BFound: Sensor Enhanced Localization for Internet of Things

Paramasiven Appavoo
National University of Singapore
Singapore
pappavoo@comp.nus.edu.sg

Hande Hong
National University of Singapore
Singapore
honghand@comp.nus.edu.sg

Mun Choon Chan
National University of Singapore
Singapore
chanmc@comp.nus.edu.sg

Anand Bhojan
National University of Singapore
Singapore
banand@comp.nus.edu.sg

ABSTRACT

Beacons, the backbone of the physical web and location-based services, are widely used to tag objects and places. However, as a beacon's wireless transmission is limited to a ranging technology, the localization information is only available when the beacon is nearby.

In this work, we propose BFound, a navigation and room level localization system that enable users to locate beacons within a building to the room level. The scheme is based on expanding the beacons' capability in two ways. First, the system builds upon crowdsourcing using data from smartphones carried by mobile users and infrastructure beacons to search for facilities within an area and navigate to the target region. Next, we utilize sensors on the beacons to further localize beacons to the room level.

In order to enable room level localization, the beacon's sensors are leveraged to generate a signature unique to the room it belongs. This is achieved using wavelet transform.

Evaluations show that BFound provides sufficient accuracy both for navigation as well as room level localization.

CCS CONCEPTS

• **Human-centered computing → Ubiquitous and mobile computing design and evaluation methods**; • **Networks → Cyber-physical networks**; **Sensor networks**;

KEYWORDS

Beacon, sensor, internet of things, navigation, wavelet, localization

ACM Reference Format:

Paramasiven Appavoo, Hande Hong, Mun Choon Chan, and Anand Bhojan. 2018. BFound: Sensor Enhanced Localization for Internet of Things. In *EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MobiQuitous '18), November 5–7, 2018, New York, NY, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3286978.3287025

MobiQuitous '18, November 5–7, 2018, New York, NY, USA

ACM ISBN 978-1-4503-6093-7/18/11.
https://doi.org/10.1145/3286978.3287025

1 INTRODUCTION

In the mobile computing context, beacons are small, battery powered wireless devices that periodically broadcast small amount of information. Beacons can enhance user experience as smartphone's applications receive contextual information and are able to provide more relevant responses. These beacons are also known as proximity beacons since they mostly use Bluetooth Low Energy (BLE) for advertisements and have limited range. However, BLE 5 can advertise with 255 bytes packet size to multiple times the range of BLE 4. While these beacons are limited in terms of functionalities, they are cheap and easy to deploy and can be easily incorporated into a Internet-of-Things (IoT) framework. There are different types of beacons, including the iBeacon [10], AltBeacon [22] and Eddystone [7]. The differences are in the format and type of data the beacons transmit.

The explosion of beacon deployments to provide contextual services is noticeable at airports [2] and shopping malls [20]. A number of case studies for the use of beacons, as proximity technology, can also be found in different applications ranging from vending machines, business cards, art exhibition projects to city scale deployment [11]. ABI research also forecasts the shipment of 400 million beacons by 2020 [13]. In the same context, there is a proliferation of location-aware services with the use of beacons to mark a location (e.g. hotel's lobby, bedroom, dining room, lecture hall, meeting room, etc.) that are mostly static (location-based beacons). Similarly, there are also beacons that tag items but do not necessarily convey location information (e.g. vending machine, water dispenser, TV monitor, projector, computer equipment or shared office equipment such as printer).

In this work, we look into the problem of localizing an object in an environment with many beacons installed, but only the locations of some of the beacons are known. For example, a set of beacons heard at a particular location could be: beacon #1 saying "I am a TV monitor in the lobby", beacon #2 saying "I am a projector in the seminar room", and beacon #3 saying "I am a coffee machine". For beacon #3, there is no direct way to localize the coffee machine to any of the known places. As illustrated in figure 1, the beacon may be able to tell us that the object being searched for is within a certain region, but there is insufficient information to localize further. Moreover, certain nearby facilities remain invisible for being beyond RF range.

In this paper, we present BFound, a light-weight navigation and localization service. BFound has two components. First, it uses a crowdsourcing approach to build up information to search and navigate to a region close to the target object. Next, it uses sensor data collected by the beacons to localize tagged objects to room level. In order to enable this feature, the sensor data from individual beacon are used to generate a signature unique to the room it belongs using *wavelet transform*.

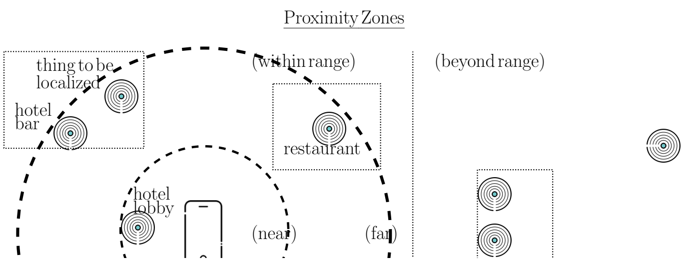

Figure 1: Problems: (1) despite location-based beacons as part of the infrastructure, things cannot be localized, (2) Other objects are nearby but out of RF range.

The main contributions of this paper are as follows:

- A system that builds a virtual map for end-to-end cross-floor navigation and searching using RF beacons from the physical world.
- A method to generate a room signature based on sensor readings that can be used for room level localization.
- We have implemented and evaluated BFound, a system that combines sensor data and RF beacons, for navigation and localization.

We have evaluated the crowdsourced navigation in a building with 100 beacons deployed and the wavelet generated room signature using a 24 hours sensor data collected from 2 buildings as well as 15 days sensor trace from the Intel/Berkeley data set [19]. Evaluation showed that BFound has high accuracy for both RF based navigation and the sensor based room localization.

The rest of this paper is organized as follows: We first present related work in Section 2. The architecture and details of the BFound system are then presented in Section 3 and 4 respectively. Next, the evaluation is discussed in Section 5. Finally, we conclude in Section 6.

2 RELATED WORK

2.1 Indoor Localization

Indoor localization has been an active area of research. Early works like the Active Badge system [26] and the Bat location sensor system [1] performed localization of people wearing special tags using infrared and ultrasound respectively. Next, with the proliferation of WiFi deployments, there has been an explosion of localization approaches using WiFi fingerprints, which started with [4, 23]. These approaches typically require labor-intensive calibration process though more recent approaches, like [17, 24, 25], have proposed mechanisms to reduce the calibration effort. Other features exploited for localization include magnetic distortions inside building [5, 25, 29] and visible light [16, 28]. Our work takes one step ahead of traditional localization techniques by leveraging beacons' sensor data to determine whether they are in same room. Wavelet transform is applied in the process.

2.2 Wavelet Transform

Wavelet transform emerged out of the main limitation of the Fourier transform, which was lacking in localization of detected frequencies. Since then, wavelet transform has been mostly used in the image processing domain. Some of the applications are in image compression (JPEG 2000), in the hiding of digital watermarks[9], in noise removal, face recognition [3] and in querying image databases [14].

2.3 Navigation

In the category of navigation with floor maps, places with high human traffic like airport and other transportation hubs can request for GoogleMaps Indoor [8]. There is also ClickLoc [27] which uses mainly images for localization and navigation. In ClickLoc, a floor map is used to map shot photos from the image space to the physical space. As for navigation without maps, there are (1) Escort [6] that uses opportunistic sensing and interaction to guide navigation, and (2) Travi-Navi [29] that is trace-driven and uses images to aid navigation.

Our work is different from existing work in the following ways. First, BFound uses data about beacons encountered, from mobile crowdsourcing, which are transformed into a graph representation; thus, it is relatively lightweight. Moreover, it uses sensor enhanced localization to enable the navigation to things of interest by potentially identifying the room in which they are.

2.4 Floor Level Detection & Lift/Staircase Detection

Previous works [21] showed the detection of staircases and lifts by using pressure difference and time difference. However, this is under the assumption that the time to change floors using the different methods differ. Moreover, the authors also concluded out that one cannot use the barometer to determine the floor level. However, BFound (1) relaxes the time constraint assumption by leveraging the step counter, and (2) uses relative pressure difference across floors for building its network, where set of beacons' encountered are expected to overlap for users walking on the similar floor levels.

3 BFOUND

The BFound architecture leverages crowdsensing and sensor enhanced localization for end-to-end navigation. As shown in figure 2, there are 4 steps/components, namely:

- Crowdsourced data collection
- Graph building on server
- Navigation
- Sensor enhanced localization

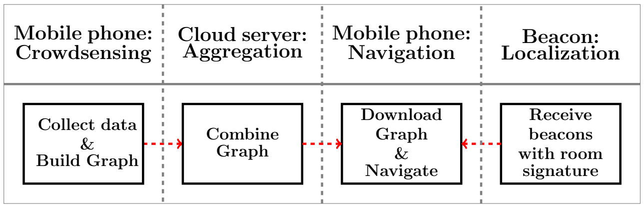

Figure 2: The architecture of BFound.

The basic idea is that beacons encountered by a user moving from one place to another can be used to guide another person moving somewhere along the same path. As such, multiple paths can be combined to create a navigation map that is eventually used to discover new paths and shortcuts relying only on existing beacons with no additional infrastructure.

To this end, crowdsensing is leveraged to build the network of beacons. The beacons encountered during a walk is used to build a local network, in the form of a graph. The challenges are (1) the beacons must be classified to the appropriate floor levels of the building, and, (2) the automatically generated graph must include floor changes, so that end-to-end cross-floor navigation is possible.

The locally generated graphs are uploaded to a Cloud server. The server combines uploaded local graphs, that may cover different parts of a building, into a global graph. The stitching process must ensure that beacons encountered on different floor levels can be combined in a consistent manner. As data from more walks are available, the paths/edges' distances of the global graph converge to the shortest paths.

The generated global graph can be used for navigation to a target beacon in the following manner. From any location, the user can listen for transmission from nearby beacons and figure out his/her approximate location. By inputting the destination, BFound uses the global graph to compute the navigation path in terms of list of beacons to be encountered, the walking directions, the turns and the floor change positions.

While the navigation graph allows the user to reach the vicinity of the destination beacon, in an indoor environment with many rooms, it is still often unclear which room the object is located. It is further challenging when obstructions/partitions in the environment change the beacon signal strength to a level making RSS not a good indicator of distance.

Hence, in order to complete the end-to-end navigation and find the room whereby the beacon is located, it is necessary to provide additional information on room level localization.

This last problem is addressed in BFound by using a sensor enhanced room level localization whereby readings from the light, temperature and humidity sensors are used to identify which room a beacon is in.

4 DESIGN

4.1 Crowdsourced Data Collection

Three kinds of data are collected through crowdsourcing on the smartphone,

- Virtual landmarks
- Relative floor change (staircases/lifts)
- Heading direction

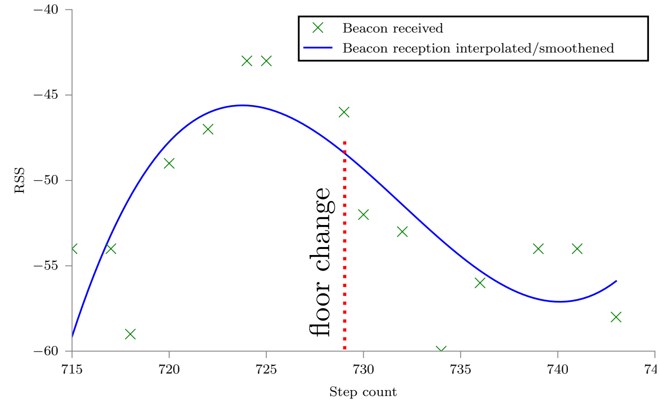

Figure 3: A beacon's location is centered at the local maxima of the RSS over the step count. The RSS is interpolated across the step count using a one dimensional smoothing spline fit.

4.1.1 Virtual Landmarks. Virtual landmarks are determined by the locations of the (strongest) RSS detected for the respective beacons. Each virtual landmark corresponds to a single beacon. However, since RSS is a relatively noisy measure and can be detected over multiple locations and floors, identifying the unique position requires further processing. In order to reduce data processing, a smartphone will consider only the strongest 10% of the RSS collected for each beacon or a distance as threshold when TxPower is available. For each beacon identified, its multiple appearances, with respect to the step count, are grouped together; an instance of a group is shown in 'beacon received' points of figure 3.

Next, we determine the closest location of the beacon to the user's step count. Figure 3 shows the RSS detected for a given beacon over 30 steps with a floor change starting at the 729^{th} step using staircases. Spline interpolation is used to get a curve that fits the RSS readings at the different step count. The beacon's position is placed at the step count corresponding to the maxima of the smooth RSS curve.

Each of the beacons detected appears as a node in the locally generated graph g.

4.1.2 Floor Change Detection. It has been observed that a floor change is easily detected by a device's barometer reading when the device moves across two floor levels [18] over a short period (say ≈ 0.4 mbar for one floor change). Hence, given that the smartphone's barometer sensor is able to detect floor changes based on pressure change, the rate of change of the pressure with respect to the step count gives us an indication about whether staircases or lifts were used. A high rate of change, nearly a vertical line, implies lifts are activated otherwise the staircases are used.

4.1.3 Detecting User's Heading Direction. Detecting a user's heading direction using the smartphone can be challenging because the phone can be hand-held, placed in a bag or kept upside-down in a pocket. Moreover, the magnetic distortions inherent in building structures makes turn detection using the absolute magnetic north unattractive. In order to deal with this, the axis on which the gravity is exerting is leveraged. The changes in the gyroscope readings are more pronounced on the same axis. As such, all relative turns are derived from that particular axis. Hence, we determine the relative change in direction by the user using the integral of the gyroscope axis onto which the gravity is pulling the accelerometer. When the device is upside down, the turn taken is negated as a result of having negative accelerometer readings with respect to the gravity.

4.1.4 Build Local Graph g. By combining and processing the RSS and sensor data, we are able to (1) locate the detected beacons based on step count and the respective floors and (2) derive the vector (direction and step count) between different nodes. The next step is to construct the local graph g. The nodes in g corresponds to the position of the beacons (virtual landmarks). Two nodes are connected in g if the distance between them is less than the window size (in step count) as shown in figure 4. An example of the graph or adjacency matrix constructed is also shown. The edge's distance is the step count between the two node being connected. Direction change detected between connected nodes are also recorded. Whenever two nodes from different floors are being connected, rate of change of the pressure with respect to step count is used to determine (1) whether it is an upward or downward movement, and (2) the mode of movement (staircases/lifts). These details are added as edges' attributes.

4.2 Graph Building on Server

These crowdsourced sub-graphs are merged to form the global navigation graph at the server. The set of common beacons between the master graph G and new sub-graph is leveraged for the merging. Information on the relative floor and pressure differences are also utilized in the merging to ensure that the relative ordering among different floors are preserved.

4.3 Navigation

Navigation is started by downloading the sub-graph of G related to beacons heard. The user search for facilities or services found in the sub-graph and can possibly states his interests. When the desired destination is known, the path, relevant to the user desired interests, is returned. The path is in the form of environment markers, like shops or rooms, that are expected encounters to the destination. While walking, the user's location is updated on the sub-graph based on the beacons' heard, such that he is always given the next expected encounter.

At times, nodes can also be encountered in slightly different orders at different instances of the crowdsensing process. This is expected as the virtual landmarks are RSS-based. The different encountered orders can be kept as separate instances of G. However, the convergence to the correct order, as that in the physical space, is expected as more data is obtained from crowdsourcing.

4.4 Sensor Enhanced Localization

The sensors available on beacons can be used to capture room level events and it is possible to determine if two beacons are in the same room if (1) trends in sensor readings are similar in the same room and different from other rooms, (2) there is some coarse time synchronization among the beacons so that sensor readings can be compared over time. In BFound, synchronization on the granularity of the sampling interval (10s) is sufficient, (3) these trends can be captured and represented efficiently using small amount of data, and (4) the computation required must be minimum since the computation is done locally on the beacon.

BFound utilizes light, humidity and temperature data to perform room level localization.

- Light events, as shown in figure 7a, are very distinctive. Even with very sparse sampling rate, sharp changes in light intensity can be recorded. However, in some settings, there may be multiple light sources in one room that are controlled separately and the changes will be less obvious. In some cases, there could be no lighting events at all (either on or off all of the time) or the light sensor could be covered. Hence, there is a need to incorporate additional sensor modalities like humidity and temperature.
- Unlike the light events, the humidity readings, as shown in figure 7c changes much more gradually and trends' similarity can only be detected over a long time period. Temperature readings exhibit similar behavior.

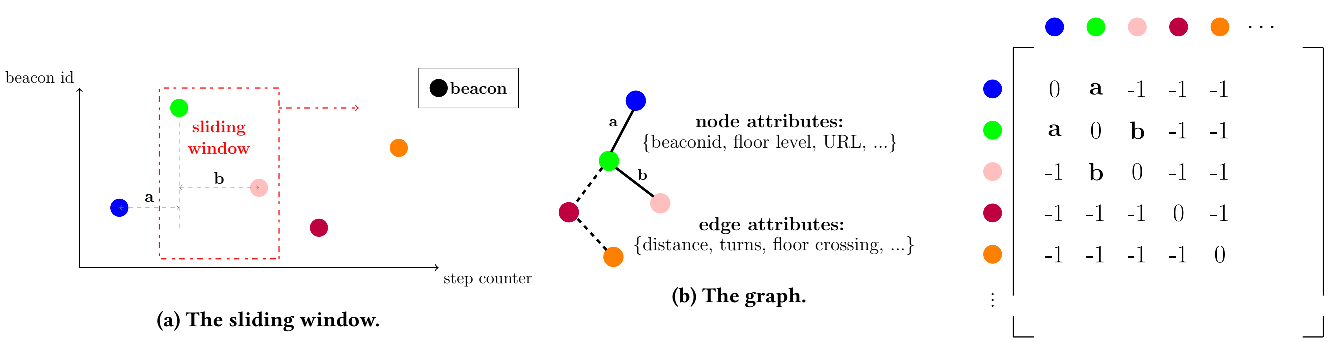

(a) The sliding window.

(b) The graph.

(c) The adjacency matrix.

Figure 4: As the sliding window moves, the graph/adjacency matrix are build up with nodes/edges' attributes. The edge's distance is the step count difference between the respective encounter of the edge's nodes.

Encoding of light, humidity and temperature poses challenges. For light readings, there is no change most of the time and the changes occur over a very short period. On the other hand, temperature and humidity changes slowly and there is a need to find a compact representation of a long time series data that can be used for comparison and classification.

The approach used in BFound is to apply wavelet transform to encode the beacons' sensors data, namely light, temperature and humidity, into a short signature.

4.4.1 Wavelet Transform. Wavelet transform [12] has been used in data/image compression, pattern recognition and noise reduction. In all cases, the transform consists of two parts,

- a low pass filter, to obtain the **approximations**,
- a high pass filter, to obtain the **detail coefficients**. The same transform can be further applied to the output as required. The crucial elements for the compression are the coefficients with highest values.

It is to be noted that (1) the time element is preserved, and (2) after each transform the number of samples from the input is halved to the approximation and coefficient domain. By using only a small set of the coefficients of the transform, a very compact presentation can be obtained. The decomposition is shown in figure 5 and figure 7 shows only the detail coefficients of decomposition.

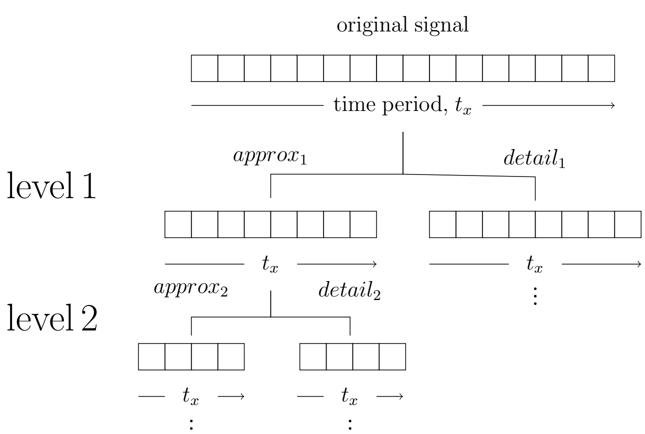

Figure 5: Signal decomposition using wavelet transform.

In BFound, wavelet transform is used to (1) reduce noise from the beacon's sensors, and (2) retain a highly compressed representation of the collected data based on the highest coefficients. As such, the outcome is a signature, in terms of lighting events and temperature/humidity trends, for the room that a beacon is located. Two beacons are likely to be in the same room if they share similar room event signature.

The number of iterations for the wavelet transform required to return a meaningful representation of the signal depends on how fast or slow the signal changes. Therefore, while few iterations are enough to consider light events, the coefficients from multiple iterations are required to enable the alignment of events in the slow varying domain of temperature and humidity. Another parameter to consider is the wavelet to be used. Although there is a panoply of wavelets, the Haar wavelet, shown in figure 6, also known as Daubechies 1 (db1), being one of the most powerful [15], is considered in this work.

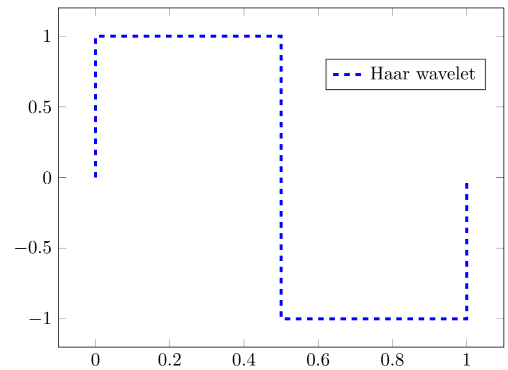

Figure 6: The Haar wavelet

Figures 7a and 7c show the raw sensor data for light and humidity respectively for three different beacons from two rooms. Visually, it is easy to tell that two of the light event changes (room 1 and room 2) are similar. Figure 7b shows the wavelet transforms that say the same.

For the case of humidity, the raw sensor values are harder to interpret since the changes are much more gradual. The corresponding wavelet transforms are shown in Figure 7d. Again, the wavelet transforms indicate that the associated beacons in the top 2 rows are in the same room though it is less obvious.

4.4.2 Room Signature with Wavelet Transform. Given that wavelet transform preserves the relevant variations/happenings of a signal by assigning higher detail coefficients, the latter are useful inputs in the constitution of the room signature. One of the property of wavelet that also make these coefficients useful is that they are localized. Unlike the Fourier Transform that only says that a particular feature (frequency) is there, wavelet transform also says where or when that feature happened. As such, the partial signature from a particular sensor is created using the respective indexes of the set of highest coefficients. In our case, these indexes also represent the times at which these events happened. Thus, it is expected that these times correlate across beacons in the same room. *Therefore, the room signature is composed of the indexes of the highest absolute detail coefficients from the multiple sensors.*

In the case shown in figure 7b, the indexes of light's highest absolute coefficients for the three beacons are:

- Beacon 1 = [462, 1719, -1741, -1742, -1743, 3187, -5745, -5746, -5747, -5748]
- Beacon 2 = [462, -1741, -1742, -1743, -1747, -1750, -2109, 3187, -5745, -5746]
- Beacon 3 = [238, 239, -1908, 2031, -2049, 2740, 2741, -6080, 6326, -6361]

For the case of humidity shown in figure 7d, they are:

- Beacon 1 = [1, -2, -3, 4, -10, 12, 13, -20, -28, 29]
- Beacon 2 = [1, -2, -3, 4, 6, -10, 16, -23, -28, 69]
- Beacon 3 = [-1, 3, 4, 12, -13, 16, -22, 23, 32, -64]

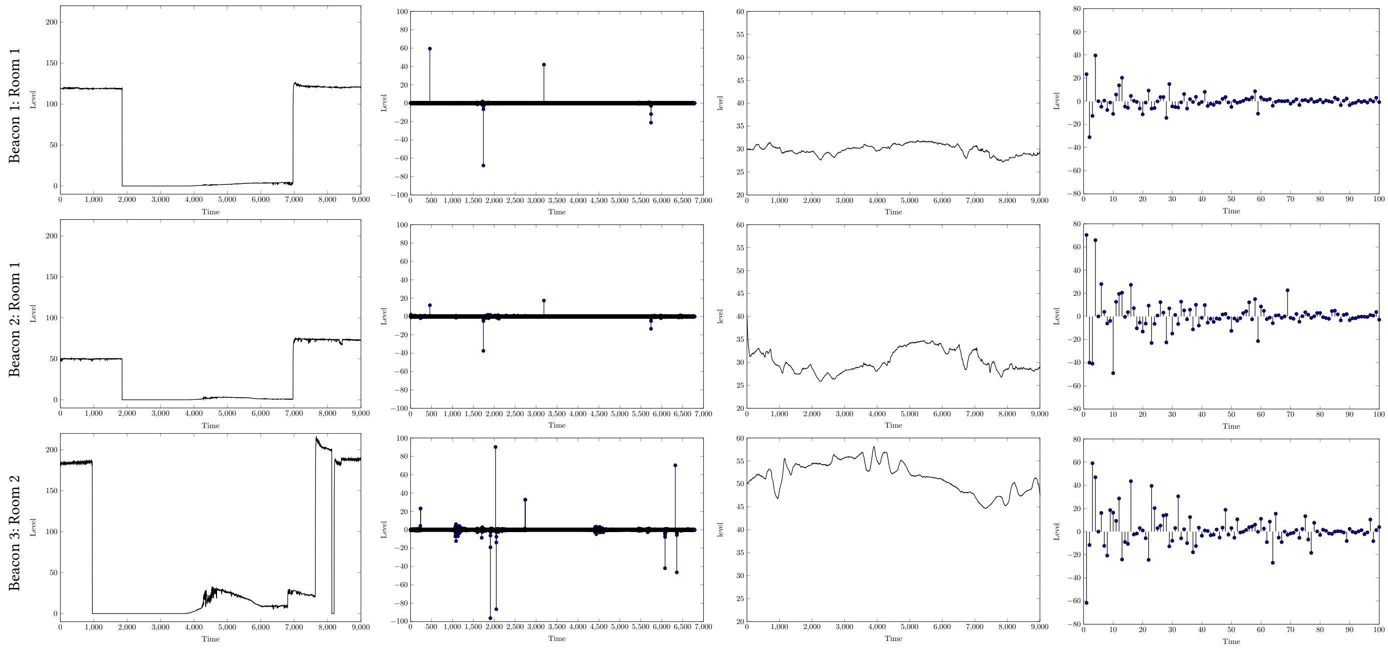

(a) Light: raw sensor data. **(b) Light: detail coefficients.** **(c) Humidity: raw sensor data.** **(d) Humidity: detail coefficients.**

Figure 7: The outcome of light and humidity sensor data going through the Haar wavelet transform at up to level 2 and 11 respectively. For lights, level 1 and level 2 coefficients are shown. For Humidity, an extract of the highest level coefficients is shown as the lower ones are close to zeros. The data from the first two rows are for beacons in the same room.

In both cases, the sign, we used to accompany each index, represents either a corresponding positive or negative coefficient value.

For the beacons in the same room (Beacons 1 and 2), the matches for the light and humidity coefficients are 7 (out of 10) and 6 (out of 10) respectively.

For the beacons in different rooms (Beacons 1/2 and 3), there is no match for the light sensors for beacons from different rooms. Noted that coefficient index with opposing sign implies the changes are dissimilarity rather than similarity. Hence, with humidity data, the match is low (only 2 out of 10) as well.

Although wavelet transform is not computationally expensive, it is worthwhile to compute only the coefficients that are needed. As the beacon is highly resource limited, less computation and less coefficients to broadcast are both important. Unlike with the light sensor values that require up to two levels of wavelet transform to detect sudden changes, the slow varying temperature and humidity data require more iterations to produce more coefficients to unveil the similarity between two beacons in the same room.

The moving up to the higher levels in the wavelet transform is similar to time compression. As such, there is the need to make sure that events are properly indexed so that the outcome from the transform represent the most appropriate time/index that an event occurred across beacons. Figure 8 illustrates the indexing process when two levels of iterations are used.

4.4.3 Localization. Figure 9 shows the decision tree used for localization. In order to localize a beacon, we first need to receive the room signature it broadcasts. This room signature is compared

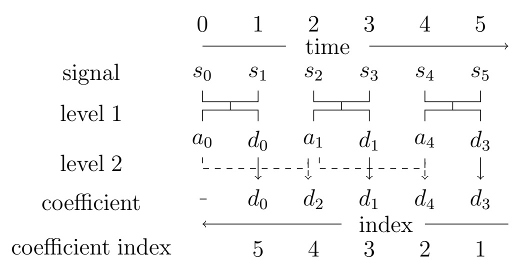

Figure 8: Indexing coefficients for proper representation of events across beacons.

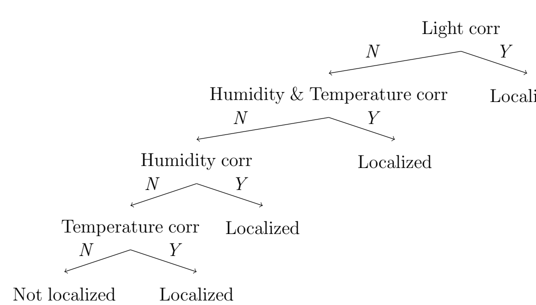

Figure 9: Localizing tree based on signature correlation.

with other beacons with known location that have been received. We localize a beacon by matching it to a beacon (with known location) that has the most similar signature, i.e. highest correlation. Since light is the strongest signature if available, it is first used to determine whether two beacons are in the same room if the signatures match.

If the light signature does not match, we consider the case of having similarities in both the temperature and humidity domain. Finally, consideration is given to the case when correlation is present in either temperature only or humidity only. If all matching fails, the beacon cannot be localized as no other beacon in the same room was found.

5 EVALUATION

There are two parts in the evaluation. First, we present results on crowdsourced data and navigation. In the second part, we present results for sensor based localization.

5.1 Data Collection and Navigation

We evaluated BFound's navigation accuracy in a building deployed with over 100 beacons in rooms, corridors, and common areas, over three floor-levels. The floors are connected by staircases and lifts. The default sliding window is set to 10 steps.

5.2 Quality of Crowdsourced Data

To build the graph G, 4 users collected data over 6 walking paths. In all 6 walks, the ground truth were recorded for comparison. We evaluate our results by comparing the graph build using crowdsensing against the ground truth. The main properties that are checked against are (1) the closest point of the beacon on the step count, (2) the attribute floor level of the node, (3) the nodes' order as they appear on the graph for traversal, and (4) the attribute turns of the respective edges.

5.2.1 Positioning Beacons' Position on Step Count. While walking and logging the beacons heard, the closest points with respect to 22 beacons were recorded. These records were used to show how close/far the beacons are positioned to the actual locations on the step count. The result in figure 10 shows that 80% of the time, the beacon is positioned within 2 step count, i.e. less than 2 meters. Almost all errors are within 3 steps. Moreover, the results are better when using the local maxima of a spline fit instead of using the highest RSS observed for a particular beacon.

5.2.2 Nodes' Ordering. For all 6 walks, the number of nodes encountered and correctness of the ordering are shown in table 1. In terms of ordering of beacons encountered, about 90% of the ordering is correct. For almost all the error cases, misordering of encountered nodes happened when the beacons are quite close to each other.

5.2.3 Expected Turns. Figure 11 shows the CDF of the percentage error in the detected degree of turns. In all walks, despite the phones' carried positions, most of the expected turns were detected without any false positives. The possible drifts of the gyroscope are properly cancelled out. The cases for the false positives are (1) two turns happening within a few steps of each other and are interpreted as one turn, and (2) turn is below the threshold set of 45°.

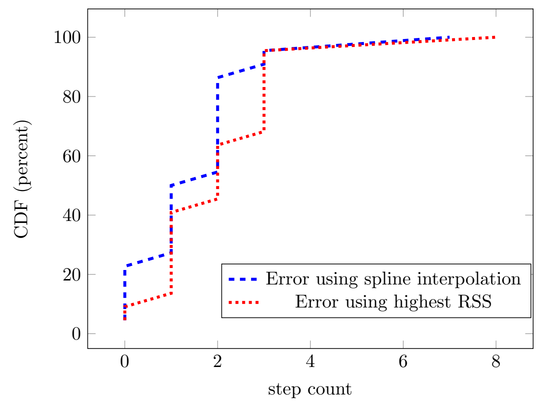

Figure 10: CDF of error in beacon's position (in step count)

Table 1: Nodes' ordering and classification

	Walk 1	Walk 2	Walk 3	Walk 4	Walk 5	Walk 6
Duration (mins)	11.12	12.39	11.13	3.50	4.70	3.30
Distinct nodes encountered	52	50	48	21	25	25
Nodes encountered	78	75	75	25	31	31
Proper ordering	70/78	70/75	68/75	19/25	26/31	22/31

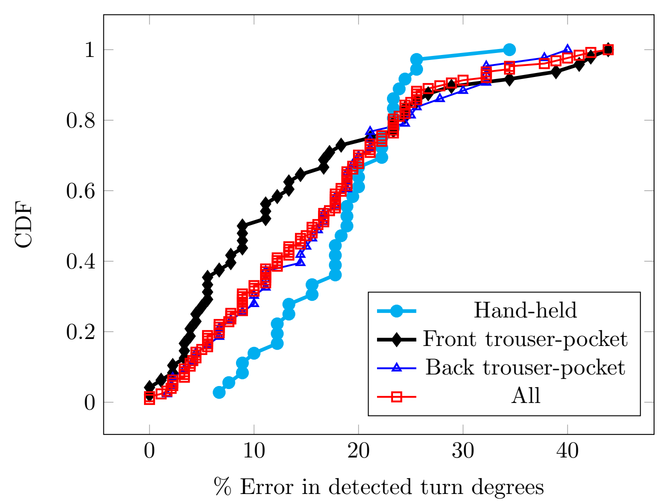

Figure 11: At 90%, the CDF of the percentage error in the detected degree turns in all walks, is less than 30%. 127 out of 130 turns were detected.

5.2.4 Navigation. For the evaluation using the generated global graph G for navigation, 6 users were asked to download G from the server. The list of beacons are annotated with appropriate descriptions such as nearby room number. Each user then chooses a destination beacon from G and asks BFound to provide a list of

Table 2: Navigation trials

No. of beacons involved	<7	7	8	9	12	>13
No. of paths	4	4	1	1	1	2
Required turns	8	19	8	3	6	6
Floor crossing	2	5	2	1	1	0
Out of order beacons (%)	6.7	10.7	0.0	11.1	25.0	24.1

expected beacon encounters along the path to reach the destination. Users were expected to use visual clues, in the form of room names, to navigate. During the navigation process, we compare the graph obtained during crowdsensing with that provided by BFound at the start of the navigation.

In total, 13 paths were taken in the evaluation. We compare the turns required, floor change and beacon ordering between the paths suggested by running a shortest path algorithm on G and the ground truth. The results are shown in Table 2. While the floor crossings were all correct, some turns were not indicated correctly because the turns were too close together. However, the visual clues (room identification) and/or user turning in opposing direction whenever required allowed the users to head the right way. Less than 15% of the nodes encountered were not in the expected order (as per the graph). But they do not affect navigation as these beacons were located close to one another. In all cases, BFound successfully brings the user to the vicinity of the target beacon.

5.3 Beacon's Localization

57 sensors (mixture of TelosB and SensorTag CC2650) with light, humidity and temperature sensors were placed in 16 rooms and corridors (over 5 floors) in two nearby buildings. The location of the beacons were either wall mounted or placed on desk. Sensor data was collected at a sampling rate of 0.1Hz over 24 hours.

5.3.1 Accuracy: Experiment 1. In the first experiment (**experiment #1**), we consider the case whereby we consider the coefficients of the wavelet transform of one node and compares the waveform to that of the other 56 nodes. This is the worst case situation since all beacons are considered.

In the evaluation, only 5, 9 and 9 indexes of highest coefficients of the transform are used for light, humidity and temperature respectively. We select the node with the most similar wavelet coefficients as the node that is in the same room (level or building). Two nodes are considered to be potentially in the same room if a majority of the coefficients matches (at least 3 and 5 coefficients out of the 5 and 9 respectively). The decision tree shown in figure 9 is applied.

The results are shown in table 3 and are grouped into three levels of accuracy, namely: building, level and room. The results also show that the light sensor is a strong indicator but with low recall. This is because lights in many of the rooms did not change at all (either always ON or OFF). On the other hand, the humidity and temperature data are less accurate but give a much higher recall.

At the building level, the accuracy is always 100% as the sensor data shows sufficient differences for easy matching. The accuracy at the floor level granularity is still good, at more than 92% when

Table 3: Experiment #1: ONE against the REST

Sensors	Accuracy			Recall
	Building	Level	Room	
Light	1	1	1	0.193
Humidity	1	0.959	0.749	0.825
Temperature	1	0.929	0.783	0.737
Combined	***1***	***0.927***	***0.818***	***0.965***
Consider rooms with 2 or more nodes				
Combined	***1***	***1***	***0.979***	***0.980***

Table 4: Experiment #2: 10 SPOTS (entrances/lobby/common spaces/rooms)

Sensors	Accuracy			Recall
	Building	Level	Room	
Combined	***1***	***1***	***0.871***	***0.884***

all sensors are combined and recall of more than 96%. The room level accuracy is also high, at 81.8% and with recall of 96.5%

As many of the inaccuracies are due to matching beacons in rooms that have only one beacon, we show the results whereby these "standalone" nodes are removed for the test set and we only consider rooms with at least 2 nodes. Under this condition, there are 11 rooms in the data set. The results shown in table 3 indicate that the accuracy at the room level improves to 97.9% with recall of 98%.

5.3.2 Accuracy: Experiment 2. In the second experiment (**experiment #2**), instead of looking at all the beacons, we consider only sensor nodes within wireless reception. This is more inline with the application scenario whereby a user, after using BFound to navigate to the vicinity, uses the wavelet coefficients broadcast by the nodes to locate the target node by matching the coefficients broadcasted by the target node to the coefficients of a node with known location.

The evaluation was performed from 10 different locations to generate 10 clusters of beacons. These locations were chosen so that most of the (57) nodes from all levels of the two buildings were included in at least one of the clusters. The minimum, average and maximum number of beacons within clusters were 5, 9 and 17 respectively. In total, 93% of the beacons were included.

The results are shown in table 4. As expected, with a smaller set of beacons to match, the accuracy improves to 87.1%

5.3.3 Accuracy: Experiment 3 (Intel Research Lab dataset). In the last experiment (**experiment #3**), we look at an open dataset [19] from Intel Berkeley Research Lab, with 54 sensors deployed in four rooms. This data set has the advantage that it was collected over multiple days. We used data from 15 consecutive days having at least 24 hours of sampling data starting from March 1, 2004.

There are some limitation in the data set. First, there are only 4 rooms whereby there were at least 2 beacons. Next, the sampling intervals were irregular, for example on one day the range is from 33s to 74s with an average sampling interval of 44s.

Table 5 shows the accuracy over 15 days. Note that for each query, we only consider coefficients computed based on data from the last 24 hours. The average room accuracy is 83.7% with a recall

Table 5: Experiment #3: Intel Research Lab (over 15 days)

Sensors	Accuracy			Recall
	Building	Level	Room	
Combined	***NA***	***NA***	***0.837***	***0.944***

of 94.4%. While the accuracy is lower than the experiments using the data we have collected, given the irregular sampling and much noisier sensor data, the accuracy is fairly good.

5.4 Power Consumption

In this section, we measured the power consumption of a beacon (SensorTag CC2650) running the BFound protocol. For each beacon, a cycle includes a periodic sensor sampling of light, humidity and temperature, followed by coefficients' computation and a beacon broadcast. The result is shown in table 6.

When the cycle repeats every 10s (0.1Hz), the average power and average current drawn are 90 μW and 40μA respectively. At this rate, a 1000 mAh battery can sustain this kind of beaconing for nearly three years.

Table 6: Power consumption

Beacon	Broadcast/ Sampling rate (Hz)	Average current (mA)	Battery life (Yrs) 1000 mAh
Broadcast only	10	0.42	0.3
	5	0.21	0.5
	1	0.06	1.9
	0.1	0.03	3.8
Prototype	1	0.14	0.8
	0.1	**0.04**	**2.9**

6 FUTURE WORKS AND CONCLUSION

In this work, both RF and sensors are used for navigation and Things' localization. In order to extend the one-hop nature of beacon, BFound was proposed to effortlessly create a network of beacons using crowdsourcing. Unlike existing heavily trace-driven indoor navigation system, it uses a lightweight graph representation for navigation and searching. The system was also successfully experimented in a building with around 100 beacons that spans over three floor levels.

Unlike any existing RF-based localization solutions, BFound bridges the navigation gap to the Thing of interest by leveraging common low-power sensors available for Internet of Things tag. With experiments conducted using beacons with mounted sensors in multiple rooms, across multiple levels of two nearby buildings, it was shown that IoT devices can be localized with high accuracy to room level. No additional infrastructure support is required. Bfound uses wavelet transform to generate room signatures from sensor data. The proposed mechanism is shown to be very energy-efficient and having high accuracy when tested over more than 15 days of sensor data.

ACKNOWLEDGMENTS

This work was supported in part by the Ministry of Education Tier 1 Grant (R-252-000-621-114), Singapore.

REFERENCES

[1] Mike Addlesee, Rupert Curwen, Steve Hodges, Joe Newman, Pete Steggles, Andy Ward, and Andy Hopper. 2001. Implementing a sentient computing system. *Computer* 34, 8 (2001), 50–56.
[2] Pavithra Babu. 2016. 10 Airports Using Beacons to Take Passenger Experience to the Next Level. https://goo.gl/HM8pGe.
[3] George Bebis, Aglika Gyaourova, Saurabh Singh, and Ioannis Pavlidis. 2006. Face recognition by fusing thermal infrared and visible imagery. *Image and Vision Computing* 24, 7 (2006), 727–742.
[4] Krishna Chintalapudi, Anand Padmanabha Iyer, and Venkata N Padmanabhan. 2010. Indoor localization without the pain. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*. ACM, 173–184.
[5] Jaewoo Chung, Matt Donahoe, Chris Schmandt, Ig-Jae Kim, Pedram Razavai, and Micaela Wiseman. 2011. Indoor location sensing using geo-magnetism. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*. ACM, 141–154.
[6] Ionut Constandache, Xuan Bao, Martin Azizyan, and Romit Roy Choudhury. 2010. Did you see bob?: human localization using mobile phones. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*. ACM, 149–160.
[7] Google. 2016. Eddystone Protocol Specification. https://github.com/google/eddystone/blob/master/protocol-specification.md.
[8] GoogleMaps. 2016. Go inside with Indoor Maps. https://www.google.com/maps/about/partners/indoormaps/.
[9] Ming-Shing Hsieh, Din-Chang Tseng, and Yong-Huai Huang. 2001. Hiding digital watermarks using multiresolution wavelet transform. *IEEE Transactions on industrial electronics* 48, 5 (2001), 875–882.
[10] iBeacon.com. 2016. What is IBeacon? A guide to beacons. http://www.ibeacon.com/what-is-ibeacon-a-guide-to-beacons/.
[11] ibeacontrends.com. 2016. iBeacon Trends. http://www.ibeacontrends.com/category/case-study/.
[12] National Instruments. 2015. Applications of Wavelet Transforms. http://www.ni.com/white-paper/3887/en/.
[13] Allied Business Intelligence. 2016. BLE Beacon Market Remains on Target to Break 400 Million Shipments in 2021. https://www.abiresearch.com/press/ble-beacon-market-remains-target-break-400-million/.
[14] Charles E Jacobs, Adam Finkelstein, and David H Salesin. 1995. Fast multiresolution image querying. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. ACM, 277–286.
[15] K Kim and K Choi. 2014. *Algorithm & SoC Design for Automotive Vision Systems*. Springer.
[16] Ye-Sheng Kuo, Pat Pannuto, Ko-Jen Hsiao, and Prabal Dutta. 2014. Luxapose: Indoor positioning with mobile phones and visible light. In *Proceedings of the 20th annual international conference on Mobile computing and networking*. ACM, 447–458.
[17] Chengwen Luo, Hande Hong, and Mun Choon Chan. 2014. Piloc: A self-calibrating participatory indoor localization system. In *Information Processing in Sensor Networks, IPSN-14 Proceedings of the 13th International Symposium on*. IEEE, 143–153.
[18] Chengwen Luo, Hande Hong, Cheng Long, Kartik Sankaran, and Mun Choon Chan. 2015. iMap: Automatic Inference of Indoor Semantics Exploiting Opportunistic Smartphone Sensing. In *IEEE International Conference on Sensing, Communication and Networking (SECON)*.
[19] Samuel Madden. 2004. Intel Lab Data. http://db.csail.mit.edu/labdata/labdata.html
[20] Shubhi Mittal. 2016. 25 Retailers Nailing it with their Proximity Marketing Campaigns. https://blog.beaconstac.com/2016/02/25-retailers-nailing-it-with-their-proximity-marketing-campaigns/.
[21] Kartik Muralidharan, Azeem Javed Khan, Archan Misra, Rajesh Krishna Balan, and Sharad Agarwal. 2014. Barometric phone sensors: More hype than hope!. In *Proceedings of the 15th Workshop on Mobile Computing Systems and Applications*. ACM, 12.
[22] Radius Networks. 2016. AltBeacon: The Open and Interoperable Proximity Beacon Specification. http://altbeacon.org/.
[23] P Pahl and Padmanabhan VN RADAR. 2000. An In-Building RF-based User Location and Tracking System [C]. *IEEE Communications Socienties* 2 (2000), 775–784.
[24] Anshul Rai, Krishna Kant Chintalapudi, Venkata N Padmanabhan, and Rijurekha Sen. 2012. Zee: zero-effort crowdsourcing for indoor localization. In *Proceedings of the 18th annual international conference on Mobile computing and networking*. ACM, 293–304.

[25] He Wang, Souvik Sen, Ahmed Elgohary, Moustafa Farid, Moustafa Youssef, and Romit Roy Choudhury. 2012. No need to war-drive: unsupervised indoor localization. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*. ACM, 197–210.

[26] Roy Want, Andy Hopper, Veronica Falcao, and Jonathan Gibbons. 1992. The active badge location system. *ACM Transactions on Information Systems (TOIS)* 10, 1 (1992), 91–102.

[27] Han Xu, Zheng Yang, Zimu Zhou, Longfei Shangguan, Ke Yi, and Yunhao Liu. 2016. Indoor localization via multi-modal sensing on smartphones. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 208–219.

[28] Zhice Yang, Zeyu Wang, Jiansong Zhang, Chenyu Huang, and Qian Zhang. 2015. Wearables can afford: Light-weight indoor positioning with visible light. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 317–330.

[29] Yuanqing Zheng, Guobin Shen, Liqun Li, Chunshui Zhao, Mo Li, and Feng Zhao. 2014. Travi-navi: Self-deployable indoor navigation system. In *Proceedings of the 20th annual international conference on Mobile computing and networking*. ACM, 471–482.