31.2	5	26	74.4	6	11	13.5

Figure 3: Example of Offline Trace Synchronization. The clock offset between beacons can be found out by matching the same messages received, assuming the messages are received at the same time.




Figure 4: Distribution of Localization Errors Under Different WiFi Data Rate. The localization error is larger when the WiFi data rate is higher.

practice, we do not need to calculate the clock offsets for every pair. Instead, we only calculate sufficient pairs of clock offsets that allow us to synchronize the whole network. The receiving tag is set as the global clock, and its clock offsets with all other beacons are calculated and used to synchronize the traces.

5. TRACE ANALYSIS

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

5.1 Localization Errors

Figure 4 depicts the cumulative distribution function (CDF) of the localization errors. The localization accuracy is pretty good with 50th percentile error, i.e., 53cm. We believe such good accuracy comes from the following reasons. First of all, DESYNC is applied on beacons. Collisions are thus reduced, and the receiving tag can receive sufficient RSSI readings to derive accurate location estimations. Secondly, the survey and the test conditions, e.g. antenna orientation and the way the receiving tag is wore, are held the same throughout the experiments. Lastly, the beacon density of the system is high. In our testbed, we place a beacon for every 5 meters.

The test results show that the localization errors are influenced by the background WiFi traffic. In the 50th percentile, the errors increases from 53cm to 81cm (53% increase) as the background WiFi traffic increases. The increase in the 80th percentile error from 160cm to 385cm (141% increase) is particularly large. This indicates 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 Impact of General Beacon Message Losses

To understand how the background traffic impacts localization accuracy, we first look into the loss of beacon messages. In Figure 5(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 received beacon message goes down under high background traffic rate; and there is a strong correlation between background traffic rate and 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 Figure 5(b). The corresponding average localization error and the variance shown in the figure indicate that fewer received beacon messages would result in increasing localization error and variation. A location estimation would be very imprecise, if it is only based on one or two beacon RSSI reading; and average errors could 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 5(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 higher background traffic would increase probability for the receiving tag to receive fewer beacon messages. This is because the background traffic interferes with, and sometimes corrupts, the delivery of beacon messages. Therefore, localization error and its variance both get worse when background traffic rate becomes higher.

5.3 Impact of Individual Beacon Message Loss


Next, we investigate the effect of losing beacon messages from difference beacons. That is, we would like to understand whether all beacons have the same influence on the localization accuracy.

We measure the importance of each beacon by defining the *beacon impact*. We first group the measurement results by the number of beacon message loss. Then, for each beacon, we further divide each group into two sets. In the first set, the *observed* set, the message from the beacon is observed; while in the other set, the *missing* set, the message from the beacon is missing. The *beacon impact* of a particular beacon is defined as the difference of the average localization error between its missing set and observed set. Intuitively, if a beacon is important, missing its RSSI value would result in a larger localization error. Then we call this beacon desired, and the beacon would have a positive beacon impact. On the other hand, if the beacon impact of a beacon is negative, including the beacon’s RSSI value may instead increase localization error and the beacon is called undesired.


In our data, we found that the beacons that are closer to the receiving tag are more likely to be desired; while the beacons that are farther away are undesired. Figure 6 shows the beacon impact values of six different beacons. Beacon 23, 24, and 25 that have higher RSSI values are beacons closer



(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 5: 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.

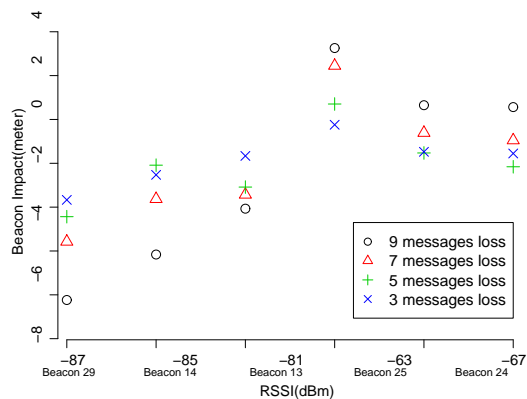


Figure 6: Impact of Individual Beacons. Beacons closer to the receiving tag (with larger RSSI) are more important to the localization system. The beacon impact values for closer beacons are positive indicates that observing their RSSI values shows smaller localization error than those missing the RSSI.

to the tag, while beacon 13, 14, and 29 are farther away from the tag and thus have lower RSSI values. In the figure, we show the beacon impact values for each beacon under different numbers of beacon messages loss, respectively, 3, 5, 7, and 9 message losses. From the figure, we can see that that it is more desirable to receive beacon messages from closer beacons, since they have 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; while the error would increase by 6 meters if an undesired beacon message is received. In addition, a larger amount of beacon message losses shows more prominent positive or negative impact on localization accuracy, since when the information is insufficient, the quality of information becomes more important. The results suggest that weaker RSSI values are more ambiguous for determining locations while stronger RSSI values are helpful for more accurate location estimation. We conjecture that this is due to the RSSI instability of the beacons that are farther away. However, this is yet to be confirmed in future studies.

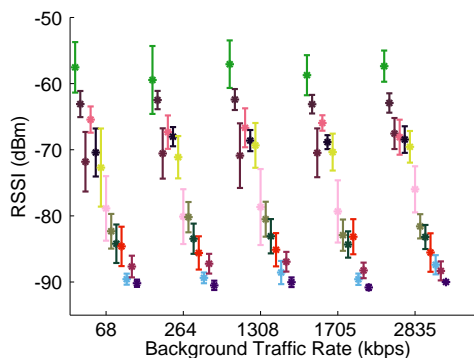


Figure 7: Effect of Background Traffic to RSSI Readings. Varying the background traffic rates does not have significant effects on beacon RSSI readings. The variation of RSSI is thus not the main cause of larger localization error in higher background traffic.

5.4 Impact of RSSI Value Instability

From the Figure 5(b), we can find that if the receiving tag receives sufficient amount of beacon messages, the localization errors remain small even under severe WiFi interference. This suggests that the background WiFi traffic does not distort the RSSI readings, even though it would cause higher packet loss rate. To verify this, in Figure 7, we show the average RSSI readings and their standard deviation for traces under different background WiFi traffic rates. The RSSI values from different beacons are slid slightly for clarity and ease of comparison. From the figure, we can see that the increase of background traffic rate doesn't have a clear impact on the deviation of RSSI readings. Therefore, background WiFi traffic seems have no significant effect on RSSI variation, which is usually a factor to cause localization error. Thus, as shown in Figure 4, the main reason behind the increase of the localization error is not the variance of RSSI readings, but the loss of beacon messages, which is caused by background WiFi traffic.




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.

6. TOWARDS ROBUST INDOOR LOCALIZATION

Our experimental results indicate that the localization system is susceptible to WiFi interference, and this is mainly due to the fact that the WiFi traffic might collide with beacon packets in the air. Note that our localization system already employs the normalized KNN algorithm to mitigate the effect of incomplete RSSI vector. The subsequent questions arise from the measurement study are: (1) How much WiFi traffic is there in everyday environment; and whether there is a need to design interference-resilient mechanisms. (2) What mechanisms might be used to detect interference and to work around the interference problem.

WiFi Traffic Presence Figure 8 shows the observed WiFi data rate during a busy working hour at the 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 previous works [22, 13]. 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 estimations. This again suggests that for a robust localization system, handling the coexistence of wireless signals is crucial. Thus, for a RSSI-based localization system to perform stably with accuracy, it might be necessary for it to be able to dynamically operate at different frequency channels. Dynamic control of frequency spectrum can be important for a robust indoor localization system that provides consistent and accurate location estimations.

Interference Detection Being aware of the interference level is important towards building robust localization systems. To detect interference, the system can monitor the link quality between beacon nodes. Beacon nodes are supposed to be static and the links among them should be relatively stable when there is no interference. Observing the abrupt change on link quality indicates the occurrence of interference. Figure 9 shows the relationship between the beacon packet reception rate of a specific link between two beacons 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. Each beacon node could track the packet reception rates from neighboring beacons. Then,

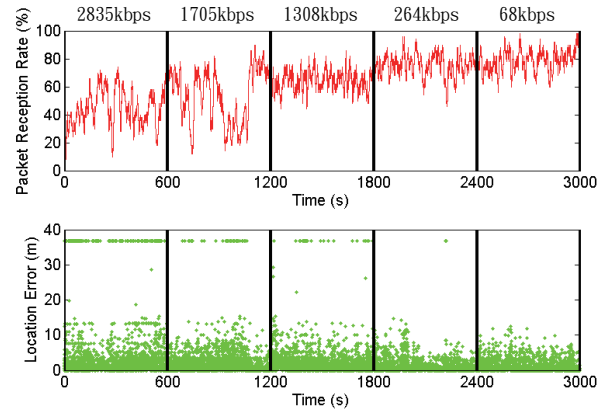


Figure 9: Packet Reception Rate (PRR) vs Localization Errors. This figure shows the PRR on a specific link between two beacons. It can be seen that there is a strong correlation between link PRR and localization error. Link PRR can potentially be an indicator on the background interference level.

from the packet reception rate observed at the neighboring nodes, the beacon node could be aware of the existence of interference.

Beacon Prioritization Furthermore, our analysis in Section 5.3 shows that 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 estimating locations, the system should give higher priority to the information provided by the closer beacons.

7. RELATED WORK

Co-channel interference is not a new problem, and the problem exists no matter whether radio standards are similar or not. Srinivasan *et al.* [25] and Zhou *et al.* [28] have observed the impact of interference between 802.11 and Zigbee signals; while Gummadi *et al.* [10] analyzed how IEEE 802.11 communication is interfered by the use of Zigbee-based devices and wireless cordless phones.

The state-of-art indoor localization techniques can be divided into two categories: range-free and range-based. The former does not localize targets based on range estimation [11]; while the latter uses absolute angle estimates [19] or point-to-point distance estimates by measuring propagation time [20, 23], signal strength degradation, or signal interference patterns [17]. In particular, the range-based systems that exploit the measured RSSI values for location estimation are either based on the decay model of distance to signal strength [20, 18, 15] or the RSSI fingerprints, also known as the RSSI signatures. The RSSI-signature-based system relates an observed set of signal strengths to ones at known locations [5]. The RSSI-signature-based techniques are further categorized into deterministic and probabilistic techniques. For deterministic method, RSSI readings are first collected at known locations; and the estimated location is the place whose RSSI signature is closest to the measured RSSI set. Probabilistic methods, on the other hand, construct a probability distribution over the target's location. The extended

Kalman filter [14] or particle filter [12, 24] based approach, though can reduce estimation error, requires predictive mobility model which is not obtainable without the aid of other sensors.

8. CONCLUSION

In this study, we build an indoor sensor network testbed and, on top of it, implement an RSSI-signature-based localization system. The unique testbed design facilitates simultaneous collection of different system components in the distributed sensor network. The measurement results allow us to examine the performance of the localization system under RF interference and provide as quantitative evidences that RF interference poses significant adverse effect on localization accuracy.

Acknowledgement

The authors would like to thank the anonymous reviewers for their constructive comments. This work was supported in part by grants from Intel Education Program, the National Science Council of Taiwan under Contract NSC 98-2221-E-002-072-MY3 and NSC 98-2220-E-002-024 and Mr. and Mrs. Chun Chiu Stanford Graduate Fellowship.

9. REFERENCES

- [1] Dumpcap man page. <http://www.wireshark.org/docs/man-pages/dumpcap.html>.
- [2] Real time locating systems 2006-2016 (RTLs). *GIExpress Market Research Report*, Jul 01, 2007. <http://www.giexpress.com/products/ix37643/>.
- [3] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *Proceedings of ACM SIGCOMM '04*, pages 121–132, 2004.
- [4] B. Bachelidor. RTLs market to grow 30 percent annually. *RFID Journal*, May 10, 2006. <http://www.rfidjournal.com/article/articleview/2325/1/1/>.
- [5] P. Bahl and V. Padmanabhan. An in building RF-based user location and tracking system. In *Proceedings of IEEE INFOCOM 2001*, April 2001.
- [6] 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 ACM SIGCOMM '06*, pages 39–50, 2006.
- [7] J. Degesys, I. Rose, A. Patel, and R. Nagpal. Desync: self-organizing desynchronization and TDMA on wireless sensor networks. In *Proceedings of ACM IPSN '07*, pages 11–20, 2007.
- [8] E. Elnahrawy, X. Li, and R. Martin. The limits of localization using signal strength: a comparative study. In *IEEE SECON 2004*, pages 406–414, 2004.
- [9] FlashFXP. <http://www.flashfxp.com/>.
- [10] 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, 2007.
- [11] 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 ACM MobiCom '03*, pages 81–95, 2003.
- [12] J. Hightower and G. Borriello. Particle filters for location estimation in ubiquitous computing: A case study. In *Proceedings of Ubicomp'04*, 2004.
- [13] A. Jardosh, K. Ramachandran, K. Almeroth, and E. Belding-Royer. Understanding link-layer behavior in highly congested ieee 802.11b wireless networks. In *Proceedings of ACM E-WIND '05*, pages 11–16, 2005.
- [14] T. Liu, P. Bahl, and I. Chlamtac. A hierarchical position prediction algorithm for efficient management of resources in cellular networks. In *Proceedings of the IEEE GLOBECOM '97*, pages 982–986 vol.2, Phoenix, AZ, USA, November 1997.
- [15] K. Lorincz and M. Welsh. Motetrack: a robust, decentralized approach to RF-based location tracking. *Personal Ubiquitous Comput.*, 11(6):489–503, 2007.
- [16] C. F. I. LTD. USB extender. <http://www.cfi.com.tw>.
- [17] M. Maróti, P. Völgyesi, S. Dóra, B. Kusý, A. Nádas, Ákos Lédeczi, G. Balogh, and K. Molnár. Radio interferometric geolocation. In *Proceedings of ACM SenSys '05*, pages 1–12, 2005.
- [18] D. Niculescu. Positioning in ad hoc sensor networks. *IEEE Networks*, 18(4):24–29, July-August 2004.
- [19] D. Niculescu and B. Nath. Ad hoc positioning system (aps) using aoa. In *Proceedings of IEEE INFOCOM'03*, volume 3, pages 1734–1743, 2003.
- [20] N. Patwari, I. Hero, A.O., M. Perkins, N. Correal, and R. O'Dea. Relative location estimation in wireless sensor networks. *IEEE Transactions on Signal Processing*, 51(8):2137–2148, Aug. 2003.
- [21] J. Polastre, R. Szewczyk, and D. Culler. Telos: enabling ultra-low power wireless research. In *Proceedings of IEEE IPSN '05*, page 48, 2005.
- [22] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan. Measurement-based characterization of 802.11 in a hotspot setting. In *Proceedings of ACM E-WIND '05*, pages 5–10, 2005.
- [23] A. Savvides, C.-C. Han, and M. B. Strivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of ACM MobiCom '01*, pages 166–179, 2001.
- [24] 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 IEEE PERCOM '05*, pages 75–84, 2005.
- [25] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. Some implications of low-power wireless to ip routing. In *Proceedings of HotNets V*, Irvine, CA, Nov 2006.
- [26] G. Werner-Allen, P. Swieskowski, and M. Welsh. Motelab: a wireless sensor network testbed. In *Proceedings of IEEE IPSN '05*, page 68, 2005.
- [27] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *ACM SenSys '03*, pages 1–13, 2003.
- [28] G. Zhou, J. Stankovic, and S. Son. Crowded spectrum in wireless sensor networks. In *Proceedings of EmNets '06*, Cambridge, MA, May 2006.