

Experimental Evaluation of Precision of a Proximity-based Indoor Positioning System

Sylvia T. Kouyoumdjieva and Gunnar Karlsson
School of Electrical Engineering and Computer Science
KTH Royal Institute of Technology, Stockholm, Sweden
Email: {stkou, gk}@kth.se

Abstract—Bluetooth Low Energy beacons are small transmitters with long battery life that are considered for providing proximity-based services. In this work we evaluate experimentally the performance of a proximity-based indoor positioning system built with off-the-shelf beacons in a realistic environment. We demonstrate that the performance of the system depends on a number of factors, such as the distance between the beacon and the mobile device, the positioning of the beacon as well as the presence and positioning of obstacles such as human bodies. We further propose an online algorithm based on moving average forecasting and evaluate the algorithm in the presence of human mobility. We conclude that algorithms for proximity-based indoor positioning must be evaluated in realistic scenarios, for instance considering people and traffic on the used radio bands. The uncertainty in positioning is high in our experiments and hence the success of commercial context-aware solutions based on BLE beacons is highly dependent on the accuracy required by each application.

Index Terms—proximity-based indoor positioning, Bluetooth Low Energy, Estimote beacons, Aruba beacons

I. INTRODUCTION

For more than a decade, academia and industry have devoted effort into the indoor localization problem; however to this date there has not been any widely accepted solution. Most existing works aim at finding a compelling alternative to the Global Positioning System for indoor usage. These new systems should be able to tell the position of a person equipped with a mobile device at any moment with as high precision as possible [1], and they are expected to be one of the major components for introducing context-awareness into our daily lives. By far the most popular and vastly researched technique for indoor positioning is based on *fingerprinting*, i.e., recorded patterns of radio signal strength measurements for a given space used to estimate a user’s location based on her device’s current received signal strength values. The technique is often based on Wi-Fi [2] but recent advances also consider Bluetooth fingerprinting (see Section VI). Fingerprinting is however expensive, mostly in terms of deployment efforts and time, as it requires engineers to map out each space by performing multiple measurements and to derive the required radio signal strength profile. It is also highly prone to disturbances not present during the recording of the fingerprints (e.g., people).

Many future smart city scenarios (for instance, locating discounted goods in a shopping mall or receiving relevant information for an exhibit in a museum) may not require a high level of precision for indoor localization. Instead,

the main concern in these scenarios is not the exact user location but *whether the user is close (enough) to a particular point*. To address these use-cases, *proximity-based indoor positioning* has been introduced. It relies on small battery-powered hardware transmitters, beacons, that are installed into a space and periodically announce their presence to nearby mobile devices, i.e., smartphones. The mobile devices are in turn able to assess the signal strength and determine how close or how far they are located from a particular beacon. A number of hardware manufacturers have already integrated the two main communication protocols for beacons, Apple’s iBeacon and Google’s Eddystone, into their products with the promise to enable real context-awareness. In actuality, however, the number of commercial implementations is limited.

In this work we explore the readiness for commercial proximity-based indoor positioning systems to be deployed on the market, and propose an algorithm for proximal indoor positioning in the presence of human mobility. To this end, we perform extensive experimental evaluations using two off-the-shelf and state-of-the-art Bluetooth Low Energy beacons. Our main contributions are as follows:

- We demonstrate that the precision of proximity-based systems for indoor localization is affected by the distance between the mobile device and the beacon as well as by the orientation of the device and the positioning of the beacon (Section III).
- In the presence of a single beacon, we show that radio wave absorption by the human body greatly affects the precision of proximity-based localization. We further show that the closer a body obstruction is to the beacon, the more corrupted are the measurements (Section III).
- In the presence of multiple beacons, we demonstrate that it is not possible to correctly evaluate which beacon is located closest to a user equipped with a mobile device, even in the case of a single body obstruction (Section IV).
- We devise an online algorithm based on moving average forecasting in order to determine the closest beacon to a mobile user without introducing much delay in the data collection and decision-making process (Section V).

The contributions together provide a reality check on how well commercially available solutions for proximity-based indoor positioning may work. Already in the limited albeit realistic settings used herein, we find that the uncertainty due




Fig. 1. (a) An Estimote proximity beacon. (b) An Aruba proximity beacon. (c) The interior of the conference room used as a setting for the experiment.

to environmental factors is high and susceptible to dynamics. These factors, such as presence of people, orientation of the device, mounting of the beacons in the infrastructure and traffic on the used radio bands, cannot be estimated by the measuring device and hence will lead to incorrect estimates of positions.

II. EXPERIMENT DESIGN

A. Bluetooth Beacons

Bluetooth Low Energy (BLE), also referred to as Bluetooth Smart, is a wireless personal area network technology aimed at novel proximity-based applications.

A Bluetooth *beacon* is a small battery-powered hardware BLE transmitter that operates in an advertisement mode and periodically broadcasts data to nearby devices, such as smartphones. Beacon advertisements are broadcasted on three separate channels, and scanning devices listen to those channels in a specific sequence. Each Bluetooth beacon is uniquely identified by a Universally Unique Identifier (UUID). Whenever a nearby device receives an advertisement packet, it may choose to undertake a specific action based on the announced UUID and other data associated with it, i.e. open a web page or invoke a mobile application. Bluetooth beacons are also capable of transmitting their Received Signal Strength Indicator (RSSI) value in addition to other data. This RSSI value is calibrated by the manufacturer of the beacon to be the signal strength of the beacon at a known distance, typically one meter. Using the known output signal strength of the beacon and the signal strength observed by the receiving device an approximation can be made about the distance between the beacon and the device.

For the purpose of this work, we use devices by two of the leading providers of Bluetooth beacons on the market: Estimote (<https://estimote.com>) and Aruba Networks, a Hewlett-Packard Enterprise company (<https://www.arubanetworks.com>). In this work, we use Estimote proximity beacons to evaluate static environments,

and Aruba beacons to evaluate the performance of our proposed online algorithm in mobile environments. This choice is primarily dictated by the greater flexibility of Estimote beacons and the levels of customization they offer in comparison to Aruba beacons which allow us to evaluate larger set of parameters for systems for proximity-based indoor positioning. We however note that the algorithm we propose is not dependent on the type of beacons used.

Estimote beacons have a maximum communication range of 70 m which is sufficient for most indoor environments. Each proximity beacon supports both Apple's iBeacon format and Google's Eddystone format, however it can only operate with one of these formats at a time. For this study we chose to work with the iBeacon format and note that the chosen format has no bearing on the performance of the indoor localization system, as we are only interested in the transmitted and received RSSI values (which are format-independent). Each beacon is also equipped with a motion sensor and a temperature sensor, which were not used during the experiments. An illustration of an Estimote beacon is shown in Fig. 1(a).

Aruba beacons have a maximum communication range of 60 m with each beacon supporting only Apple's iBeacon format. A beacon can be configured to work either in a location-based mode providing direction, or in a proximity-based mode providing proximity-aware notifications. In this work we only configure beacons in proximity-based mode. Aruba beacons allow for significantly less customization than Estimote beacons, with the only tunable parameters being the device name and the transmission power level. Such limitations do not affect our evaluation, however should be taken into account if one would like to deploy beacons at scale. An illustration of an Aruba beacon is shown in Fig. 1(b).

B. Experiment Setup: Static Environment

The experiment is conducted in a conference room, shown in Fig. 1(c). We mount one (or more) beacons on a wall, and we use the length of the conference room to collect measurements at different distances from each beacon. We do not remove objects from the conference room, as we want to evaluate the performance of beacons in a realistic setting. Also, the radio environment is not controlled and there is simultaneous traffic on the 2.4 GHz band.

Fig. 2 illustrates the experiment setup in the presence of a single beacon, positioned at 1 m height. We evaluate the following three cases:

- **No body obstruction:** There is only one person in the conference room, who is holding a mobile device horizontally and is facing a beacon mounted on the wall, see Fig. 2(a). For brevity, we refer to this person as an *observer* and the distance between the observer and the beacon is denoted with $d(O)$.
- **One body obstruction:** There are two people in the conference room. The observer is at a distance $d(O)$ facing the beacon. The second person, which we refer to as a *body obstruction* is at a distance $d(BO)$ from




Fig. 2. Illustration of the experiment setup in the presence of a single beacon, placed 1 m above ground: (a) without body obstruction, (b) in the presence of a single body obstruction facing the beacon and (c) in the presence of two body obstructions between the beacon and the mobile device.

the beacon, facing the beacon, see Fig. 2(b), where $d(O) > d(BO)$.

- **Two body obstructions:** There are two people in the conference room. The observer is at a distance $d(O)$ and is facing in the opposite direction of the beacon thus creating a second obstruction with his own body. As before, the body obstruction (the second person) is at a distance $d(BO)$ from the beacon, see Fig. 2(c), where $d(O) > d(BO)$.

In all of the above cases, we evaluate the precision of proximity-based localization when the beacon is mounted at different heights in the conference room, as well as when the observer and the body obstruction are located at different distances from the beacon (Section III). We further conduct experiments with two beacons, mounted on opposite walls, in the presence of no obstructions and with a single body obstruction at different distances from the beacons (Section IV).

C. Experiment Setup: Mobile Environment

The experiment is conducted in a room, similar to the one shown in Fig. 1(c). We mount Aruba beacons at approximately 1 m height on opposite walls, and we use the length of the room to collect measurements as one user equipped with a mobile device is walking back-and-forth between the beacons at a leisurely pace. No body obstructions are introduced along the path of the mobile user.

D. Data Collection

To collect data, we develop an Android application which is able to constantly listen to the signal advertised by one or more beacons located in the vicinity of a mobile device, and to log every RSSI value received by the Bluetooth interface. The application runs on Android 6.0.1, and for the experiments we use a Huawei Honor 8 device except when stated otherwise. In the application design process, we did not use any manufacturer-specific or open source Bluetooth APIs, as we wanted to evaluate the performance of our proximity-based localization system without any software bias. For instance, custom APIs often abstract the listening behaviour and provide the application directly with an average RSSI value, which can be calculated from a sample size anywhere between 10 and 30 samples. As we are only interested in evaluating the precision of the proximity-based localization, and are currently not concerned with preserving power at the mobile end, we use Android's `SCAN_MODE_LOW_LATENCY` mode which allows the mobile device to constantly listen to advertisement packets.

Typically, Estimote beacons are suggested to operate at a 950 ms advertising interval in order to prolong battery lifetime while simultaneously providing fast beacon discovery. However, to speed up the data collection process, we reconfigure each beacon with an advertising interval of 100 ms. (We explore the effect of different advertising intervals in Section III-A.) The transmit power is set to -8dBm which provides a maximum communication range of approximately 30 m, sufficient for our indoor setup described above. Aruba beacons, on the other hand, do not offer flexibility when choosing the advertising interval. It is preset to 500 ms, however some randomness seems to be introduced to that interval in order to avoid collisions on the wireless medium whenever two or more beacons are within communication range to a mobile device. The transmit power may be chosen between three available levels; for this work it is set to 'high'.

During each experiment in the static environment setup we collect at least 100 RSSI samples per beacon. The duration of experiments however vary, as not all beacon advertisements are received by the mobile device at all times. To assure that experiments have a common ground for comparison, we only interrupt an experiment once the threshold of 100 RSSI samples is reached. Each experiment is done statically, i.e. neither the person holding the mobile phone, nor the person acting as an obstruction move during the data collection. Samples for each measurement point are logged into a file on the mobile device, and then post-processed. All figures that present mean values also show the 95% confidence intervals.

However collecting at least 100 RSSI samples per beacon may result in large delays when the advertising interval is non-adjustable or when it is set to a larger value in order to prolong battery lifetime. For an advertising interval of 500 ms (Aruba beacons) collecting 100 RSSI samples results in a delay of at least 50 s before deciding which beacon is closer to a user. Such delays are not acceptable in mobile environments. Thus in Section V we propose an online algorithm which is able to quickly determine the closest beacon to a user.

III. EXPERIMENT RESULTS IN STATIC ENVIRONMENT: ONE BEACON

A. Effect of advertising interval

We first examine the effect of the advertising interval on the measurement accuracy. For this experiment, a beacon and a mobile device are positioned at a distance 1 m on a flat surface (a table). We vary the advertising interval at the beacon. For

each received beacon, the application at the mobile device collects information about the RSSI and estimates in real-time the distance (according to FRIIS equation [3]) and the perceived sampling interval at the receiver end. The results are summarized in Table I.

TABLE I
EFFECT OF ADVERTISING INTERVAL ON ESTIMATED DISTANCE AND ESTIMATED SAMPLING INTERVAL AT RECEIVING NODE. ACTUAL DISTANCE BETWEEN BEACON AND RECEIVER IS SET TO 1 M.

Advertising interval [ms]	Mean sampling interval [ms]	RSSI [dB]	Estimated distance [m]
939	1784.6	-68.3 ± 3.7	1.1 ± 0.5
759	1131.9	-67.9 ± 3.9	1.1 ± 0.5
497	932.5	-68.6 ± 4.0	1.2 ± 0.6
248	269.1	-69.2 ± 3.5	1.2 ± 0.6
100	105.1	-68.8 ± 3.2	1.2 ± 0.5

The estimated distance between the beacon and the receiver stays approximately constant across different advertising intervals. The mean sampling interval, at which advertisement packets are received at the mobile device, increases significantly the higher the advertising interval. This is due to the unreliability of the wireless channel (remember that the radio environment is not controlled). We note that for applications which require responsiveness, i.e., applications that rely on real-time estimates, the advertising interval should always be kept small while still providing sufficient energy conservation. Evaluating this trade-off is out of the scope of the current work. In this experiment we are not concerned with the battery lifetime, and we set the advertising interval to 100 ms to speed up the data collection process.

B. Effect of device rotation

Next, we are interested in studying how the rotation of the mobile device in the horizontal plane affects the accuracy of the estimate. As before, we position a mobile device and a beacon at a distance 1 m apart on a table, but we now change the angle between the beacon and the mobile device from 0° to 180° . The results for mobile devices of two manufacturers, a Samsung Galaxy S6 and a Huawei Honor 8, are presented in Fig. 3. When the top of the phone is directed towards the beacon (rotation of 0°), the RSSI estimate is closest to the announced baseline of the beacons. Surprisingly, for both devices when the bottom of the phone is directed towards the beacon (rotation of 180°) the RSSI measurements significantly underestimates the distance to the beacon, reporting a mean value of 0.3 m. This behavior may be related to the position of the Bluetooth antenna on the device. However, it shows that the unreliability of location estimates with BLE beacons is further enhanced by the user behavior and the device being used. Results obtained in controlled environment studies may thus be biased, and may present an overly optimistic evaluation of the performance of BLE beacons. For the rest of our experiments, we always position the device such that its top points towards the beacon; although we are aware that this behavior introduces a bias, it also allows us to evaluate other




Fig. 3. Effect of device rotation on average RSSI for two different mobile phones, Samsung Galaxy S6 and Huawei Honor 8.

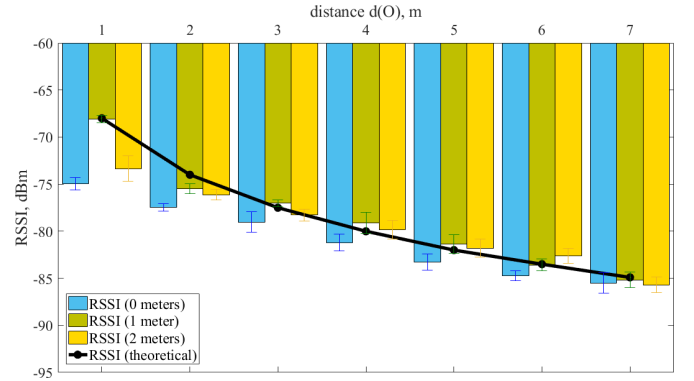


Fig. 4. Effect of beacon positioning on average RSSI. The beacon is consecutively mounted at 0 m, 1 m and 2 m above ground.

effects without being affected by many external factors at the same time.

C. Effect of beacon positioning

We now examine whether the height at which a beacon is mounted affects the precision of localization. The setup is shown in Fig. 2(a) with the observer holding the mobile device at approximately 1 m above the floor. The beacon is consecutively mounted at the floor level and at 1 m and 2 m above the floor. The observer collects RSSI measures at different distance from the beacon, from 1 to 7 m. To compare our measurements to the theoretical RSSI value at each distance, we use the equation defined in [4]:

$$RSSI = -10n \log_{10} d - A \quad (1)$$

where n is the signal propagation constant (usually $n = 2$), d is the distance in meters between the beacon and the mobile device, and A is the reference received signal strength in dBm measured when the receiver and the transmitter are 1 m apart.

The results are presented in Fig. 4. Independent of the height, at which the beacon is mounted, the RSSI value decreases with increased distance between the beacon and the mobile device. However, we notice that placing the beacon close to the ground produces slightly worse results than mounting it on the wall. We also note that placing the beacon at approximately the same height at which the mobile device

is held (in this case, at 1 m) produces results which are closest to the theoretical RSSI values. Although in a realistic scenario it is highly unlikely to always meet this requirement (as people often hold their mobile devices at different heights and angles while moving through an area), we will utilize this result in our further experiments.

D. Effect of body obstruction

Next, we conduct a set of experiments to assess the effect of body obstructions on the precision of proximity-based localization. The setup is shown in Fig. 2(b). We explore two cases: one in which the body obstruction is positioned in close proximity to the beacon, $d(\text{BO}) = 0.5 \text{ m}$, and one in which the body obstruction is located at half distance between the beacon and the observer, $d(\text{BO}) = 0.5d(\text{O})$. The observer is facing the body obstruction and is holding the mobile device at approximately 1 m above the floor. Fig. 5 shows the results. We see that the position of the body obstruction has a significant effect on the received signal strength. When the body obstruction is located in close proximity to the beacon ($d(\text{BO}) = 0.5 \text{ m}$), the received signal strength is lower than if the body obstruction is positioned further away from the beacon. An explanation for this phenomenon is that due to the small distance between the body obstruction and the device, multipath propagation is shielded in a broader sector. Although the relative error is larger when the body obstruction is located in close proximity to the beacon, $d(\text{BO}) = 0.5 \text{ m}$, we see a clear trend in decreasing RSSI values with increase of the distance. Thus, even if accurate absolute positioning is impossible, relative positioning of the observer is still achievable.

In Fig. 6 we further show the actual vs. estimated distance to the observer in the case of a single body obstruction. When the body obstruction is positioned in close proximity to the beacon, $d(\text{BO}) = 0.5 \text{ m}$, the absolute estimation error is larger, especially at shorter distances between the observer and the beacon. However we see a clear trend for increase in the estimated distance with increase in the actual distance. This is not the case when the body obstruction is located half distance between the beacon and the observer, $d(\text{BO}) = 0.5d(\text{O})$. Although the absolute estimation error is smaller by approximately a factor of two, it is harder to determine the actual position of the observer in comparison to previous measurements at other locations due to the lack of a trend in the collected values. (This is also illustrated in Fig. 5 where a clear trend is not present for $d(\text{BO}) = 0.5d(\text{O})$.) In other words, depending on the application scenario at hand, the estimation error may or may not be a defining factor for what is an acceptable performance of a context-aware service. If we are only interested in relative positioning, then being able to determine whether the observer has moved closer or further away from the beacon may be sufficient, independent of the actual estimated distance. However, we should always consider that determining even relative positioning may not be trivial depending on the presence and position of potential body obstructions.




Fig. 5. Effect of positioning of body obstruction on precision of proximity-based localization: a single body obstruction is placed at a distance of $d(\text{BO}) = 0.5 \text{ m}$ and at $d(\text{BO}) = 0.5d(\text{O})$ from the beacon.




Fig. 6. Actual vs. estimated distance to observer in the case of a single body obstruction positioned at $d(\text{BO}) = 0.5 \text{ m}$ and at $d(\text{BO}) = 0.5d(\text{O})$ from the beacon.

We then examine how the number of the body obstructions affects the precision of the indoor positioning system. To create the second body obstruction, the observer is facing in a direction opposite the body obstruction, see Fig. 2(c), while again holding the mobile device at approximately 1 m above the floor. Fig. 7 shows the results. We present measurements from a single body obstruction and two body obstructions and compare them to the theoretical reference values for RSSI, obtained by Eq. (1). Although there is a significant difference in the performance of the system in the presence of a single obstruction, adding a second obstruction to the system when the first one is at $d(\text{BO}) = 0.5 \text{ m}$ does not dramatically decrease the performance. Instead, the RSSI measurements in the case of two body obstructions are relatively close to those, obtained with only one body obstruction in close proximity to the beacon. We do see however that the impact of the number of body obstructions is more visible in the case when the body obstruction is positioned half distance between the observer and the beacon. In this case, we experience a decrease of additional -3 dBm on average across all sampling distances. Finally, we compare our experimental results to the theoretical reference RSSI values. We see that when the obstruction is positioned close to the beacon, the difference between the reference and the measured values is approximately -15 dBm .




Fig. 7. Effect of number of body obstructions on precision of proximity-based localization.

Although this difference decreases as the observer moves further away from the beacon, even at 7 m we still experience around -6 dBm lower attenuation than the theoretical value. Assessing the performance of the system if more than two body obstructions are in the area is out of scope, however, we could expect that the general performance will be degrading as we add more obstructions to the experiment. Thus, a general conclusion we could draw is that evaluating schemes for improving accuracy of BLE-based indoor positioning systems should never be done in an idealized environment, as this may produce misleading results when there are people present in the space.

IV. EXPERIMENT RESULTS IN STATIC ENVIRONMENT: TWO BEACONS

We are further interested in evaluating whether proximity-based systems are able to correctly recognize which is the closest beacon to an observer in the case of multiple beacons in an area. We believe this is an important question, as it is closely related to applications in different proximity marketing scenarios, for instance when locating goods in shopping malls or receiving information about art works in a museums. For this experiment, we install two beacons on the opposite walls of the conference room, Fig. 1(c), at a distance of 7 m apart. Both beacons are mounted 1 m above ground, based on the findings in Section III-C. We conduct two sets of experiments. In the first set, the observer does not block the line-of-sight between the two beacons, and she is positioned perpendicularly to the beacons. In the second set, the observer faces one of the beacons, and her body serves as an obstruction towards the beacon mounted on the opposite wall. In both cases, the observer holds a mobile device at approximately 1 m height above ground, i.e. the mobile device is at the same height as both beacons. The results are presented in Fig. 8.

We would like to note the existence of two x-axes in both graphs in Fig. 8 with each x-axis representing the distance between the observer and the corresponding beacon. (Observe that the scale of the top axis is decreasing.) Fig. 8(a) shows some promising results about the performance of the proximity-based indoor localization system when there are no obstructions between the beacons. As the observer moves, the




Fig. 8. Precision of proximity estimation in (a) the absence of obstructions and (b) the presence of a single obstruction between two beacons positioned 7 m away from each other.

system is able to detect sufficiently well which of the two beacons is closer to her. However, such an uninterrupted scenario is too optimistic and very unlikely to happen in reality. The results in Fig. 8(b), which shows the performance in the presence of a single body obstruction, reveal the shortcomings of the system: it is not possible to detect the correct distance to the obstructed beacon (beacon #2) almost until the last meter. Thus, the observer is led to believe that the unobstructed beacon is located closer to her for almost the whole duration of the experiment. (Observe that the difference in the RSSI values for the unobstructed beacon in Fig. 8(a) and 8(b) is due to the fact that the radio environment is not controlled.) We note that since the performance of a proximity-based indoor localization system is so poor in the presence of only a single obstruction (the observer's own body), adding more obstructions as well as working with a more complex topology will only further deteriorate the performance, ultimately leading to entirely untrustworthy results. This is a clear indication that the success of proximity-based indoor localization is highly dependent on the precision required by potential applications.

V. EXPERIMENT RESULTS IN MOBILE ENVIRONMENT

In the previous sections we looked at static environments in which is it possible to tolerate long delays in the data collection process in order to determine which beacon is closer to a user. However, a more realistic scenario would be one in which the

user is moving between two (or more) beacons at walking speed. For instance, such setup is highly relevant to museums that aim to provide an augmented reality experience to visitors by equipping art installations with one or more beacons. The goal of such system would be to deliver relevant information for the art installation closest to the user. Thus, as the user moves through space, the information delivered to her should change in real-time with respect to her position. We note that we do not aim to *precisely* position the user with respect to the beacons but simply to determine what is her position *relative to* the beacons. In the case of providing augmented reality for museum visitors, the precision is sufficient when the information provided by the beacon is associated with an artwork that the user is able to see from her current location (similarly to Section IV but with mobility).

On one side, the above requirements for a proximity-based indoor positioning system do not allow for long delays in the data collection process (i.e., 50 s for a typical advertising interval of 500 ms for Aruba). On the other side, the RSSI measurements, received by a mobile device, are usually too noisy to provide correct estimation of the closest beacon when taken as-is. Fig. 9(a) shows the fluctuations in RSSI measurements when a user is walking at a constant speed towards a beacon placed in one end of the room and then turning around and moving away from it; the user walks the route two times and is only facing the beacon whenever she walks towards it. The high level of noise in the consecutive measurements makes it hard to determine the direction of movement of the user, i.e. whether she is approaching or going away from the beacon. To address this issue, we propose an online algorithm for smoothing noisy RSSI measurements, based on moving average forecasting; such method is typically used for trend analysis of financial data. We here extend the concept, and propose an algorithm consisting of two phases. The first phase begins whenever the mobile device receives an advertising frame from a new beacon in its vicinity (i.e., a beacon it has not heard of in the previous advertising interval). During this phase, the mobile device collects RSSI data in each advertising interval and computes a *simple moving average* across the obtained samples, as follows: $s_i = \sum_{k=1}^i m_k / i$ where s_i is the value of the smoothed curve for sample i , and m_k is the value of a past measurement ($k = 1..i$). The first phase has a duration of 10 consecutive samples. In the second phase, the device continues to collect RSSI measurements however it now evaluates the value of the smoothed curve using an *exponential moving average*: $s_i = (m_i - s_{i-1})w + s_{i-1}$ where s_i is the value of the smoothed curve for sample i , m_i is the value of the RSSI measurement, s_{i-1} is the previously calculated value of the smoothed curve (i.e., for the previous advertising interval), and w is a weighted constant. For the purpose of the algorithm $w = 2/11 = 0.1818$. We note that we use exponential moving average in the second phase of the algorithm, as it gives more weight to recent measurements as compared to weighted moving average. The smoothed curve obtained by the proposed algorithm is shown in Fig. 9(a). Compared to the raw RSSI measurements, the smoothed curve

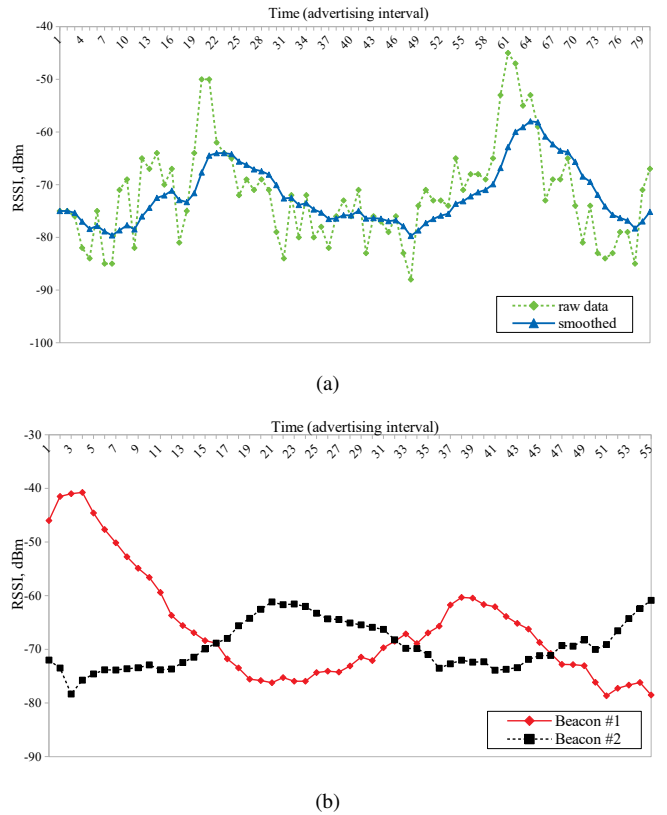


Fig. 9. Performance of an online algorithm based on moving average forecasting for (a) trend line computation under noisy RSSI measurements and (b) proximity estimation in the presence of human mobility.

better depicts the trend in movement of the mobile user.

Finally, we apply the proposed algorithm to a scenario in which the mobile user is moving between two beacons located in both ends of a room. The user starts next to the first beacon and walks at a constant speed towards the second beacon; once she reaches it, she turns around, goes back to the first beacon, turns around once more and again walks to the second beacon. Fig. 9(b) shows the smoothed RSSI curves for both beacons. We note that by applying the proposed online algorithm it is possible to determine the direction of movement of the user without introducing any delay by the data collection process. Moreover, based on the obtained trend lines, it is not only possible to deliver information relevant to the beacon located closest to the user (the beacon which trend line has the biggest negative value at any given time) but also to dynamically switch between information whenever the user is moving towards another beacon. To do this, a simple threshold value can be used. A beacon is then considered closest whenever the value of its smoothed curve is less than the threshold.

VI. RELATED WORK

Although most works on indoor positioning concentrate on working with Wi-Fi, a few recent studies have attempted to realize precise indoor localization only with the help of Bluetooth Low Energy beacons [5], [6], [7]. While most works attempt to improve the performance of indoor positioning

systems, few efforts are done in exploring and characterizing the shortcomings of BLE-based solutions.

A prototype of an indoor localization system that relies on BLE fingerprinting is presented in [5]. The setup requires beacons to be mounted on the ceiling, and uses a combination of RSSI measurement collection and a trilateration algorithm to determine the position of a user. The experiment is conducted in a controlled environment, and there are no obstacles.

In [6] the authors present two BLE-based schemes for indoor localization which rely on collecting RSSI measurements to outline a region in which the user is guaranteed to be found. In [7] the authors propose a two-stage RSSI-based BLE positioning method, including offline training and online locating. Neither of these two works is evaluated experimentally, and they consider only idealized control environments.

While most works attempt to improve the performance of indoor positioning systems, few efforts are done in exploring and characterizing the shortcomings of BLE-based solutions.

One of the first works on BLE beacon accuracy for indoor positioning is presented in [8]. The authors compare the performance of BLE and Wi-Fi fingerprinting, and devise a multipath mitigation scheme (in essence, smoothing batches of measurement data) in order to reduce the measurement error in the Bluetooth signal. The scheme is then validated on a testbed of 19 beacons, installed on a single floor, with a user traversing a predefined path to collect measurements. One of the main shortcomings of this work is the fact that experiments are conducted without other people on the floor, and with no other ongoing communication over the 2.4GHz channel.

A more recent study [9] relies on Estimote hardware to evaluate an indoor positioning system consisting of four beacons. The authors also adopt a fingerprinting technique and attempt to establish in real time the location of a single user within an area based solely on the measured RSSI in the user's device. The proposed algorithm produces promising results, however suffers from the disadvantage of only being evaluated in a controlled environment with minimal obstacles and no additional body obstructions.

As opposed to previous studies, in this work we focus on experimentally evaluating the performance of off-the-shelf and state-of-the-art BLE beacons for proximity-based indoor positioning in the presence of obstructions. The goal is not to provide a new algorithm for BLE-based indoor positioning but instead to demonstrate the need for testing existing and future indoor positioning algorithms in realistic environments.

VII. SUMMARY AND CONCLUSION

Bluetooth Low Energy beacons are inexpensive devices with long battery life that are easy to install. They are therefore highly attractive for providing proximity-based context-aware services. Previous work on indoor beacon-based localization have evaluated the systems in controlled and insulated environments. However, if such systems are to be deployed in public spaces, they will be exposed to different obstructions, both human and physical, as well as exogenous radio disturbances. These are known in principle, but in this work we have

quantified the effects of: beacon height, device orientation and presence of human bodies.

In summary, we have measured the performance of a proximity-based indoor positioning system built with off-the-shelf hardware (Estimote and Aruba beacons) in a realistic office environment. We find that the accuracy of the system depends highly on the orientation of the mobile device and its distance to the beacon, and it depends slightly on the height at which the beacon is mounted. Furthermore, we show that the presence of body obstructions (both the position of the body obstruction and the number of body obstructions in an area) introduce errors that significantly deteriorate the performance of the positioning system. Remark that the environmental settings change dynamically in space and time, and their effects on the system performance cannot easily be estimated by the measuring device. Hence, the performance of a positioning system will depend on the person holding the device (height, orientation and movements), the plan for mounting beacons in the infrastructure, as well as the population of the space. A solution tested without regard for these external factors might give disappointing results when deployed in production.

We also proposed an online algorithm for quick prediction of the closest beacon for mobile users. It illustrates what can be afforded in the design space limited by both a low frequency of beacon advertisements and the acceptable latency of the application for an estimate.

To conclude, our findings indicate that the use cases for position-based services might be circumscribed by the large uncertainty in localizing users with respect to beacons. This is also a message to the research community that continued work on indoor localization should consider realistic settings for the evaluation, not to give overly positive results. We would like to encourage further work on these environmental factors and mitigation of their effects on the performance of indoor positioning. We also suggest to determine the needed accuracy for some generic use cases for such systems, in order to set targets for the continued improvements.

REFERENCES

- [1] D. Lymberopoulos, J. Liu, X. Yang, R. R. Choudhury, S. Sen, and V. Handziski, "Microsoft indoor localization competition: Experiences and lessons learned," *GetMobile: Mobile Comp. and Comm.*, vol. 18, no. 4, pp. 24–31, Jan. 2015.
- [2] S. He and S. H. G. Chan, "Wi-fi fingerprint-based indoor positioning: Recent advances and comparisons," *IEEE Communications Surveys Tutorials*, vol. 18, no. 1, pp. 466–490, 2016.
- [3] C. A. Balanis, *Antenna Theory: Analysis and Design, 4th Edition*.
- [4] Q. Dong and W. Dargie, "Evaluation of the reliability of rssi for indoor localization," in *Proc. ICWCUCA*, Aug 2012, pp. 1–6.
- [5] M. E. Rida, F. Liu, Y. Jadi, A. A. A. Algawhari, and A. Askourih, "Indoor location position based on bluetooth signal strength," in *Proc. ICISCE*, April 2015, pp. 769–773.
- [6] Y. Wang, Q. Ye, J. Cheng, and L. Wang, "Rssi-based bluetooth indoor localization," in *Proc. MSN*, Dec 2015, pp. 165–171.
- [7] Z. Jianyong, L. Haiyong, C. Zili, and L. Zhaohui, "Rssi based bluetooth low energy indoor positioning," in *Proc. IPIN*, Oct 2014, pp. 526–533.
- [8] R. Faragher and R. Harle, "Location fingerprinting with bluetooth low energy beacons," *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 11, pp. 2418–2428, Nov 2015.
- [9] P. C. Deepesh, R. Rath, A. Tiwary, V. N. Rao, and N. Kanakalata, "Experiences with using ibeacons for indoor positioning," in *Proc. ISEC*, 2016, pp. 184–189.