Enhancing Indoor Smartphone Location Acquisition using Floor Plans

Niranjini Rajagopal
Carnegie Mellon University
niranjir@andrew.cmu.edu

Patrick Lazik
Carnegie Mellon University
plazik@andrew.cmu.edu

Nuno Pereira
Instituto Superior de Engenharia do Porto
nap@isep.ipp.pt

Sindhura Chayapathy
Carnegie Mellon University
schayapa@andrew.cmu.edu

Bruno Sinopoli
Carnegie Mellon University
brunos@ece.cmu.edu

Anthony Rowe
Carnegie Mellon University
agr@ece.cmu.edu

ABSTRACT

Indoor localization systems typically determine a position using either ranging measurements, inertial sensors, environmental-specific signatures or some combination of all of these methods. Given a floor plan, inertial and signature-based systems can converge on accurate locations by slowly pruning away inconsistent states as a user walks through the space. In contrast, range-based systems are capable of instantly acquiring locations, but they rely on densely deployed beacons and suffer from inaccurate range measurements given non-line-of-sight (NLOS) signals. In order to get the best of both worlds, we present an approach that systematically exploits the geometry information derived from building floor plans to directly improve location acquisition in range-based systems. Our solving approach can disambiguate multiple feasible locations taking into account a mix of LOS and NLOS hypotheses to accurately localize with significantly fewer beacons.

We demonstrate our geometry-aware solving approach using a new ultrasonic beacon platform that is able to perform direct time-of-flight ranges on commodity smartphones. The platform uses Bluetooth Low Energy (BLE) for time synchronization and ultrasound for measuring propagation distance. We evaluate our system's accuracy with multiple deployments in a university campus and show that our approach shifts the 80% accuracy point from $4-8m$ to $1m$ as compared to solvers that do not use the floor plan information. We are able to detect and remove NLOS signals with 91.5% accuracy.

CCS CONCEPTS

• **Information systems** → **Location based services**; **Sensor networks**; Mobile information processing systems; • **Computer systems organization** → *Embedded systems*;

KEYWORDS

Indoor localization, location acquisition, range-based localization, floor plan integration, ray tracing, NLOS detection

ACM Reference Format:
Niranjini Rajagopal, Patrick Lazik, Nuno Pereira, Sindhura Chayapathy, Bruno Sinopoli, and Anthony Rowe. 2018. Enhancing Indoor Smartphone Location Acquisition using Floor Plans . In *Proceedings of International Conference on Information Processing in Sensor Networks (IPSN'18).* ACM, Porto, Portugal, Article 4, 12 pages. https://doi.org/None

1 INTRODUCTION

Over the past decades, we have seen several localization systems emerge that show promise for mobile phones. Some of these systems can achieve sub-meter accuracy, but require a mapping or calibration process that changes over time and often rely on Inertial Measurement Unit (IMU) data to slowly converge on a location estimate. The key to these approaches often comes down to how the systems prune away the state space of possible locations by tracking the user over time. In comparison, beacon-based ranging systems based on time-of-flight (TOF), time-difference-of-arrival (TDOA) or angle-of-arrival (AOA) are capable of rapid location acquisition across a building without *a priori* information. Unfortunately, this often comes at the cost of high beacon densities that can be expensive in terms of hardware and deployment costs. In this paper, we explore a technique that exploits constraints from floor plan information to directly improve range-based systems.

At lower densities, beacon-based ranging systems struggle to perform accurately due to errors from NLOS signals. Figure 1(a) shows an example deployment with high beacon-density and pure LOS that is an ideal scenario often assumed by system installers. Figure 1(b) shows a more realistic scenario with NLOS ranges where typical location solvers will fail to produce the correct location

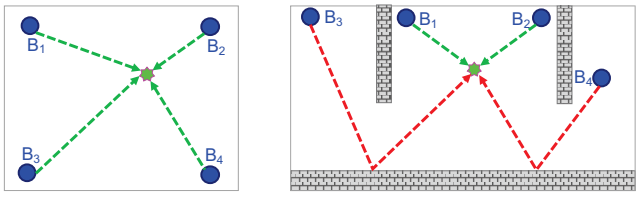

(a) Typically assumed scenario: High beacon-density with LOS

(b) Realistic problematic scenario: Low beacon density with NLOS

Figure 1: Problem Illustration

0-7695-6377-5/18/$31.00 ©2018 IEEE
DOI 10.1109/IPSN.2018.00056

estimate. Common solutions to deal with low beacon density and NLOS fuse the information from beacons with inertial sensors or constrain the motion within the floor plan using particle-filters or similar approaches. Unfortunately, these approaches require the user to walk around and explore the space before they can acquire an accurate lock. In many applications such as augmented reality, way-finding and targeted advertising, quick location acquisition is critical for keeping users engaged with the application.

We present an approach that leverages the geometry of the floor plan to both reduce the density of beacons required and to localize in the presence of NLOS. The floor plan information along with the coverage model of the beacons gives us information about which beacons we expect to be in LOS and NLOS in different regions of the building. The proposed location solver (1) considers feasible hypotheses of LOS and NLOS beacons among the received ranges and solves for locations under each hypothesis, (2) checks for consistency between the estimated location and the assumed hypothesis against the predetermined coverage information (3) selects the most likely hypothesis-location pair. Also, if sufficient beacons are deployed in order to localize without any ambiguity with all LOS beacons, then we can maintain the same accuracy even with additional NLOS signals. The intuition is that NLOS ranging errors caused by reflected signals are positively biased, and if the location is estimated assuming the hypothesis that an NLOS measurement from a beacon is an LOS measurement, the location will be inconsistent. Likewise, a true hypothesis would produce a location estimate that is consistent with the coverage information.

Table 1 compares various indoor localization technologies in terms of their ability to operate on commodity smartphones, acquisition time and accuracy. The accuracy varies widely across literature and system implementations, but insights can be gained in head-to-head deployments like at the Microsoft Indoor Localization Competition [1, 18] where acoustic/ultrasonic systems have demonstrated high accuracy. These systems are compatible with smartphones and can be made inaudible if operated above the human hearing frequency, and are capable of relatively high update rates of 4Hz in [12], 150Hz in [2], resulting in quick acquisition. Building upon the ALPS platform [12], we further reduce the beacon density by designing a time synchronization mechanism using a stream of BLE packets in order to enable TOF ranging (as opposed to TDOA). This requires fewer beacons to cover the same area. Our enhanced ultrasonic platform also provides coarse-grained angle-of-arrival data from a sectored transmitter that can disambiguate certain areas on a map with a single beacon. Figure 2 shows the overall architecture of our system with a gateway that synchronizes multiple beacons that can transmit to one or more mobile smartphones. In this paper, we focus on ultrasonic ranging, but many of the principles generalize to RF systems. TOF RF technologies built on WiFi, BLE and Ultra-wideband (UWB) will eventually be integrated into smartphones.

In summary, the main improvement of our approach over state-of-the-art systems is that we leverage the coverage information of beacons derived from the floor plan to localize with range-based beacons in the presence of NLOS and low beacon density thus enabling accurate location acquisition. Our floor-plan aware techniques can be applied to localization systems that have LOS and NLOS error models similar to acoustic systems (zero mean for LOS,

Technology	Smartphone compatible?	Acquisition <1sec?	Accuracy <1m?
WiFi	y	y/n	n
IMU	y	n	y/n
UWB	n	y	y
Lidar	n	y/n	y
BLE	y	y	n
Acoustic/ Ultrasonic	y	y	y

Table 1: Suitability of technologies for accurate location acquisition on smartphones

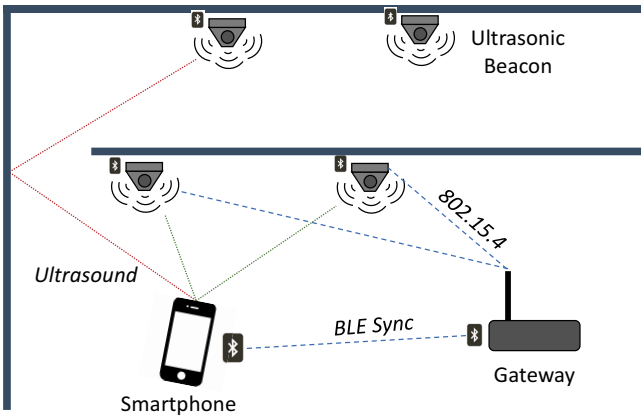

Figure 2: System architecture

positively biased for NLOS), which is the case for multiple emerging TOF RF ranging systems. In support of this approach we provide the following contributions:

(1) A floor-plan aware solver that is capable of working with low-density beacon deployments and improving robustness to NLOS signals
(2) The design and evaluation of an enhanced ultrasonic speaker array platform that can perform TOF ranging with angle information to smartphones.

2 RELATED WORK

This work focuses on the integration of floor plan information, coping with NLOS signals and the design of a new ultrasonic platform. In Section 2.1 we discuss work related to incorporating floor plans and highlight how the problem they solve and their approach is fundamentally different from our work. In Section 2.2 we compare and contrast several acoustic/ultrasonic systems. Finally, in Section 2.3 we discuss prior approaches for NLOS detection and highlight how our work differs.

2.1 Integration of Floor Plan

Approaches that integrate building floor plans in to localization systems broadly do so with Bayesian estimation methods by eliminating possible hypotheses of locations over a period of time as the user traverses the indoor space. They exploit the asymmetric nature of indoor spaces and traces walked. For instance, a particle-filter

System year (20xx)	Ranging type	Smartphone?	Multi-room?	Num. beacons for 2D loc.	NLOS resistance
Bat[2] '01 Cricket[21] '05	TOF	no	yes	3	Limited range, high density outlier in time
BeepBeep[20] '07	TOF	yes (audible)	no	3	Empirical, system specific
SpiderBat[19] '11	TOF + AOA	no	no	1	Angle consistency
[14] '12, [17] '13	TDOA	yes	no	4	None
Guoguo[16] '13	TDOA	yes	no	9	Channel stats
ALPS[12] '15	TDOA for first fix TOF subsequently	yes	yes	4 for first fix, 3 later	Machine learning BLE, acoustic signal statistics
ASSIST[5] '15	TDOA	yes	no	4	None
[6] '15	TOF	no	no	6	Empirical, peak finder system specific
[26] '16	TDOA	yes	no	4	None
Proposed	TOF	yes	yes	1 for corridors, 2 otherwise	Using floor plan

Table 2: Comparison of acoustic/ultrasonic localization approaches

based approach that integrates the information of walls and doors can eliminate unlikely positions of the user [7]. Another approach is to discretize the floor plan, apply a probabilistic model for transitioning between two locations in the floor plan and integrate the motion with a particle filter [8, 27]. These approaches are not suitable for the acquisition problem since they converge on a location estimate over extended periods of time. In contrast, the problem we are solving is estimating location using a single set of range measurements from beacons without having the user to walk.

Another class of floor-plan based schemes use complex 3D ray-tracing (accounting for direct, reflected, transmitted, diffracted path, dielectric constant of materials) to model the signal propagation [10, 11, 25]. In our work, we use ray tracing to determine the likelihood of location being in direct LOS or not direct LOS (NLOS) from a beacon. When a NLOS signal is received, the signal can reflect off any of the walls, the floor or ceiling or any other obstruction not modeled by the floor plan. We do not make assumptions on the path of the reflected signal, only that a NLOS path has a positive bias error.

2.2 Acoustic/Ultrasonic-based Localization

Table 2 compares various acoustic/ultrasonic localization systems. An interested reader can refer to [5, 9, 20] for a more comprehensive coverage. These systems have beacons or devices deployed at known locations that perform ranging to the target device. The target device's location is estimated using trilateration (TOF) or multilateration (TDOA) or triangulation (AOA). The problem we are solving is estimating location using a single set of range measurements from beacons since the goal is to localize as quickly as possible. In Section 4.2 we show that we are able to perform TOF ranging between ultrasonic beacons and off-the-shelf smartphones. Most of the systems have been demonstrated to work in single room setups and assume a deployment and scenario as shown in Figure 1(a) where all regions are covered by three or more beacons. In practice, covering all regions of a building with three or more beacons can render installation and maintenance prohibitively expensive. A low-density deployment where regions are covered by one or two beacons is desirable, but we can no longer uniquely perform trilateration. In addition, these systems suffer from NLOS measurements in multi-room setups as shown in Figure 1(b), where there are two LOS measurements and two inaccurate NLOS measurements. If all beacons are assumed to be in LOS, the estimated location would be incorrect if conventional location solvers are used.

2.3 NLOS Detection

The last column of Table 2 shows various techniques adopted by these systems for combating NLOS. Among the approaches in Table 2, [6, 20] use variants of peak detection based on empirical methods that worked well for their system in terms of detecting LOS and NLOS. These approaches use the received signal time series. Another approach to the NLOS problem is to extract statistics from a full Channel Impulse Response (CIR) waveform like Kurtosis, RMS delay spread, and mean excess delay [16, 24]. Unfortunately, these statistics are highly dependent on the environment. Filtering of measurements over time can be performed, but significantly increases acquisition times and message passing overhead. In contrast, our solution does not assume that the ranging technology has access to the entire CIR or that several measurements are averaged over time.

Residue-based approaches attempt to detect and mitigate NLOS purely based on the range measurements [3, 4, 15]. They iterate through all possibilities of signals and estimate the locations as a function of the residual error derived from each possibility. Examples of functions include the location with least residue, or average of the locations weighed by the residue. Each of these approaches require three or more LOS beacons or the number of LOS ranges to be much higher than NLOS in order for only the true location to

produce the least residue. They fail in low beacon density deployments with high NLOS. This motivates the need to develop location solving techniques to cope with localizing in scenarios such as the case in Figure 1(b), where all received range measurements are treated with equal confidence.

3 COVERAGE-AWARE LOCATION SOLVER

The acquisition problem is defined as follows: Given a set of range measurements
$[B_r = \{B_i, B_j, ..B_k\}, R = \{R_i, R_j, ..R_k\}]$, where R_i is the range received from beacon B_i, and given the location of all beacons in the floor plan B_{all}, we need to estimate the location of the device receiving ranging signals. For instance in the two scenarios shown in Figure 1, we are given the set of range measurements from the four beacons and must determine the most likely location of the receiver.

3.1 Ray-Tracing on Floor Plans

Our approach is based on the observation that we can predict the coverage of acoustic signals with a ray-tracing model due to the inability of sound to pass through walls. Figure 3(a) illustrates the simple ray-tracing coverage model we assume. The beacons are labeled as B_i, and the walls are represented by solid lines. The coverage information is generated automatically by a ray-tracing algorithm operating on floor plans represented in the form of polygons with holes. The algorithm assigns a point to be in LOS with a beacon if the line joining the beacon and the point does not intersect the floor plan polygon. Figure 3(b) validates the ray-tracing model experimentally using data from real-world environments described in Section 5. It shows the distributions of the ranging error between beacons and test points where each range measurement is classified as LOS or NLOS based on the assumed coverage model. In general as expected, the LOS errors are about zero and the NLOS errors are positively biased and environment-dependent. The model is not perfect and we see some inaccuracy described below. First, we see that in the LOS case, there is a slight positive bias in the error due to not capturing all obstructions in the environment such as cubicle partitions, etc. that are not in the floor plan. Second, in the NLOS case, we see that there are locations with close to zero error since we have not accounted for diffraction around doorways and corners. In future work, we would like to take into account a more realistic modeling of the environment, but for now this model is simple and closely resembles the true coverage of the beacons.

We extend this ray-tracing model to a multi-beacon deployment by combining the coverage information of the individual beacons. The information we get is the mapping between all the regions of the floor plan and the beacons in LOS of these regions. Figure 3(c) illustrates the resulting coverage. We see that the floor plan gets partitioned into six disjoint zones based on which beacons are in coverage. For example, the blue color shaded region marked with $H_3 = \{N, N, L, L, N\}$ is the region where B_3 and B_4 are in LOS and B_1, B_2 and B_5 are in NLOS. We store this information offline as Ψ_i: set of all X that are covered by the beacons in LOS in the set H_i.

3.2 Intuition for Proposed Approach

We illustrate our approach with two example scenarios in Figure 4 where we are required to localize in low beacon density in the presence of NLOS. This floor plan has a combination of hallways that are covered by single beacons and open areas that are covered by two beacons. In the first scenario in Figure 4(a) two LOS and two NLOS ranges are received. In the second scenario in Figure 4(b), one LOS and one NLOS ranges are received. The true location X_{true} and the minimum means square error location assuming all received ranges are LOS signals, $X_{minerror}$ are marked on the figures. We can clearly see that utilizing all the ranges as LOS results in incorrect location estimates.

The intuition behind leveraging the floor plan information is that we can check for consistency between the received measurements and the beacon coverage information. In the scenario in Figure 4(a), the resulting location estimate assuming all the beacons are in LOS is $X_{minerror}$. However, $X_{minerror}$ is in NLOS of B_3 and B_4.This produces an inconsistency with the hypothesis that yielded this location estimate. Another hypothesis we can consider is that B_1 and B_3 are in LOS and B_2 and B_4 are in NLOS. However, according to the coverage information, there is no region of the floor plan that is in LOS of both B_1 and B_3, and hence we can dismiss this hypothesis. Next, say there exists some hypothesis that results in the location being estimated at the point marked as X_4. We see that the distance between X_4 and B_2 is much higher than the range received from B_2, and B_2 is in NLOS of X_4. This implies that the NLOS range measurement has a negative error. This is highly unlikely given that acoustic NLOS measurements have a positively biased error. Hence the hypothesis that estimated X_4 as the true location generates an inconsistency. In the remainder of this section, we show a systematic approach for integrating the coverage information by checking and validating multiple hypotheses to yield a location estimate.

3.3 Localization Algorithm Preliminaries

In this section, we define the notation, state our assumptions and then discuss the conditions that the algorithm uses for checking for consistency with the coverage.

Notation:

- B_{all}: Set of all beacons in the floor plan
- B_r, $\widetilde{B_r}$: Set of beacons from which measurements are received, not received respectively; $B_r \cup \tilde{B_r} = B_{all}$
- $B^L(X)$, $B^N(X)$: Set of beacons in LOS, NLOS of location X
- H_i: Hypothesis i. A hypothesis is a partition of the beacons in to a set of LOS and NLOS beacons.
  $H_i : B_{all} = H_i(B^L) \cup H_i(B^N)$
- $H_i(B^L)$: Set of LOS beacons in hypothesis H_i
- $H_i(B^N)$: Set of NLOS beacons in hypothesis H_i
- Ψ_i: Set of all locations in a floor plan that have the exact set of LOS and NLOS beacons as the partition under hypothesis H_i
- $X_{minerror}$: The minimum mean square error estimated location assuming all beacons are in LOS
- X_{true}: The true location of the receiver

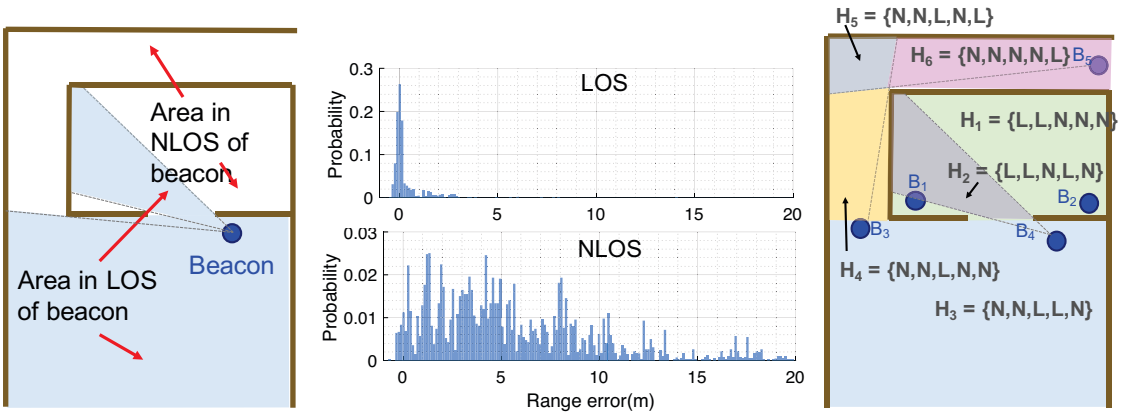

(a) Coverage model of single beacon

(b) Experimental accuracy of model

(c) Coverage model of multiple beacons

Figure 3: Ray tracing coverage model

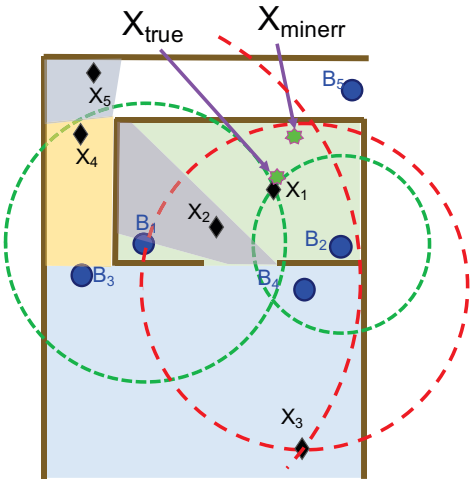

(a) Scenario 1 (B_1, B_2: LOS; B_3, B_4: NLOS)

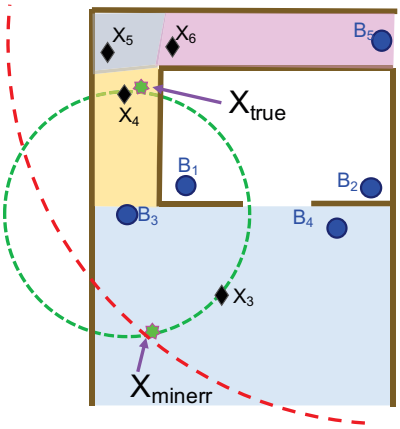

(b) Scenario 2 (B_3: LOS; B_5: NLOS)

Figure 4: Examples to illustrate coverage-aware solver

Assumptions:

- A NLOS measurement has a positive error. This is because NLOS signals are caused by reflections which always take a longer path than the true path, as experimentally validated in Figure 3(b).
- A LOS measurement has zero-mean error.
- The beacon coverage model is deterministic. A location is either in LOS or NLOS of a beacon.

Conditions for satisfying consistency :
If X_i is a location estimated under a hypothesis H_i, in order for this hypothesis-location pair to produce consistency between the beacon coverage model and the received range measurements, the following conditions should be satisfied.

- C_1: Consistency with beacon coverage
  $H_i(B^L) \cap B_r \subseteq B^L(X_i)$
  $H_i(B^N) \cap B_r \subseteq B^N(X_i)$
- C_2: Consistency with NLOS error model
  $\forall k$: $B_k \in H_i(B^N) \cap B_r$, $[R_k - \parallel B_k - X_i \parallel > 0]$

The first two conditions under C_1 check that the LOS/NLOS beacons as assumed by the hypothesis are in LOS/NLOS of location X_i. Since we only receive measurements from the set B_r, we perform an intersection with this set. There can be additional beacons in LOS or NLOS of X_i from which no range measurements are received, hence we use subset and not strict equality in these two conditions. The second condition under C_2 checks that the NLOS errors are positively biased as per our assumption. We check this condition for all the beacons from which a range is received that are also in NLOS according to the hypothesis H_i.

3.4 Localization Algorithm

We describe two location solver algorithms for implementing the approach. The first coverage-aware location solving algorithm evaluates all feasible hypotheses (*CA-all*) and then selects the most likely location-hypothesis pair among the ones that are consistent with the beacon coverage model. While evaluating all hypotheses, in general the hypotheses that assume the short-range beacons to be in NLOS get eliminated. This motivates our second algorithm

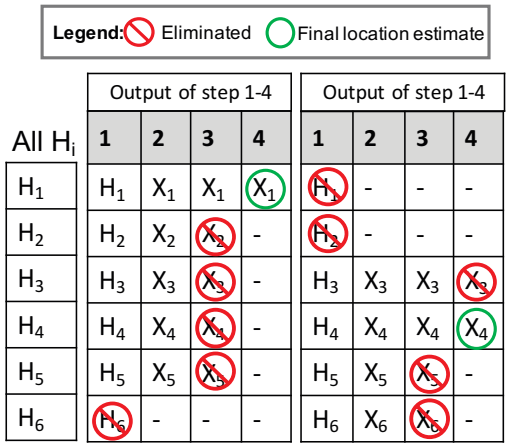

(a) Scenario 1 (b) Scenario 2

Figure 5: Step-by-step results of proposed *CA-all* solver for the scenarios in Figure 4 and hypotheses from Figure 3(c)

that ranks all the hypotheses such that a hypothesis with a shorter range LOS beacon gets a higher rank. It then iteratively checks the ranked hypotheses in sequence till one of them satisfies a likelihood criteria. We refer to this as the *CA-short* algorithm since it prioritizes *shorter* ranges over longer ranges and evaluates fewer hypotheses, and as a result, is *shorter* than the *CA-all* algorithm.

Algorithm 1: *CA-all* **:**

- **Step 1**: Enumerate all feasible hypotheses
- **Step 2**: For each hypothesis in Step 1, find the most likely location that is consistent with the coverage model
- **Step 3**: For each location estimated in Step 2, check if it is consistent with NLOS error model
- **Step 4**: If more than one hypotheses are consistent after Step 3, select the most likely hypothesis based on a likelihood that is a function of the residue LOS error and number of LOS beacons from which no measurements were received.

Next, we elaborate on the implementation of the algorithm with the examples in Figure 4(a) and Figure 4(b). The tables in Figure 5(a) and Figure 5(b) show the results of the algorithm after each of the four steps.

Step 1: The first step is to find all feasible hypotheses. Any hypothesis that includes at least one of the received range measurements as an LOS beacon is a feasible hypothesis. For Scenario 1, since measurements are received from four beacons, $B1$-$B4$, among the 6 possible hypotheses, all except H_6 are feasible. For Scenario 2, any hypothesis that includes B_3 or B_5 as an LOS beacon is feasible. The hypotheses H_3, H_4, H_5, H_6 are selected for Step 2.

Step 2: In Step 2, we estimate the most likely location for each hypothesis. The location has to satisfy consistency with the coverage model (condition C_1). For each hypothesis, the measurements from the beacons in LOS as per this hypothesis is used to estimate the location with minimum error. While solving under hypothesis H_i, we only need to evaluate the locations that belong to the set Ψ_i. For example, in Step 2 of Case 1, while solving for hypothesis H_3, which has only B_3 and B_4 in LOS, any location outside the blue region marked by H_3 would not satisfy either C_1 or C_2. In other words, each location in the floor plan is evaluated only once, ie. if and when the hypothesis that it satisfies is being evaluated. For hypothesis H_i, the minimum error location estimated is

$$X_i = \underset{X}{\text{argmin}}\left(\frac{1}{N}\sum_k (R_k - \parallel B_k - X \parallel)\right)$$
$$\forall k : B_k \in H_i(B^L) \cap B_r, X \in \Psi_i$$

The black diamond-shaped markers in Figure 4 show the most likely estimate for each hypothesis.

Step 3: In Step 3, we check for consistency with the NLOS error model (consistency check C_2). In Scenario 1, consider the location X_4 under hypothesis H_4. It violates the consistency check with respect to beacon B_2 since $B_2 \in H_4(B^L) \cap B_r$ and the received range measurement $R_2 = \parallel B_2 - X_{true} \parallel$ is lesser than the distance between X_4 and B_2 ie. $\parallel B_2 - X_4 \parallel$. In the same manner, X_2, X_3 and X_5 are also eliminated. In Scenario 2, at Step 3, hypothesis H_6 with location estimate X_6, and hypothesis H_5 with location estimate X_5 are eliminated.

Step 4: Practically, the probability of missing a measurement from an LOS beacon is very low. It would occur if the beacon was temporarily blocked, or the data was dropped. Though it depends on the system, we can safely assume that the probability of missing a range from an LOS beacon is lower than the probability of receiving a range from an LOS beacon. Though we cannot estimate this probability since it depends on the environment, we introduce a metric $p(missing_{LOS})$ and assign it a low value to penalize a location with a higher number of missing LOS measurements. We use the number of missing LOS measurements $\#missing_{LOS}$, which is found by the number of beacons in the set $H_i(B^L) \cap \widetilde{B^r}$ for hypothesis H_i. We empirically assign $p(missing_{LOS})$ a value of 0.1 in our implementation. This implies, a location estimate with 2 missing LOS beacons is assigned a likelihood 0.01 times another location that has no missing LOS beacons. The second empirical metric is the residue LOS error. The residue error for location X_i which is estimated under hypothesis H_i is:

$$residue_{LOS}(X_i) = \frac{1}{N}\sum_k (R_k - \parallel B_k - X_i \parallel)$$
$$\forall k : B_k \in H_i(B^L) \cap B_r$$

We assign the likelihood of a location X as being proportional to $e^{-residue^2_{LOS}(X)}$, which is equivalent to assuming that the true residue is drawn from zero-mean Gaussian process. This is the empirical metric we adopt to weigh the LOS residual error and the number of missing LOS beacons. Ideally, the number of beacons used for estimation, the geometry of beacons, the likelihood of a beacon being blocked based on distance from beacon, and the environmental factors should be taken into account for accurate modeling.

At the end of Step 3, if we have more than one consistent hypothesis, we assign a likelihood for the locations and select the hypothesis-location pair with highest likelihood:

$$L(X) = \exp^{-residue_{LOS}{}^2(X)} \times \#missing_{LOS}(X)^{p(missing_{LOS})}$$

In Scenario 2, by the end of Step 3, we have two consistent hypotheses X_3 and X_4. Both have zero residue error. Among them, X_3 is

less likely since it has a missing LOS range measurement from B_4, and hence we select X_4.

Algorithm 2: *CA-short* **:**
The second algorithm we propose is a variant that is motivated by the observation that several hypotheses get eliminated during Step 3. In Figure 4(a), we see that beacon B_2 has the shortest range and the four hypotheses that got eliminated in Step 3 were due to producing an inconsistency with this range measurement. This motivates us to sort the hypotheses based on the range measurements since the hypotheses with the *short*est range beacon in LOS are more likely than the hypotheses with the longest range beacon to be in LOS. In terms of computational complexity, we have the additional step of sorting the hypotheses but we end up evaluating fewer hypotheses for Step 2 and this algorithm is *short*er than *CA-all*.

- **Step 1a**: Enumerate all feasible hypotheses
- **Step 1b**: Rank them based on received ranges
- **Step 2**: For highest rank hypothesis, solve for location
- **Step 3**: For the location estimated in Step 2, check if it is consistent with NLOS error model and satisfies the likelihood criteria. If not, then go to Step 2 and evaluate the hypothesis next in rank, and iterate till stop criteria satisfied.

We sort the hypotheses such that the hypotheses with the shortest range beacon in LOS has the highest rank. Among them, the hypotheses with second shortest range beacon in LOS have the highest rank, and this repeats. However, we do end up with multiple hypotheses with the same rank. Among them, we give higher rank to the hypothesis with least number of LOS beacons that are not in the received range set. The intuition for this is, the locations with higher number of missing LOS beacons get penalized while estimating the likelihood. Hence we use this criterion while sorting. Finally, we check whether each hypothesis satisfies a likelihood criteria. In our implementation, we assigned a value of 0.91 as a stop condition, since a location with $30cm$ residue and no missing LOS beacons gives a likelihood of 0.91 based on our likelihood function $L(X)$.

4 ULTRASONIC BEACON DESIGN

In this section, we discuss the design of a new ultrasonic beacon platform that complements our solver to reduce beacon density in two ways. First, it provides a sectored acoustic array that the location solver can use to select among multiple feasible locations. Second, we show that it is possible to recover a clock with enough accuracy using precisely timed BLE advertisement packets to perform direct TOF ranging as opposed to TDOA that is more commonly used in multiple receiver systems.

Figure 6(a) shows the hardware design of our beacons and gateway board. The platform, shown on the left, is based on a multi-standard BLE and IEEE 802.15.4 SoC (TI CC2650) connected to a $192kHz$ audio codec (running at $48kHz$), a MEMS microphone (for inter-beacon ranging or for future uses like ultrasonic motion detection) and an array of speakers connected to two Class D speaker amplifiers. The hardware can transmit two arbitrary sound waveforms up to $80kHz$ from one to four speakers simultaneously. The gateway board, shown on the right, contains similar hardware with the addition of a IEEE 802.15.4/BLE power amplifier and an FTDI USB-to-serial interface for connecting with a computer. The beacon nodes are synchronized using IEEE 802.15.4 from the gateway devices and then broadcast BLE packets that can be used to trigger normal Bluetooth proximity services on mobile phones. These wakeup signals can in turn begin decoding ultrasound for improved accuracy.

4.1 Sectored Speaker Array

Sectored transmitters have long been used to provide AOA information to location tracking systems. Based on the relatively long wavelength of our acoustic signals, beam-forming approaches for accurately determining angles would require a transmitter baseline on the order of meters which is quite obtrusive for most installations. Instead of beamforming, we use a sectored array that solves two problems: (1) we can more easily generate uniform ultrasonic coverage and (2) we can accurately estimate the coarse direction to the beacon just based on signal strength metrics.

Figure 6(b) shows the horizontal beam pattern of our speaker obtained in an anechoic chamber. The beam pattern shows the receiver correlation magnitude of $20.0 - 21.5kHz$ ultrasonic chirps transmitted at various angles in the horizontal plane to the microphone. This particular speaker was selected for its wide beam patterns across our desired frequency range supporting approximately a 90-degree sector. Upon testing, we realized that transmitting the same signal causes interference (combing) between adjacent speakers which significantly distorts the signal at boundaries. This indicated that we needed to orthogonally code or time multiplex neighboring speakers to avoid potential interference. For this reason, selecting four sectors in our array provides the best uniform coverage with enough angular resolution to enable hypothesis pruning while maximizing update rate.

We adapt the signal and demodulation approach from [12] that utilized ultrasonic chirps that are just above human hearing range, but can still be detected by mobile devices like smartphones, tablets and computers. Since upchirps and downchirps are primarily orthogonal, we can transmit both at the same time without interference. Each chirp in our system is sized to be $110ms$ in length with $2.6ms$ between successive chirps. To support four sectors, we can shift one pair of up/down transmissions in time with a neighboring beacon. Each beacon is given a unique time slot that is $350ms$ in length to transmit its ranging signal as described in the next section. It is also worth noting that the overlapping transmissions are not only orthogonally coded, but also transmitted in opposite directions in the horizontal plane. Figure 6(c) evaluates the performance of detecting the sectors in a real-world deployment with 6 beacons. The test points are marked in grey, and are connected to the sector-beacon pair with highest RSSI. We see that in general it selects the closest sector of the beacon. We observe that due to variation in RSSI with angle of sector, it does not always select the closest beacon. In our experiments, we accurately selected the correct 90° sector 80% of the time and were within 180° of the center of the correct sector 98% of the time.

Incorporating the sector information into the solver alters the beacon coverage model, but does not alter the localization algorithm. With sectors for each beacon, each beacon's coverage is partitioned into disjoint regions. However, while localizing, we

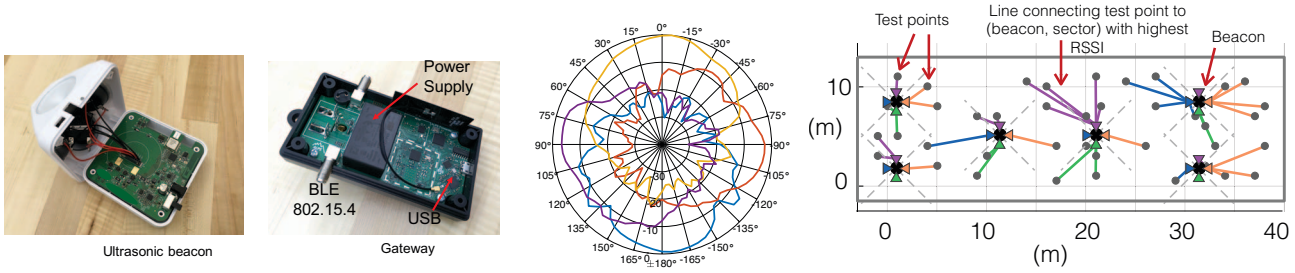

(a) Beacon and gateway hardware (b) Speaker array beam pattern (c) Real world deployment to characterize sectors

Figure 6: Ultrasonic beacon and performance of sectors

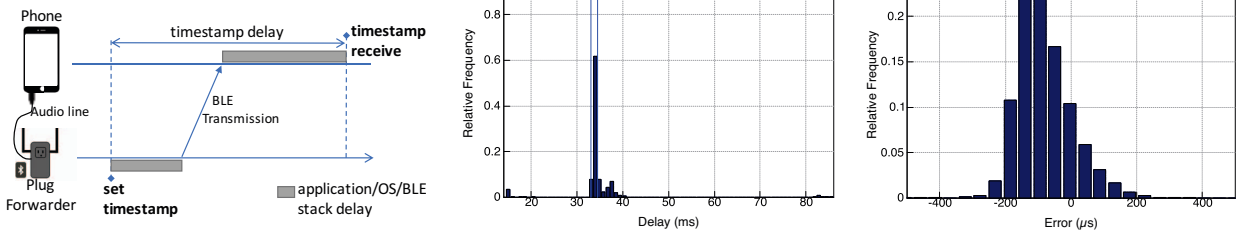

(a) BLE Synchronization experimental setup (b) Timestamp delay (c) Synchronization error

Figure 7: BLE synchronization experiments with iOS device

first detect the sector based on the highest RSSI among the four speakers. Subsequently, we only need to consider the hypothesis corresponding to the LOS coverage of that sector. Though the total number of feasible hypotheses increases, since the sectors of a beacon are disjoint, a much smaller region of the floor plan will be searched while estimating the location. In our real-world evaluation in Section 5, we did not include the sector information since our objective is to evaluate the floor-plan aware solver, and the sectors information is specific to our platform. The benefit of sectors can be seen in deployments much sparser than the deployments we have evaluated in this paper. For instance, small rooms or long hallways would only require a single beacon placed in the center of the room.

4.2 BLE time synchronization

Time synchronization between beacons and the smartphone is challenging due to the timing uncertainties in the smartphone. The authors in [13] developed an acoustic TDOA synchronization scheme that achieved an average accuracy of $720\mu s$ when a smartphone is in LOS to several acoustic beacons and the synchronization accuracy is in the order of several milliseconds when NLOS signals are present. In this work, we look at directly using controlled BLE advertisement packet arrival timing to tightly synchronize between an application running on a smartphone and beacons. We design and evaluate a software Phase Loop Lock (PLL) that is able to recover the clock from our beacons after just a few cycles. The shift from TDOA to TOF significantly simplifies the design and reduces the beacon density.

In order to evaluate BLE synchronization, we first profile the jitter in message arrival and then evaluate the absolute delay. We modified the firmware on our beacons to transmit BLE advertisement packets at a precise $50ms$ interval. As shown in Figure 7(a), we connected the output of the BLE beacon over a wire directly into the audio input of an iPhone 6. By processing the audio signal on the phone, we can measure the time from the timestamp to reception with a maximum error of $42\mu s$ (two audio samples), which validates the claims made in [13]. Figure 7(b) shows the delay distribution of 2000 update requests, where we can see that 75% of the values are within $[33, 35]ms$, moreover, by analyzing the values, we found that the interquartile range is $[33.596, 34.496]ms$. Next, we implement a PLL in sofware using BLE time stamps and the operating system time within iOS. The first exchange between the phone and the forwarder allows us to get an initial estimate of the clock offset between the two devices. We configure the forwarder to respond to data read requests using a connection interval of $50ms$. Next, we periodically perform timestamp requests (as shown in Figure 7(a)), and update our estimate of the offset based on a error-proportional adjustment. After we reach a steady-state (determined by a minimum number of iterations of these adjustments and achieving an estimated error below 1 ms), we start a median filter on the offset estimation. To test our clock synchronization, we used the previous setup with the GPIO connected to the audio as a ground truth of the timestamp timing. Based on this, we computed the difference between the time at which the timestamp was set and the estimated time using the median offset. Figure 7(c) shows this resulting error during the steady state of our clock synchronization scheme. The synchronization error is well under $1ms$, and 96% is within $[-200, 200]\mu s$. According to data from [13] the clocks on

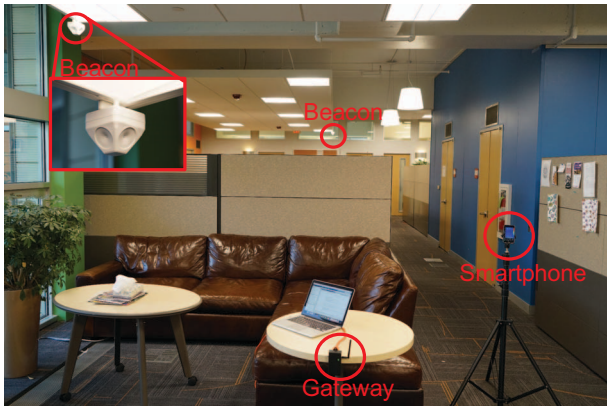

Figure 8: Experimental setup for data collection, with beacons, smartphone and gateway.

iPhones are stable enough that they can remain synchronized to below one millisecond for tens of minutes.

In this section, we evaluate the performance (localization accuracy and the accuracy of detecting NLOS) of the floor-plan aware solver in real-world deployments. We also discuss practical aspects related to the proposed approach such as the trade-off between increasing LOS range and decreasing the amount of NLOS, the number of hypotheses evaluated and effect of environmental factors.

Description of deployments: The experimental setup is shown in Figure 8. The nodes were installed below ceiling tiles and the phone was placed on a tripod. We deployed up to 7 ultrasonic nodes in four different environments on Carnegie Mellon University's campus. Figure 9 shows the four environments (labeled $F1$-$F4$) with their floor plan, the position of the beacons (green circles) and test locations (red dots).

5 EVALUATION

The modeled floor plan is outlined by solid grey lines and the regions with low beacon density (1 or 2 beacons) are shaded in grey. The beacon locations were determined manually. We adopted guidelines from [22], which shows a systematic approach for beacon placement and tried to optimize for high coverage, good geometry, and low number of beacons. Table 3 shows some of the characteristics of the deployments, which inform us about the complexity of the floor plan, beacon coverage density and amount of NLOS. $F1$ includes an open lounge and kitchen area (in the center), a large classroom (on the right), open areas leading to offices and hallways. Among the floor plans this has the maximum number of beacons leading to the maximum number of feasible hypotheses. It also has a high amount of NLOS (33%). $F2$ has four connected hallways with low geometrical complexity and as few as 7 feasible hypotheses. Among the floor plans, this has the highest beacon coverage density and all regions could be localized without ambiguity in the absence of NLOS. $F3$ has long hallways that are mostly covered by a single beacon and includes a kitchen area which had areas of 2 or 3 beacon coverage. Among the floor plans this has the least beacon density. Floor plan $F4$ includes a large conference room and open broad hallways around it.

Solvers evaluated: We evaluate and compare the following solvers:
(1) **GD:** Gradient Descent: This is a common method to estimate location and is most efficient in terms of computation.
(2) **GS:** Grid Search: This solver estimates the minimum mean square error location in a rectangular box surrounding the floor plan. This serves as our baseline solver in the absence of floor plan information.
(3) **FP:** Grid Search in Floor Plan: This approach searches for the minimum error solution within the floor plan. This is the baseline for comparing against using the floor plan to constrain the location without using the beacon coverage information.
(4) **CA-all:** Proposed coverage-aware solver which checks for all hypotheses and then selects one.
(5) **CA-short:** Proposed coverage-aware solver that sorts the hypotheses based on the ranges.

5.1 Localization Accuracy

Figure 10 shows the Cumulative Distribution Function (CDF) of the localization accuracy of each solver across all floor plans. Overall, using the floor plan to confine the solutions (*FP*) improves the performance over not using the floor plan (*GD, GS*), and integrating the coverage information (*CA-all* and *CA-short*) outperforms the other solvers across all deployments. Some key observations are:
(1) *The coverage-aware solver is able to maintain 80-percentile error of 1m under different environments.* This can be seen in the four CDFs.
(2) *Under high-NLOS conditions, the coverage-aware solver significantly improves performance.* $F1$ has wide open areas with beacons located around corners, resulting in a third of the measurements to be in NLOS, as seen in Table 3. We see that the 80-percentile localization accuracy improved from 6m to 1m by integrating the coverage information.
(3) *In low-density deployments, the coverage-aware solver significantly improves performance.* Among the floor plans, *F3* has the lowest beacon density and most areas cannot be localized without ambiguity even under pure-LOS conditions. We see that the 80-percentile localization performance improved from 8m to 1m with the coverage-aware solver.
(4) *Under low NLOS and sufficient beacon density to localize with LOS without ambiguity, the coverage-solver may not provide much improvement over using only the floor plan.* We observe that F4 did not benefit much from the coverage-aware solver as compared to using only the floor plan. This is because it had few NLOS measurements in reality. The classification of 22% of measurements as NLOS is due to inaccuracy in our ray-tracing model by not accounting for diffraction around corners. This floor plan has two thin partitions that we modeled as solid walls and the NLOS measurements through these had low error and the *FP* solver benefited from using them and the coverage aware solvers eliminated these measurements.

5.2 NLOS/LOS Detection Accuracy

The *TP* and *TN* columns of Table 3 show the accuracy of detecting NLOS and LOS correctly. The *CA-all* solver has a *TP* of 91.5% and *TN* of 90.5% on average. The *CA-short* solver is biased towards a higher *TN* of 97.8% at the expense of a lower *TP* of 77%. This is

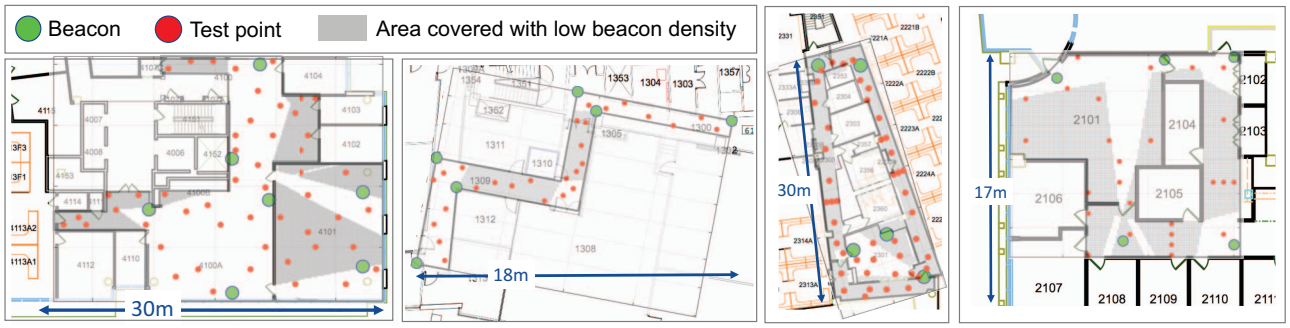

(a) F1 (b) F2 (b) F3 (b) F4

Figure 9: Real-world deployments

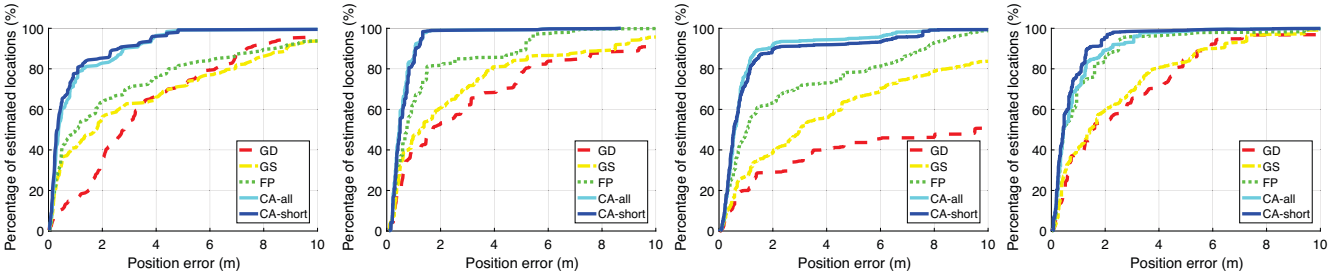

(a) F1 (b) F2 (c) F3 (d) F4

Figure 10: Localization performance in real-world deployments

because the *CA-short* solver checks for the hypotheses in iteration by giving higher weight to hypotheses that have higher number of LOS beacons. Hence, it is biased towards detecting signals as LOS. We also observe that the performance of of *CA-short* to be sometimes better than *CA-all*, which is unexpected since *CA-all* evaluates all hypotheses and *CA-short* evaluates only a subset of the hypotheses. This happens when both approaches have selected an incorrect hypothesis. This occurs when (1) there is only one LOS beacon in range located around a corner such there is a NLOS region closer to the beacon than the LOS location (2) an NLOS signal from another beacon diffracts around a corner/doorway causing NLOS range error close to zero. In these situations, the *CA-short* is biased towards misclassifying the NLOS as LOS, which provides lower error than *CA-all*, which converges on a consistent hypothesis with close to zero residue by misclassifying the LOS range as NLOS, discarding it and localizing with a single NLOS beacon. These cases tend to be rare in practice and their occurrence can be reduced with better placement of beacons.

5.3 Trade-off between LOS and NLOS Performance

In this section, we discuss the practical trade-off between increasing LOS range and decreasing the amount of NLOS signals. Several range-based systems estimate the time-of-arrival (TOA) in either the received channel impulse response or the result of matched filtering of received signal with transmitted signal. The TOA is estimated by selecting a peak in the received signal based on the RSSI using algorithms such as first peak above noise floor, highest peak above noise floor, first peak within a window before the highest peak, etc. A detailed survey of the approaches can be found in [5, 16, 23]. However, a persistent challenge is to determine the threshold of RSSI for peak detection. This is illustrated in Figure 11(a) with three types of channels. Channel A represents a short range LOS path. Channel B represents an NLOS path of the same range. Channel C represents a long range LOS path, and has comparable signal strength to Channel B. If a high threshold for peak detection is selected, the NLOS signal in Channel B and the

Characteristics of real-world deployments					Performance of solver in detecting NLOS				
Floor Plan	# vertices	# beacons	# feasible hypotheses	% of NLOS	CA-all		CA-short		
					TP	TN	TP	TN	# hypotheses evaluated
F1	33	7	39	33	0.94	0.86	0.88	0.95	23
F2	10	6	7	24	0.96	0.97	0.82	0.99	2
F3	73	5	19	20	0.83	0.91	0.56	0.98	5
F4	23	5	17	22	0.93	0.88	0.83	0.99	5

Table 3: Characteristics of the deployments and NLOS/LOS detection performance. (Legend: #: Number of)

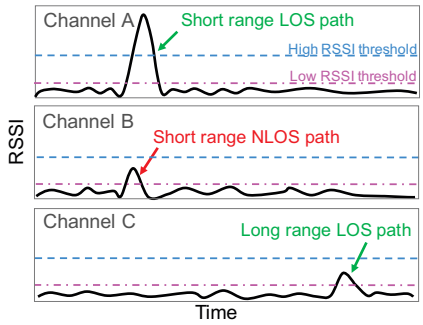

(a) Illustration of channel

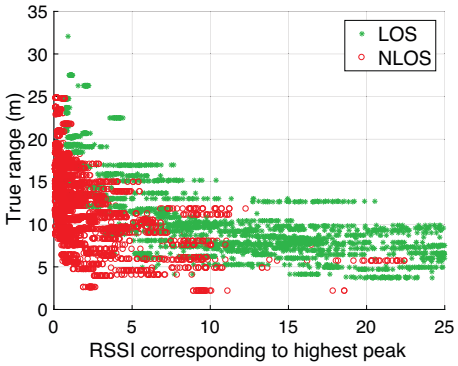

(b) Variation in RSSI with range

Figure 11: NLOS signal characteristics

long range LOS signal in Channel C would not be detected. On the other hand, if a low RSSI threshold is selected, we would detect both signals. Hence there is a trade-off between the amount of NLOS signals detected and the maximum range of LOS signals that are detectable. Figure 11(b) shows this trade-off from real-world data collected using our beacon platform in several environments. If the threshold is set to be high, we can eliminate most of the NLOS but the LOS range is limited. Since our localization system is robust to NLOS signals under the assumption that NLOS signals are positively biased, we set the threshold to a low value of 1 and allow NLOS signals in order to have a long range of 25m for LOS. This further reduces beacon density as it allows us to cover a larger area with fewer beacons as compared to systems with lower LOS range.

5.4 Complexity and Number of Hypotheses

We first make a distinction between the number of hypotheses feasible for a deployment and the number of hypotheses evaluated while solving for location. The number of feasible hypotheses for a floor plan grows with the number of beacons and is highly dependent on the geometry of the floor plan. If we consider all combinations of N beacons, we can get a maximum of 2^N combinations of hypotheses. The number of feasible hypotheses is much lower due to the structured nature of floor plans. As a comparison, the floor plan in Figure 4 has five beacons, and hence 32 possible hypotheses, but the number of feasible hypotheses for the floor plan is six. For building-scale large floor plans, the number of feasible hypotheses can be high but among the feasible hypotheses, only the hypotheses containing one of the beacons in range to be in LOS will be evaluated. The last column of Table 3 shows the number of hypotheses evaluated on an average across all test points for *CA-short* before the algorithm converged. If we notice the difference between this and the total number of hypotheses feasible for this floor plan (fourth column), we see that the *CA-short* approach evaluates fewer hypotheses. The cost of considering a hypothesis is the cost of computing the mean square error across all locations in that hypothesis. The complexity at worst compares with performing a grid search (min mean square error estimation) which in practice works well even on large floor plans.

5.5 Environmental Effects

Our ray-tracing model assumes that the NLOS signals are caused by signals reflecting off walls. Two environmental factors that challenge this assumption are the presence of doors (doors can be open or close but the ray-tracing model is deterministic) and the presence of people (people move around and reflect/obstruct signals). In this section, we discuss the impact of both these factors.

Effect of doors: For the system to work in a region when a door is closed, a sufficient number of beacons should be deployed to provide coverage assuming the scenario of the door being closed irrespective of the localization approach. During ray-tracing, we model the path between a location and a beacon across a doorway as LOS. If the door is closed, the true hypothesis is penalized for having a missing LOS range. However, this hypothesis would still rank higher than the others since sufficient LOS beacons would be in range even without the range from the beacon across the doorway. When the door is open, we receive a LOS range from the beacon which is utilized for localization and the hypothesis assuming this is a LOS beacon will be selected. If a door is generally left open when the environment is in use, we recommend assuming the doorway is always open in order to deploy fewer beacons.

Effect of people: To study the impact of people in the environment on the received signals from the beacons, we perform an experiment in the environment shown in Figure 8 with the help of 10 participants. A person was holding the test device in hand and was stationary. In the first test scenario, the participants were crowded around the test device either standing still or walking, within a 0.5m radius. In the second test scenario, the participants were walking around the test device within a 4m radius. The error in range measurements from four beacons deployed in the test area is shown in Figure 12. We have shown the median error and the 85 percentile error. With people walking around there is little impact because the beacons are deployed at ceiling level and people are unlikely to block the signals unless they are very close to the device. The performance degrades when people are crowded around the phone. Though the ranging error increases due to the diffraction of sound around people, they do not completely block the direct path and the errors are less pronounced compared to wall obstructions.

6 CONCLUSIONS

In this paper, we presented a localization approach for range-based beaconing systems that leverages floor plan information and beacon coverage models to improve accuracy even with low beacon densities. Our key insight is that the combination of a well-defined

coverage-model and floor plan information can be used to reason about the consistency of LOS and NLOS beacon combinations. While our approach can be used with any range-based beacon technology, it performs best with technologies that have a well-behaved coverage model that is predictable with respect to walls in a floor plan. We capitalize upon these properties by designing a sectored ultrasonic localization platform where each beacon can transmit unique signals from each of four quadrants. Our platform is able to synchronize mobile phones using periodic BLE advertisements with enough accuracy to perform TOF ranging. We experimentally evaluate our system in four floor plans and see the 80% accuracy point shift from $4-8m$ to $1m$ as compared to standard gradient descent and grid search approaches. As future work, we intend to focus on the problem of beacon-placement, automatic beacon mapping and explore approaches for heading acquisition in addition to location acquisition.

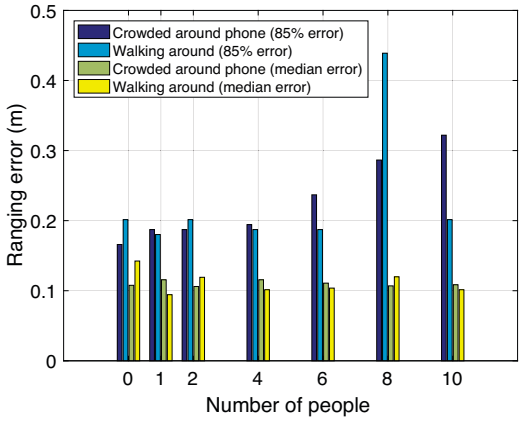

Figure 12: Effect of presence of people on measured range

ACKNOWLEDGMENT

This work was supported in part by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA, and by the NIST Public Safety Communications Research division.

REFERENCES

[1] https://www.microsoft.com/en-us/research/event/microsoft-indoor-localization-competition-ipsn-2017/ (viewed 10/02/2017).

[2] 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.

[3] György Balogh, A Ledeczi, Miklós Maróti, and Gyula Simon. 2005. Time of arrival data fusion for source localization. In *Proceedings of The WICON Workshop on Information Fusion and Dissemination in Wireless Sensor Networks (SensorFusion 2005)*.

[4] Pi-Chun Chen. 1999. A non-line-of-sight error mitigation algorithm in location estimation. In *Wireless Communications and Networking Conference, 1999. WCNC. 1999 IEEE*, Vol. 1. IEEE, 316–320.

[5] Alexander Ens, Fabian Höflinger, Johannes Wendeberg, Joachim Hoppe, Rui Zhang, Amir Bannoura, Leonhard M Reindl, and Christian Schindelhauer. 2015. Acoustic self-calibrating system for indoor smart phone tracking. *International Journal of Navigation and Observation* 2015 (2015).

[6] F Hammer, M Pichler, H Fenzl, A Gebhard, and C Hesch. 2015. An acoustic position estimation prototype system for underground mining safety. *Applied Acoustics* 92 (2015), 61–74.

[7] JC Herrera, A Hinkenjann, PG Ploger, and J Maiero. 2013. Robust indoor localization using optimal fusion filter for sensors and map layout information. In *Indoor Positioning and Indoor Navigation (IPIN), 2013 International Conference on*. IEEE, 1–8.

[8] Sebastian Hilsenbeck, Dmytro Bobkov, Georg Schroth, Robert Huitl, and Eckehard Steinbach. 2014. Graph-based data fusion of pedometer and WiFi measurements for mobile indoor positioning. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 147–158.

[9] Faheem Ijaz, Hee Kwon Yang, Arbab Waheed Ahmad, and Chankil Lee. 2013. Indoor positioning: A review of indoor ultrasonic positioning systems. In *Advanced Communication Technology (ICACT), 2013 15th International Conference on*. IEEE, 1146–1150.

[10] Yung-Hoon Jo, Joon-Yong Lee, Dong-Heon Ha, and Shin-Hoo Kang. 2006. Accuracy enhancement for UWB indoor positioning using ray tracing. In *Position, Location, and Navigation Symposium, 2006 IEEE/ION*. IEEE, 565–568.

[11] Jongdae Jung and Hyun Myung. 2011. Indoor localization using particle filter and map-based NLOS ranging model. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 5185–5190.

[12] Patrick Lazik, Niranjini Rajagopal, Oliver Shih, Bruno Sinopoli, and Anthony Rowe. 2015. ALPS: A Bluetooth and Ultrasound Platform for Mapping and Localization. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems (SenSys '15)*. ACM, New York, NY, USA, 73–84. DOI: https://doi.org/10.1145/2809695.2809727

[13] Patrick Lazik, Niranjini Rajagopal, Bruno Sinopoli, and Anthony Rowe. 2015. Ultrasonic Time Synchronization and Ranging on Smartphones. In *Proceedings of the 21st IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2015) (RTAS '15)*. IEEE, IEEE, 10.

[14] Patrick Lazik and Anthony Rowe. 2012. Indoor Pseudo-ranging of Mobile Devices Using Ultrasonic Chirps. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems (SenSys '12)*. ACM, Toronto, Ontario, Canada, 99–112. DOI: https://doi.org/10.1145/2426656.2426667

[15] Xinrong Li. 2006. An iterative NLOS mitigation algorithm for location estimation in sensor networks. *Proceedings of the15 th IST Mobile and Wireless Communications Summit* (2006), 1–5.

[16] Kaikai Liu, Xinxin Liu, and Xiaolin Li. 2013. Guoguo: Enabling fine-grained indoor localization via smartphone. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*. ACM, 235–248.

[17] Kaikai Liu, Xinxin Liu, Lulu Xie, and Xiaolin Li. 2013. Towards accurate acoustic localization on a smartphone. In *INFOCOM, 2013 Proceedings IEEE*. IEEE, 495–499.

[18] Dimitrios Lymberopoulos, Jie Liu, Xue Yang, Romit Roy Choudhury, Vlado Handziski, and Souvik Sen. 2015. A realistic evaluation and comparison of indoor location technologies: Experiences and lessons learned. In *Proceedings of the 14th international conference on information processing in sensor networks*. ACM, 178–189.

[19] Georg Oberholzer, Philipp Sommer, and Roger Wattenhofer. 2011. Spiderbat: Augmenting wireless sensor networks with distance and angle information. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*. IEEE, 211–222.

[20] Chunyi Peng, Guobin Shen, Yongguang Zhang, Yanlin Li, and Kun Tan. 2007. Beepbeep: a high accuracy acoustic ranging system using cots mobile devices. In *Proceedings of the 5th international conference on Embedded networked sensor systems*. ACM, 1–14.

[21] Nissanka Bodhi Priyantha. 2005. *The cricket indoor location system*. Ph.D. Dissertation. Massachusetts Institute of Technology.

[22] Niranjini Rajagopal, Sindhura Chayapathy, Bruno Sinopoli, and Anthony Rowe. 2016. Beacon placement for range-based indoor localization. In *Indoor Positioning and Indoor Navigation (IPIN), 2016 International Conference on*. IEEE, 1–8.

[23] BJ Silva and Gerhard P Hancke. 2015. Practical challenges of IR-UWB based ranging in harsh industrial environments. In *Industrial Informatics (INDIN), 2015 IEEE 13th International Conference on*. IEEE, 618–623.

[24] Bruno Silva and Gerhard P Hancke. 2016. IR-UWB-Based Non-Line-of-Sight Identification in Harsh Environments: Principles and Challenges. *IEEE Transactions on Industrial Informatics* 12, 3 (2016), 1188–1195.

[25] Abdelhamid Tayebi, Josefa Gomez Perez, Francisco Manuel Saez de Adana Herrero, and Oscar Gutierrez Blanco. 2009. The application of ray-tracing to mobile localization using the direction of arrival and received signal strength in multipath indoor environments. *Progress In Electromagnetics Research* 91 (2009), 1–15.

[26] Yu-Ting Wang, Rong Zheng, and Dongmei Zhao. 2016. Towards Zero-Configuration Indoor Localization Using Asynchronous Acoustic Beacons. *Embedded and Ubiquitous Computing (EUC)* (2016).

[27] Zhuoling Xiao, Hongkai Wen, Andrew Markham, and Niki Trigoni. 2014. Lightweight map matching for indoor localisation using conditional random fields. In *Information Processing in Sensor Networks, IPSN-14 Proceedings of the 13th International Symposium on*. IEEE, 131–142.