

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

---

7-2018

### Experiences & challenges with server-side WiFi indoor localization using existing infrastructure

Dheryta JAISINGHANI

Singapore Management University, dherytaj@smu.edu.sg

Rajesh Krishna BALAN

Singapore Management University, rajesh@smu.edu.sg

Vinayak NAIK

Birla Institute of Technology and Science (BITS)

Archan MISRA

Singapore Management University, archanm@smu.edu.sg

Youngki LEE

Singapore Management University, YOUNGKILEE@smu.edu.sg

Follow this and additional works at: [https://ink.library.smu.edu.sg/sis\\_research](https://ink.library.smu.edu.sg/sis_research)



Part of the [OS and Networks Commons](#), and the [Software Engineering Commons](#)

---

#### Citation

JAISINGHANI, Dheryta; BALAN, Rajesh Krishna; NAIK, Vinayak; MISRA, Archan; and LEE, Youngki. Experiences & challenges with server-side WiFi indoor localization using existing infrastructure. (2018). *MobiQuitous '18: Proceedings of the 15th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services, New York, November 5-7*. 1-10. Available at: [https://ink.library.smu.edu.sg/sis\\_research/4249](https://ink.library.smu.edu.sg/sis_research/4249)

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylids@smu.edu.sg](mailto:cherylids@smu.edu.sg).

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/328964184>

# Experiences & Challenges with Server-Side WiFi Indoor Localization Using Existing Infrastructure

Conference Paper · November 2018

DOI: 10.1145/3286978.3286989

CITATIONS

0

READS

70

5 authors, including:



**Dheryta Jaisinghani**

Indraprastha Institute of Information Technology

9 PUBLICATIONS 6 CITATIONS

SEE PROFILE



**Vinayak Naik**

BITS Pilani, K K Birla Goa

63 PUBLICATIONS 1,471 CITATIONS

SEE PROFILE



**Archan Misra**

Singapore Management University

313 PUBLICATIONS 6,608 CITATIONS

SEE PROFILE



**Youngki Lee**

Korea Advanced Institute of Science and Technology

56 PUBLICATIONS 767 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Ph.D. Thesis [View project](#)



WHO project on the development of Vaccine against Kala Azar [View project](#)

# Experiences & Challenges with Server-Side WiFi Indoor Localization Using Existing Infrastructure

Dheryta Jaisinghani  
Indraprastha Institute of Information  
Technology - Delhi  
India  
dherytaj@iiitd.ac.in

Rajesh Krishna Balan  
School of Information Systems,  
Singapore Management University  
Singapore  
rajesh@smu.edu.sg

Vinayak Naik  
Birla Institute of Technology and  
Science-Pilani, Goa  
India  
vinayak@goa.bits-pilani.ac.in

Archan Misra  
School of Information Systems,  
Singapore Management University  
Singapore  
archanm@smu.edu.sg

Youngki Lee  
Seoul National University,  
Seoul  
Korea  
youngkilee@snu.ac.kr

## ABSTRACT

Real-world deployments of WiFi-based indoor localization in large public venues are few and far between as most state-of-the-art solutions require either client or infrastructure-side changes. Hence, even though high location accuracy is possible with these solutions, they are not practical due to cost and/or client adoption reasons. Majority of the public venues use commercial controller-managed WLAN solutions, that neither allow client changes nor infrastructure changes. In fact, for such venues we have observed highly heterogeneous devices with very low adoption rates for client-side apps.

In this paper, we present our experiences in deploying a scalable location system for such venues. We show that server-side localization is not trivial and present two unique challenges associated with this approach, namely *Cardinality Mismatch* and *High Client Scan Latency*. The “Mismatch” challenge results in a significant mismatch between the set of access points (APs) reporting a client in the offline and online phases, while the “Latency” challenge results in a low number of APs reporting data for any particular client. We collect three weeks of detailed ground truth data ( $\approx 200$  landmarks), from a WiFi setup that has been deployed for more than four years, to provide evidences for the extent and understanding the impact of these problems. We propose heuristics to alleviate them. We also summarize the challenges and pitfalls of real deployments which hamper the localization accuracy.

## CCS CONCEPTS

• **Networks** → **Network measurement; Location based services;** • **Human-centered computing** → **Empirical studies in ubiquitous and mobile computing;**

## KEYWORDS

WiFi, Localization, Server-side, Device-agnostic, Large-scale Measurements.

## 1 INTRODUCTION

There has been a long and rich history of WiFi-based indoor localization research [1–3, 6, 9, 12, 13, 15, 18–20, 22, 23, 27–30, 32, 33, 35, 36, 39, 42–47, 49, 50]. However, in spite of several breakthroughs, there are very few real-world deployments of WiFi-based indoor localization systems in public spaces. The reasons for this are many-fold, with three of the most common being – (a) the high cost of deployment, (b) arguably, the lack of compelling business use, and (c) the inability of existing solutions to seamlessly work with all devices. In fact, current solutions impose a tradeoff between universality, accuracy, and energy, for example, client-based solutions that combine inertial-based tracking with WiFi scanning offer significantly better accuracy but require a mobile application which will possibly drain energy faster and which will be downloaded by only a fraction of visitors [53].

In this paper, we present our experiences with deploying and operating a WiFi-based indoor localization system across the entire campus of a small Asian university. It is worth noting that the environment is very densely occupied, by  $\approx 10,000$  students and 1,500 faculty and staff. The system has been in the production for more than four years. It is deployed at multiple venues including two universities (Singapore Management University, University of Massachusetts, Amherst), and four different public spaces (Mall, Convention Center, Airport, and Sentosa Resort) [21, 40]. These venues use the localization system for various real-time analytics such as group detection, occupancy detection, and queue detection while taking care of user privacy.

Our goal is to highlight challenges and propose easy to integrate solutions to build a universal indoor localization system – one that can spot localize all WiFi enabled devices on campus without any modifications whether client or infrastructure-side. The scale and the nature of this real environment, presents unique set of challenges – (a) infrastructure *i.e.* controller and APs do not allow any changes, (b) devices cannot be modified in any way *i.e.* no explicit/implicit participation for data generation, no app download allowed, and no chipset changes allowed, and (c) only available data is RSSI measurements from APs, which are centrally controlled by the controller, using a *Real-Time Location Service* (RTLS) interface [8]. It is worth noting that within the face of these challenges

we have to rule out more sophisticated state-of-the-art schemes, such as fine-grained CSI measurements [28], Angle-of-Arrival [9], Time-of-Flight [29], SignalSLAM [37], or Inertial Sensing [14].

Given the challenges, we adopt an offline fingerprint-based approach to compute each device’s location. Fingerprints have been demonstrated to be more accurate than model-based approaches in densely crowded spaces [26] and hence widely preferred. Our localization software processes the RSSI updates using well-known “classical” fingerprint-based technique [36]. Given the wide usage of this approach, our experiences and results apply to a majority of the localization algorithms.

Our primary contribution is to detail the cases where such a conventional approach succeeds and where it fails. We highlight the related challenges for making the approach work in current, large-scale WiFi networks, and then develop appropriate solutions to overcome the observed challenges. We collect three weeks of detailed ground truth data ( $\approx 200$  landmarks) in our large-scale deployment, carefully construct a set of experimental studies to show two unique challenges – *Cardinality Mismatch* and *High Client Scan Latency* associated with a server-side localization approach. The three weeks of data is representative of our four years of data.

(a) *Cardinality Mismatch*: We define cardinality as the set of APs reporting for a client located at a specific landmark. We first show that the cardinality, during the online phase, is *often* quite different from the cardinality in the offline phase. Note that this divergence is in the *set* of reporting APs, and not just merely a mismatch in the values of the RSSI vectors. Intuitively, this upends the very premise of fingerprint-based systems that the cardinality seen at any landmark is the same during the offline and online phases. This phenomenon arises from the dynamic power and client management performed by a centralized controller in all commercial grade WiFi networks (for example, those provided by Aruba, Cisco, and other vendors) to achieve outcomes such as (i) minimize overall interference (shift neighboring APs to alternative channels), (ii) enhance throughput (shift clients to alternative APs), and (iii) reduce energy consumption (shut down redundant APs during periods of low load).

(b) *High Client Scan Latency*: Most localization systems use client-side localization techniques where clients actively scan the network when they need a location fix. However, when using server-side localization, the location system has no way to induce scans from client devices. Hence, the system can only “see” clients when clients scan as part of their normal behavior. However, as we show in Section 4, scanning frequency of clients is low for lower RSSI.

These phenomena do not exist in small-scale deployments often used in the past pilot studies, where each AP is configured independently. In large-scale deployments, where it is fairly common to use controller-managed WLANs with a large number of devices, these phenomena invariably persist to a great extent. To exemplify, we noticed 57.30% instances of cardinality mismatch in 2.4 GHz and 30.60% in 5 GHz in our deployment. We saw 90<sup>th</sup> percentile of client scan interval to be 20 minutes. While localizing with fingerprint-based solutions in such environments, these phenomena translate to either *minimal* or even worse *no* matching APs, resulting in substantial delays between client location updates and “teleporting” of clients across the location.

It is important to note that not only the schedule of these algorithms is non-deterministic but also their distribution during offline and online phases. This is attributed to the fundamental fact that the dynamics of WiFi networks such as load and interference, is non-deterministic in most of the cases and that the controller algorithm is a black-box to us. Furthermore, the differences in signal propagation and scanning behavior of 2.4 and 5 GHz contribute to these problems. We believe that we are the first to present the challenges of server-side localization as well as their mitigation. Our proposals are device-agnostic, simple, and easily integrable with any large-scale WiFi deployment to efficiently localize devices.

### Key Contributions:

- We identify and describe a couple of novel and fundamental problems associated with a server-side localization framework. In particular, we provide evidence for the (i) “Cardinality Mismatch” and (ii) “High Client Scan Latency” problems, explain why these problems are progressively becoming more significant in commercial WiFi deployments. We discuss the reasons why these problems are non-trivial to be solved given the challenge of no client/infrastructure-side allowed. Our entire analysis is for both frequency bands – 2.4 and 5 GHz.
- We provide valuable insights about the causes of these problems with extensive evaluations based on the ground truth data collected over three weeks for 200 landmarks. We propose heuristics to improve the accuracy of the localization in the face of these problems. We see an improvement from a minimum of 35.40% to a maximum of 100%. We show an improvement in the higher percentiles over SignalSLAM [37]. This shows that our lessons learned have the potential of improving the existing localization algorithms.
- We describe our experiences with deploying, managing, and improving a fingerprint-based WiFi localization system, which has been operational, since 2013, across the entire campus of Singapore Management University. We not only focus on the final “best solution” that uses RTLS data feeds, but also discuss the challenges and pitfalls encountered over the years.

**Paper Organization:** We discuss the related works in Section 2. We present the system architecture and the details of data collection in Section 3. We introduce the challenges, their evidences, and propose the solutions in Section 4. We discuss the challenges of localizing clients in real world deployments and the limitations of our proposed solutions in Section 5. We conclude in Section 6.

## 2 RELATED WORK

In this section, we discuss existing solutions and their limitations for indoor localization.

**Fingerprint vs. Model-based Solutions:** One of the oldest localization techniques use either a fingerprint-based [2, 19, 27, 30, 33, 36, 42, 55] or model-based [1, 10, 16, 22] approach, or a combination of both [24]. Overall, fingerprint-based solutions tend to have much higher accuracies than other approaches albeit with a high setup and maintenance cost [26]. The fingerprint-based approach was pioneered by Radar [36] and has spurred numerous follow-on research. For example, Horus [33] uses a probabilistic


technique to construct statistical radio maps, which can infer locations with centimeter level accuracy. PinLoc [42] incorporates physical layer information into location fingerprints. Liu *et al.* [27] improved accuracy by adopting a peer-assisted approach based on p2p acoustic ranging estimates. Another thread of research in this line is to reduce the fingerprinting effort, for example, using crowd-sourcing [2, 6, 52, 55] and down-sampling [19]. The mathematical signal propagation model approach [1, 22] has the benefit of easy deployment (no need for fingerprints) although its accuracy suffers when the environment layout or crowd dynamics change [10]. Systems, such as EZ, improve the accuracy by additionally using GPS to guide the model construction.

**Client vs. Infrastructure-based solutions:** There is a rich history of client-based indoor location solutions, to name a few, SignalSLAM [37], SurroundSense [30], UnLoc [15], and many others [2, 3, 6, 35, 39, 44]. All of them share some commonalities in that they extract sensor signals (of various types) from client devices to localize. The location algorithms usually run on the device itself; however, it is also possible to run the algorithm on a server and use the signals from multiple clients to achieve better performance [27]. Overall, client-based solutions have very high accuracy (centimeter resolution in some cases [33, 42]). An alternative would be to pull signal measurements directly from the WiFi infrastructure, similar to what our solution does. The research community has only lightly explored this approach since it requires full access to the WLAN controllers, which is usually proprietary. Our main competitors are the commercial WiFi providers themselves. In particular, both Cisco [7] and Aruba [4] offer location services. These solutions use server-side tracking coupled with model-based approaches (to eliminate fingerprint setup overhead).

**Other Solutions:** There are several other solutions, complementary to the signal strength-based technique. Time-based solutions [12, 29, 32, 43, 45] use the arrival time of signals to estimate the distance between client and AP, while angle-based solutions [9, 18, 20, 28, 46] utilize angle of arrival information, estimated from an antenna array, to locate mobile users. Recently, the notion of passive location tracking [13, 23, 47, 49, 50] has been proposed, which does not assume people carry devices. In large and crowded venues, however, the feasibility and accuracy of such passive tracking is still an open question. Other systems like light-based localization [25, 34, 38] and acoustic-based localization [31, 41, 48, 54].

**Limitations of above solutions:** These solutions can achieve higher accuracy, but they have at least one of the following limitations – (a) need of a customized hardware, which cannot be implemented in large-scale deployments, (b) installation of client application, making them hard to scale, (c) rooting client OS - Android or iOS, which limits their generalizability, (d) energy savvy, (e) high error rates in dense networks, and (f) proprietary and expensive to deploy (especially, solutions from vendors like Cisco and Aruba).

To summarize, even though several wonderful solutions are available, their scalability is still a question. Therefore, we advocate using server-side localization approach with fingerprints. Our aim is not to compare the efficacy of different approaches, but to address the challenges of practical and widely deployed device-agnostic indoor localization using today’s WiFi standards and hardware, for example, use of 5 GHz band and controller-based architecture.



**Figure 1: Block diagram of Indoor Localization System. Note that lines between AP and C denote the coverage area of AP and not the association. Legend: AP - Access Point, C - Client, WLAN - Wireless LAN, RTLS - Real-Time Location System**

**Table 1: Details of RTLS data feeds**

Field	Description
Timestamp	AP Epoch time (milliseconds)
Client MAC	SHA1 of original MAC address
Age	#Seconds since the client was last seen at an AP
Channel AP	Band (2.4/5 GHz) on which client was seen MAC address of the Access Point
Association Status	Client’s association status (associated/unassociated)
Data Rate	MAC layer bit-rate of last transmission by the client
RSSI	Average RSSI for duration when client was seen

### 3 SYSTEM ARCHITECTURE AND DATA COLLECTION

In this section, we present the details about system architecture and the dataset.

#### 3.1 Background & Deployment

This work began in 2013 when we started deploying a WiFi-based localization solution across the entire campus. It has since gone through much major and minor evolutions. However, in this paper, we focus our evaluation and results on just one venue – a university, as we have full access to that venue.

Our university campus has seven schools in different buildings. Five buildings have six floors, remaining two have five and three floors respectively, with a floor area of  $\approx 70,000 m^2$ . Landmarks, characterized by water sprinklers are deployed every three meters, on a given floor denote a particular location. There are 3203 landmarks across thirty-eight floors of seven schools. WLAN deployment includes 750+ dual-band APs, centrally controlled by eleven WiFi controllers, with  $\approx 4000$  associated clients per day.

#### 3.2 System Architecture

Figure 1 represents the primary building blocks of the system. The system is bootstrapped with APs configured by the WLAN controller to send RTLS data feeds every 5 seconds to the RTLS server. Most commercial WLAN infrastructures allow such a configuration. Once configured, APs bypass WLAN controller and report RTLS

data feeds directly to our Location Server. Table 1 presents all the fields contained in an RTLS data feed per client. The reported RSSI value is not on a per-frame basis, but a summarized value from multiple received frames. The Location Server analyzes these RTLS data feeds for the signal strengths reported by different APs to estimate the location of a client. Note that the APs do not report the type of frames. They gather information from their current channel of operation and scan other channels to collect data. Vendors have microscopic details of what APs measure [5], however as an end-user we do not have access to any more information than what is specified. Nevertheless, even this information at large-scale gives us a view of the entire network from a single vantage point.

### 3.3 Recording of the Fingerprints

We define a fingerprint as a vector of RSSI from APs for a given client. We consider two types of fingerprints – offline and online. An offline fingerprint is collected and stored in a database before the process of localization is bootstrapped, while an online fingerprint is collected in real-time.

**Offline Fingerprinting** A two-dimensional offline fingerprint map is prepared for each landmark on the per-floor basis. The client devices used for fingerprinting were dual-band Android phones, which were associated with the network, and they actively scanned for APs. For each landmark, the device collected data for 5 minutes. While the clients scan their vicinity, APs collate RSSI reports for the client and send their measurements as RTLS data feeds to the Location Server. For a given landmark  $L_i$ , an offline fingerprint takes the following form:

$$\langle L_i, B, AP_1 : RSSI_1; \dots; AP_n : RSSI_n; \rangle \quad (1)$$

We maintain fingerprints for both 2.4 and 5 GHz frequency bands. Band  $B$ , in the above equation, takes a value of band being recorded. The vectors are stored in a database on the Location Server.

**Online Fingerprinting** Localization of a client is done with online fingerprints. An online fingerprint takes the same syntax as offline fingerprints in Equation 1, except the landmark, as shown below:

$$\langle B, AP_1 : RSSI_1; \dots; AP_m : RSSI_m; \rangle \quad (2)$$

Now, we match this online fingerprint with offline fingerprints of each landmark to calculate the distance in signal space, as discussed in [36]. The landmark with minimum distance in signal space is reported as the probable location of the client.

### 3.4 Pre-processing of the Data

Now, we present the details of data collection and its processing.

**3.4.1 Collection of the Ground Truth.** We collect the ground truth data for online fingerprints. We want to correlate the data collection with real-world usage scenarios. Therefore, we choose four most common states of WiFi devices as per their WiFi association status and Data transmission. The states are – (i) Disconnected, (ii) WiFi Associated – (ii.a) Never actively used by user, (ii.b) Intermittently used, and (ii.c) Actively used. These states implicitly modulate the scanning frequency. We use a separate phone for each state; thus, we use 4 Samsung Galaxy S7 phones to record ground truth for each landmark.

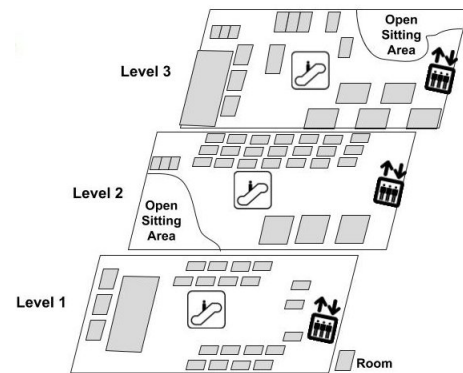



Figure 2: Floor Map of the school where we collected ground truth data.

State (a) client is disconnected. In this state, WiFi is turned on but not associated with any AP and screen remains off throughout the data collection. Therefore, only traffic generated from this client is scanning traffic and no data traffic. We ensure that this client did not follow MAC address randomization, which most latest devices follow in the unassociated state [17]. State (b) client is associated but inactive. In this state, WiFi is turned on, it is associated but the screen remains off throughout the data collection. State (c) client is associated and intermittently used. In this state, WiFi is on, the client is associated and user intermittently uses the device. This is one of the most common state for mobile devices and previous research [11] states scanning is triggered whenever screen of the device is lit up. State (d) client is associated and actively used. In this state, WiFi is on, the client is associated, and a YouTube video plays throughout the data collection. This state generates most data traffic, *i.e.* non-scanning traffic. Each client stayed at a landmark for about a minute before it moved to the next landmark. We manually noted down start time and end time for every landmark at the granularity of seconds. We did this exercise for 3203 landmarks of our university, collected 86 hours worth of data, that accounts for 54,096 files carrying 274 GB of data. The amount of time to localize a client is 40 seconds. Processing the entire dataset would take  $\approx 100$  days. Therefore, given the size of the entire dataset, we present our analysis of 200 landmarks, which accounts for 3121 files with 15.3 GB of data. Figure 2 shows the floor map of one of the schools whose data we refer for our analysis.

Our aim is to demonstrate the challenges associated with fingerprint-based localization. These challenges apply to all the solutions that employ fingerprint-based localization, irrespective of the type of device present in the network. The variation of RSSI with device heterogeneity is well known [51] and that will further exacerbate the problems identified by this paper. We collect ground truth with only one device so that we can highlight issues without any complications added by heterogeneous devices.

**3.4.2 Pre-processing of the RTLS Data Feeds.** Our code reads every feed to extract the details of APs reporting a particular client. RTLS data feeds, may obtain stale records for a client. Therefore, we filter the raw RTLS data feed for the latest values, with age less than or equal to 15 seconds, and the RSSI should be greater than or equal to  $-72$  dBm. The threshold for age is a heuristic to take the



**Figure 3: Cardinalities observed during the offline and online phases. For both the phases, the cardinalities are lower for 5 GHz. During the online phase, there is a substantial decrease in the cardinality for both bands as compared to the offline phase.**

most recent readings. The threshold for RSSI is decided based on the fact that a client loses association when RSSI is below  $-72$  dBm.

For our analysis, we classify MAC layer frames in two classes (a) *Scanning Frames* – high power and low bit rate probe requests and (b) *Non-Scanning Frames* – all other MAC layer frames. Offline fingerprints are derived from the scanning frames, which are known to provide accurate distance estimates as they are transmitted at full power. In the offline phase, a client is configured to scan continuously. However, in the online phase we have no control over the scanning behavior of the client, resulting in a mix of scanning and non-scanning frames. Therefore, while localizing with the fingerprints, RSSIs available for matching are from the different categories of frames. RTLS data feeds do not report the type of frame and do not have a one-to-one mapping of MAC layer frames to the feeds. Therefore, we devise a probabilistic approach to identify these frames.

We design with a set of controlled experiments, where we configured the client in one of the two settings at a time (a) send scanning frames only and (b) send non-scanning frames only. These two settings are mutually exclusive. We collected the traffic from the client with a sniffer as well as the corresponding RTLS data feeds. Then, we compare both the logs – sniffer and RTLS, to confirm the frame types and the corresponding data rates.

Our analysis reveals that when a client is associated and sending non-scanning frames, the AP to which it is associated reports the client as *associated*. The data rates of the RTLS data feeds vary among various 802.11g rates, e.g. 1, 2, 5.5, ..., 54 Mbps. Even though, our network deployment is dual-band and supports the latest 802.11 standards including 802.11ac, still the rates reported in the RTLS data feeds follow 802.11g. We do not have any visibility in the controller’s algorithm to deduce the reason for this mismatch in the reported data rates. However, when a client sends scanning frames, all the APs that could see the client report the client as *unassociated* and the data rates reported is fixed at either 1, 6, or 24 Mbps, as per the configured probe response rate.

We use these facts to differentiate non-scanning and scanning RTLS data feeds. We believe this approach correctly infers scanning frames because (a) the data rates are fixed to 1, 6, or 24 Mbps, (b) when an associated client scans, other APs report that client as

unassociated, and (c) an unassociated client can only send either scanning or association frames. However, our approach may still incorrectly identify a scanning frame as non-scanning in the following cases – (a) When an associated client scans and the AP, to which it is associated, reports. This AP reports the client as associated and its data rate as 1, 6, or 24 Mbps. In this case, these rates may also be because of the non-scanning frames. We identify such feeds as non-scanning. (b) When an unassociated client sends association or authentication frames. In this case also, the rates overlap with the scanning data rates and the association status is reported as unassociated. Here, we incorrectly identify non-scanning frames as scanning frames. However, these cases are rare. For other cases, our approach is deterministically correct.


## 4 CHALLENGES DISCOVERED

In this section, we give evidence of the issues, namely Cardinality Mismatch and High Client Scan Latencies. We compare the severity of these issues for both the frequency bands. We identify the causes behind these issues and measure their impact on the issues.


### 4.1 Evidence of the Issues

The Cardinality Mismatch arises from the dynamic power and client management performed by a centralized controller as well as the client-side power management. Given the dynamic nature of these management policies, it is not possible to estimate their implications on the Cardinality Mismatch, and thereby on the localization errors. We take an empirical approach to see whether (a) we can find out the severity of these implications on the Cardinality Mismatch and the localization error and (b) identify the tunability of the implicating factors.

Figure 3 plots the differences in cardinality between the offline and online phases for 2.4 and 5 GHz. Figure 3a shows the cardinality observed in our offline fingerprints. Figure 3b shows the cardinality observed during the online phase. While the maximum cardinality is 16 during the offline phase, it is merely 6 in the online phase. This shows the spectrum of the Cardinality Mismatch. In the online phase, 80% of the time only 1 AP reports for a client in 5 GHz while 40% in 2.4 GHz. Any fingerprint-based algorithm will be adversely affected by such a big difference in the cardinality. For



**Figure 4: Variations in RSSI for scanning and non-scanning frames in two scenarios – (a) client close to the AP and (b) client far from the AP. For both cases the RSSI from scanning frames vary far lesser than the non-scanning frames.**



**Figure 5: Frequency of scanning in both band increases as RSSI reduces. 2.4 GHz experiences higher scan frequencies.**

each band, we find out how much is the extent of the Cardinality Mismatch. Overall, across all the cardinalities, 2.4 GHz has 57.30% mismatches and 5 GHz has 30.6% mismatches. The 5 GHz band is more adversely affected by the Cardinality Mismatch issue because it experiences lower cardinality, which increases the chances of a mismatch. Overall, 2.4 GHz always sees higher cardinality than 5 GHz, both during the offline and online phases. This is because signals in 2.4 GHz travel farther than that of 5 GHz. However, it is not the only reason. The other reason is the number of scanning frames transmitted. Unlike the data frames, the scanning frames are broadcasted and hence heard by more number of APs. As the number of scanning frames increases, more APs hear them and revert, thereby increasing the cardinality.

Besides, the RSSI variation for the scanning frames is lesser compared to that of the data frames. To validate, we perform a controlled experiment with a stationary client and collect client’s traffic using a sniffer. The client has an ongoing data transmission and periodic scanning is triggered every 15 seconds. From the sniffed packet capture, we extract per-frame RSSI. The experiment is repeated for two scenarios- (a) the client is close to the AP and (b) the client is far from the AP. With these two scenarios, we simulate the client behavior for low and high RSSI from the AP.

Figure 4 shows the RSSI measurements in two cases. In the first scenario, when the client is close to an AP, RSSI of the scanning frames varies by up to 10 dB and for non-scanning frames it varies

by up to 50 dB. Similarly, in the second scenario, when the client is far from AP, RSSI of scanning frames varies by up to 5 dB and for non-scanning frames it varies by up to 30 dB. Both our experiments validate that the RSSI from scanning frames vary far lesser than the non-scanning frames. This means the online RSSI from scanning frames match more closely and is a much more reliable indicator of the client’s position. We want to study how clients in their default configuration behave in *real* networks; therefore we do not modify the default behavior of client driver in any way. We repeated the experiment with devices of Samsung, Nexus, Xiaomi, and iPhone.

Next, we study the effect of the band on the frequency of scanning. We collect WiFi traffic with sniffers listening on the channels in operation at that time in both the bands for 6 hours. Data from 200 WiFi clients is recorded. Figure 5 shows the plot. For both 2.4 and 5 GHz bands, the frequency increases as RSSI reduces. Overall, the frequency is lesser for 5 GHz, even though most ( $\approx 2X$ ) of the clients in our network associate in 5 GHz. More the frequency of scanning, lesser is the chance of Cardinality Mismatch. Our comparative analysis of the two bands revealed that the instance of frame losses and poor connection quality, which cause scanning, are much lower in 5 GHz due to lower interference. The analysis of the scanning behavior of our clients reveals that – (a) 90<sup>th</sup> %ile values of scanning intervals is in the order of few 1000 seconds, which is a lot for fingerprint based solutions, (b) 5 GHz is the least preferred band of scanning, and (c) clients rarely scan both the frequency bands. Hence, we rule out the possibility of the reduced range of 5 GHz resulting in lesser scanning frames.

## 4.2 Causes Behind the Issues

Next, we study the combined effect of frequency of scanning, *i.e.*, number of scans per hour, and transmission distance on the cardinality. For this, we consider clients configured in one of the four states as discussed in Section 3.4.1. Note that each state implicitly controls the amount of scanning. We do not manually control scanning behavior to imitate the real-world. In the absence of client scans, APs get only non-scanning frames. For each of the four states of the client, we study how many APs report that client *i.e.* the cardinality. With this analysis, we are able to compare the cardinality in the presence and absence of scans, for both 2.4 and 5 GHz bands.



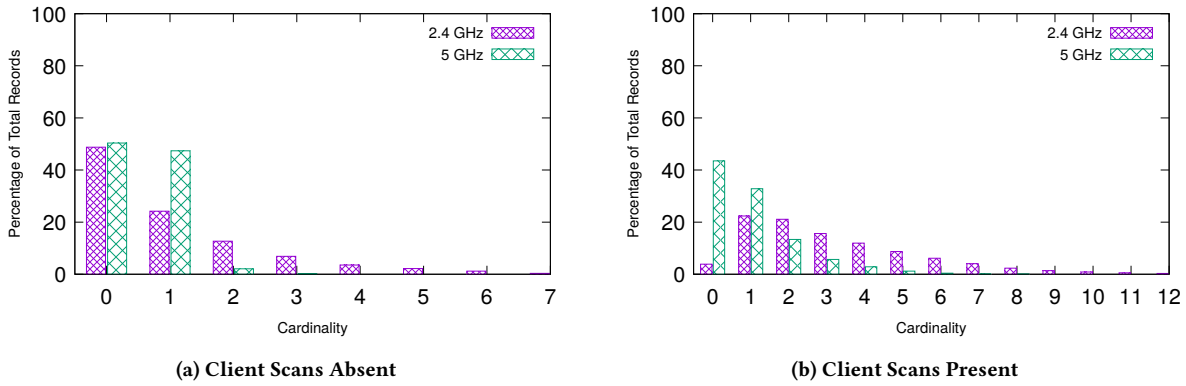


Figure 6: Cardinality in the absence and presence of client scans. More APs report during scanning.




Figure 7: Variations in RSSI in 2.4 GHz and 5 GHz for a stationary client as measured at the server. Notice that 5 GHz is relatively stable than 2.4 GHz.

We see that as the frequency of scanning increases, more number of APs respond and the cardinality increases. Due to lack of space, we don't show the results for individual states of the clients. Figure 6 shows the aggregated results. The cardinality is consistently higher for 2.4 GHz than that of 5 GHz. This implies that higher frequency of scanning possibly reduces the Cardinality Mismatch and vice-versa.

However, a downside for 2.4 GHz is that the frames (both scanning and non-scanning) have higher variation in RSSI. This means, even though the extent of the Cardinality Mismatch is lower, the RSSI will differ more in 2.4 GHz. To confirm this, we analyze the RSSI from a stationary client by enabling its association in one band at a time and disabling the other band altogether. We use the RSSI recorded at the RTLS server for a duration of 1 hour. As shown in Figure 7, even with more scanning information available 2.4 GHz is more prone to RSSI fluctuations than 5 GHz. The reasons for this behavior are (a) the range of 2.4 GHz is almost double than that of 5 GHz and (b) a lesser number of non-overlapping channels makes it susceptible to interference. Therefore, RSSI from 2.4 GHz results in predicting distant and transient locations. We validate this in different locations with devices of four other models.

To summarize, there is a significant extent of Cardinality Mismatch and High Client Scan Latency. There is a difference in the extent of the issues for the two classes of frames and the two bands

Table 2: A summary of the causes and their impact (✓- Reduces Localization Errors, ✗- Increases Localization Errors). Causes conflict with each other, making server-side localization non-trivial.

Frames	Transmission Distance	RSSI Variation	Frequency of Transmission
Scanning	High - ✓	Low - ✓	Low - ✗
Non-Scanning	Low - ✗	High - ✗	High - ✓

Band	Transmission Distance	RSSI Variation	Frequency of Scanning
2.4 GHz	High - ✓	High - ✗	High - ✓
5 GHz	Low - ✗	Low - ✓	Low - ✗

of operation. While scanning frames has a longer distance of transmission and less variation in RSSI, they are not often sent by the clients. The factors favoring 2.4 GHz are longer distance of transmission and higher frequency of scanning. However, low variation in RSSI works in favor of 5 GHz. We summarize these observations in Table 2.


### 4.3 Impact of Causes on Localization Errors

We now evaluate the impact of the causes on the localization errors. We implemented a server-side localization using well known fingerprint-based method [36]. Since we use server-side processing, we do not require any client-side modification. Our proposals do not make assumptions about hardware or OS of the clients or the controller. Although each adaptation of fingerprint-based technique from the existing body of work may result in different errors, our exercise gives us a baseline that cuts across all the adaptations. The 2.4 and 5 GHz bands differ in distance of transmission, variation in RSSI, and frequency of scanning. We measure the localization errors for both the bands.

We report localization errors for each value of the cardinality in online phase to understand how the error varies as a function of the cardinality. We measure the errors in terms of (a) Different Floor and (b) Same Floor errors. Different Floor error is the percentage

**Table 3: Localization errors with different floor detection heuristics ( $C$  - Cardinality, SF - Same Floor [ $85^{th}$ ile (meters)], DF - Different Floor [%], NA - Not Applicable). Cardinalities  $>3$  are not applicable to 5 GHz due to cardinality mismatch. Lowest localization errors obtained using AP of Association.**

$C$	Baseline				Maximum Number of APs				AP with Maximum RSSI				AP of Association			
	2.4 GHz		5 GHz		2.4 GHz		5 GHz		2.4 GHz		5 GHz		2.4 GHz		5 GHz	
	SF	DF	SF	DF	SF	DF	SF	DF	SF	DF	SF	DF	SF	DF	SF	DF
1	>50	20.00	29.50	03.10	>50	48.00	29.60	03.49	49.20	14.70	28.30	02.09	32.30	04.90	27.60	01.28
2	>50	18.30	26.80	02.00	>50	33.33	27.00	05.60	24.20	05.60	27.00	03.46	21.60	00.70	25.80	01.48
3	>50	33.15	20.12	00.00	>50	40.00	24.00	00.00	36.00	14.70	20.12	00.00	22.80	05.50	20.12	00.00
4	>50	18.49	NA	NA	33.00	13.40	NA	NA	24.00	11.00	NA	NA	18.90	04.40	NA	NA
5	19.20	10.25	NA	NA	26.00	14.60	NA	NA	17.49	00.00	NA	NA	17.49	00.00	NA	NA
6	23.00	04.17	NA	NA	22.00	00.00	NA	NA	16.00	00.00	NA	NA	22.00	00.00	NA	NA



**Figure 8: Localization errors with three floor detection heuristics – (a) Maximum Number of APs, (b) AP with Maximum RSSI, and (c) AP of Association at Cardinality=1 for 2.4 GHz. AP of Association performs the best for both bands. Maximum Number of APs performs worse than the Baseline, while the other two significantly reduced the errors.**

of total records for which a wrong floor is estimated. These are seen at the higher percentiles. For the rest of the records, the Same Floor error is the distance in meters between the actual and the predicted landmark on the floor. The errors at the higher percentiles are essential for security applications, for example, an error by a floor in localizing a suspect can make or break the evidence. We want to minimize both the errors.

The columns under Baseline in Table 3 show the results. The errors are high for the low cardinalities. We see that the errors for 2.4 GHz are more significant compared to that of 5 GHz. This is despite the fact that more APs hear clients and frequency of transmission of scanning frames is more for 2.4 GHz. This means that the variation in RSSI, including that induced by the transmission power control, has a significant impact on the cardinality and therefore on the localization errors.

#### 4.4 Reducing the Localization Errors

As seen in Table 2, the causes are conflicting to each other. Neither 2.4 nor 5 GHz has all the causes in its favor. Therefore, getting rid of the two issues is not trivial. For improvement, we take a position to make the best use of whatever RSSI we receive during the online

phase. We use heuristic to select the APs from the online phase to reduce localization errors.

We know the location of each AP. We use this information to shortlist the APs from the online fingerprint. The algorithm first selects a floor and then shortlists all the APs that are located on the same floor. We use the shortlisted APs to find a match with offline fingerprints. For selecting the floor, we explore three heuristics – (a) Maximum Number of APs - the floor for which the maximum APs are reporting the client, (b) AP with Maximum RSSI - the floor from which the strongest RSSI is received, and (c) AP of Association - the floor of AP to which the client is presently associated with.

Table 3 shows how the localization errors vary for the three heuristics for  $85^{th}$  percentile values. There is clearly an improvement for both Same Floor and Different Floor errors. Floor detection with Maximum Number of APs gives the least improvement. In fact, until cardinality 4 it performs worse than the Baseline. A cause behind this is that the distant APs, specially in 2.4 GHz that has longer transmission distance, respond and thus localization errors increase. Next, is the floor detection with the AP with Maximum RSSI and AP of Association. The AP with Maximum RSSI or the AP of Association are typically closest to the client, except when the controller does load balancing and transmit power control. There is marginal improvement for 5 GHz. Since the Table 3 only showed data for  $85^{th}$  percentile, we plot the CDF of error for cardinality=1 in Figure 8 for 2.4 GHz. We see that the error reduces for all the percentiles. We see similar results for other cardinalities, but we don't include them due to space constraint.

We compare our results with Signal SLAM [37] which is deployed in a public space like mall since we also have similar deployments. We have similar observations in other venues too. We find their  $90^{th}$  percentile is about 15 meters. We perform similar. In fact, their AP visibility algorithm has  $90^{th}$  percentile as 24.3 meters. We perform better than this in 5 GHz. Given the complexity of the algorithm Signal SLAM incorporates and the amount of sensing it needs, we believe even with few meters of accuracy our approach is better; particularly because its simple and scalable.

## 5 DISCUSSION

Now, we discuss the practical challenges encountered while localizing clients in real deployments and limitations of our solution.

## 5.1 Challenges Of Real Deployments

Real deployments have myriad of practical challenges that hamper the efficiency and the accuracy of an empirical study. For instance, there can be sudden and unexpected crowd movement which is known to increase signal variations. Furthermore, as and when required network administrators either replace old APs or deploy new APs. These administrative decisions are not under our control. However, such changes severely affect the offline fingerprints and change the floor heat maps that ultimately affect location accuracy. Preparing fingerprints for an entire campus with several thousands of landmarks is already tedious, such developments make the process of iterations even more cumbersome.

Beyond insufficient measurement and latency issues, various contextual dynamics makes the fingerprint-based system erroneous. The primary reason is that such dynamic changes results in significant fluctuation in RSSI measurements, which affects the distance calculation of the localization algorithms. These fluctuations can happen quite frequently as there are many different factors affecting RSSI between an AP and its clients, such as crowds blocking the signal path, AP-side power-control for load balancing, and client-side power control to save battery. In Section 4 we shed light on most of these factors. However, we leave a full evaluation for future work. Lastly, all MAC addresses in our system are anonymized. We do not do a device to person mapping to preserve user privacy.

## 5.2 Limitations

A major limitation of this work is that we have not considered an exhaustive set of devices. Given a multitude of device vendors, it is practically impossible to consider all set of devices for this kind of in-depth analysis. We did cover the latest set of devices, though, including iPhone and Android devices. The second limitation is that even though we collected the data for both lightly (semester off, very few students on campus) and heavily loaded (semester on, most students on campus) days. We tested our observations on the lightly loaded dataset but, only on a subset of heavily loaded days. We do not yet know the behavior of system during heavily loaded days, in its entirety. Specifically, the load, concerning the number of clients and traffic is expected to increase interference and thus, signal variations. However, this study is still a part of future work. The third limitation of this work is that we do not consider the effect of MAC address randomization algorithms which make clients intractable. Although there is an active field of research that suggest ways to map randomized MAC to actual MAC [17], but given its complexity we do not employ these.

## 6 CONCLUSION

To conclude, we presented two major issues that need to be addressed to perform server-side localization. We validated these challenges with a huge data from a production WLAN deployed across a university campus. We discussed the causes and their impact on these challenges. We proposed heuristics that handle the challenges and reduce the localization errors. Our findings apply to all the server-side localization algorithms, which use fingerprinting techniques. Most of this work provides real-world evidence of “where” and “what” may go wrong for practically localizing clients in a device agnostic manner.

## REFERENCES

- [1] Abhishek Goswami, Luis E. Ortiz, and Samir R. Das. 2011. WiGEM: A Learning-based Approach for Indoor Localization (*CoNEXT*). <http://doi.acm.org/10.1145/2079296.2079299>
- [2] Anshul Rai, Krishna Kant Chintalapudi, Venkata N. Padmanabhan, and Rijurekha Sen. 2012. Zee: Zero-effort Crowdsourcing for Indoor Localization (*MobiCom*). <http://doi.acm.org/10.1145/2348543.2348580>
- [3] Anthony LaMarca, Yatin Chawathe, Sunny Consolvo, Jeffrey Hightower, Ian Smith, James Scott, Timothy Sohn, James Howard, Jeff Hughes, Fred Potter, Jason Tabert, Pauline Powledge, Gaetano Borriello, and Bill Schilit. 2005. Place Lab: Device Positioning Using Radio Beacons in the Wild (*Pervasive Computing*). [https://doi.org/10.1007/11428572\\_8](https://doi.org/10.1007/11428572_8)
- [4] Aruba. 2017. <https://www.arubanetworks.com>.
- [5] Aruba. 2017. Deep dive Radio technologies for indoor location. <https://community.arubanetworks.com/aruba>.
- [6] C. Wu, Z. Yang, Y. Liu, and W. Xi. 2013. WILL: Wireless Indoor Localization without Site Survey (*IEEE TPDS*).
- [7] Cisco. 2017. <https://www.cisco.com>.
- [8] Coleman and Westcott. 2009. *CWNA Certified Wireless Network Administrator Official Study Guide*. John Wiley & Sons.
- [9] D. Niculescu and Badri Nath. 2003. Ad hoc positioning system (APS) using AOA (*Infocom*).
- [10] D. Turner, S. Savage, and A. C. Snoeren. 2011. On the empirical performance of self-calibrating WiFi location systems (*LCN*).
- [11] Dheryta Jaisinghani, Vinayak Naik, Sanjit K. Kaul, and Sumit Roy. 2017. Sniffer-based inference of the causes of active scanning in WiFi networks (*NCC*).
- [12] Domenico Giustiniano and Stefan Mangold. 2011. CAESAR: Carrier Sense-based Ranging in Off-the-shelf 802.11 Wireless LAN (*CoNEXT*). <http://doi.acm.org/10.1145/2079296.2079306>
- [13] Donny Huang, Rajalakshmi Nandakumar, and Shyamnath Gollakota. 2014. Feasibility and Limits of Wi-Fi Imaging (*SenSys*). <http://doi.acm.org/10.1145/2668332.2668344>
- [14] Fan Li, Chunshui Zhao, Guanzhong Ding, Jian Gong, Chenxing Liu, and Feng Zhao. 2012. A reliable and accurate indoor localization method using phone inertial sensors (*UbiComp*).
- [15] He Wang, Souvik Sen, Ahmed Elgohary, Moustafa Farid, Moustafa Youssef, and Romit Roy Choudhury. 2012. No Need to War-drive: Unsupervised Indoor Localization (*MobiSys*). <http://doi.acm.org/10.1145/2307636.2307655>
- [16] Hyuk Lim, LuChuan Kung, Jennifer C. Hou, and Haiyun Luo. 2010. Zero-configuration Indoor Localization over IEEE 802.11 Wireless Infrastructure (*Wireless Networks*). <http://dx.doi.org/10.1007/s11276-008-0140-3>
- [17] Jeremy Martin, Travis Mayberry, Collin Donahue, Lucas Foppe, Lamont Brown, Chadwick Riggins, Erik C. Rye, and Dane Brown. 2017. A Study of MAC Address Randomization in Mobile Devices and When it Fails (*CoRR*). <http://arxiv.org/abs/1703.02874>
- [18] Jie Xiong and Kyle Jamieson. 2013. ArrayTrack: A Fine-Grained Indoor Location System (*NSDI*). <https://www.usenix.org/conference/nsdi13/technical-sessions/presentation/xiong>
- [19] John Krumm and John Platt. 2003. *Minimizing Calibration Effort for an Indoor 802.11 Device Location Measurement System*. Technical Report. NIPS.
- [20] Jon Gjengset, Jie Xiong, Graeme McPhillips, and Kyle Jamieson. 2014. Phaser: Enabling Phased Array Signal Processing on Commodity WiFi Access Points (*MobiCom*). <http://doi.acm.org/10.1145/2639108.2639139>
- [21] Kasthuri Jayarajah, Rajesh Krishna Balan, Meera Radhakrishnan, Archan Misra, and Youngki Lee. 2016. LiveLabs: Building In-Situ Mobile Sensing & Behavioural Experimentation TestBeds (*MobiSys*). <http://doi.acm.org/10.1145/2906388.2906400>

- [22] Krishna Chintalapudi, Anand Padmanabha Iyer, and Venkata N. Padmanabhan. 2010. Indoor Localization Without the Pain (*MobiCom*). <http://doi.acm.org/10.1145/1859995.1860016>
- [23] Lei Yang, Qiongzhen Lin, Xiangyang Li, Tianci Liu, and Yunhao Liu. 2015. See Through Walls with COTS RFID System! (*MobiCom*). <http://doi.acm.org/10.1145/2789168.2790100>
- [24] Liqun Li, Guobin Shen, Chunshui Zhao, Thomas Moscibroda, Jyh Han Lin, and Feng Zhao. 2014. Experiencing and Handling the Diversity in Data Density and Environmental Locality in an Indoor Positioning Service (*MobiCom*). <http://doi.acm.org/10.1145/2639108.2639118>
- [25] Liqun Li, Pan Hu, Chunyi Peng, Guobin Shen, and Feng Zhao. 2014. Epsilon: A Visible Light Based Positioning System (*NSDI*). <http://dl.acm.org/citation.cfm?id=2616448.2616479>
- [26] H. Liu, H. Darabi, P. Banerjee, and J. Liu. 2007. Survey of Wireless Indoor Positioning Techniques and Systems (*IEEE SMC*).
- [27] Hongbo Liu, Yu Gan, Jie Yang, Simon Sidhom, Yan Wang, Yingying Chen, and Fan Ye. 2012. Push the Limit of WiFi Based Localization for Smartphones (*MobiCom*). <http://doi.acm.org/10.1145/2348543.2348581>
- [28] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. 2015. SpotFi: Decimeter Level Localization Using WiFi (*SIGCOMM*). <http://doi.acm.org/10.1145/2829988.2787487>
- [29] Alex T. Mariakakis, Souvik Sen, Jeongkeun Lee, and Kyu-Han Kim. 2014. SAIL: Single Access Point-based Indoor Localization (*MobiSys*). <http://doi.acm.org/10.1145/2594368.2594393>
- [30] Martin Azizyan, Ionut Constandache, and Romit Roy Choudhury. 2009. SurroundSense: Mobile Phone Localization via Ambience Fingerprinting (*MobiCom*). <http://doi.acm.org/10.1145/1614320.1614350>
- [31] Md Abdullah Al Hafiz Khan, H M Sajjad Hossain, and Nirmalya Roy. 2015. Infrastructure-less Occupancy Detection and Semantic Localization in Smart Environments (*MOBIQUITOUS*). <http://dx.doi.org/10.4108/eai.22-7-2015.2260062>
- [32] Moustafa Youssef, Adel Youssef, Chuck Rieger, Udaya Shankar, and Ashok Agrawala. 2006. PinPoint: An Asynchronous Time-Based Location Determination System (*MobiSys*). <http://doi.acm.org/10.1145/1134680.1134698>
- [33] Moustafa Youssef and Ashok Agrawala. 2005. The Horus WLAN Location Determination System (*MobiSys*). <http://doi.acm.org/10.1145/1067170.1067193>
- [34] Naveed Ul Hassan, Aqsa Naeem, Muhammad Adeel Pasha, Tariq Jadoon, and Chau Yuen. 2015. Indoor Positioning Using Visible LED Lights: A Survey. *ACM Comput. Surv.* (2015). <http://doi.acm.org/10.1145/2835376>
- [35] Nilanjan Banerjee, Sharad Agarwal, Paramvir Bahl, Ranveer Chandra, Alec Wolman, and Mark Corner. 2010. Virtual Compass: Relative Positioning to Sense Mobile Social Interactions (*Pervasive Computing*). [http://dx.doi.org/10.1007/978-3-642-12654-3\\_1](http://dx.doi.org/10.1007/978-3-642-12654-3_1)
- [36] P. Bahl and V. N. Padmanabhan. 2000. RADAR: an in-building RF-based user location and tracking system (*Infocom*).
- [37] P. Mirowski, T. K. Ho, Saehoon Yi, and M. MacDonald. 2013. SignalSLAM: Simultaneous localization and mapping with mixed WiFi, Bluetooth, LTE and magnetic signals (*IPIN*).
- [38] Pan Hu, Liqun Li, Chunyi Peng, Guobin Shen, and Feng Zhao. 2013. Pharos: Enable Physical Analytics Through Visible Light Based Indoor Localization (*HotNets*). <http://doi.acm.org/10.1145/2535771.2535790>
- [39] Rajalakshmi Nandakumar, Krishna Kant Chintalapudi, and Venkata N. Padmanabhan. 2012. Centaur: Locating Devices in an Office Environment (*MobiCom*). <http://doi.acm.org/10.1145/2348543.2348579>
- [40] Rajesh Krishna Balan, Archan Misra, and Youngki Lee. 2014. Live-Labs: Building an In-situ Real-time Mobile Experimentation Testbed (*HotMobile*). <http://doi.acm.org/10.1145/2565585.2565597>
- [41] Sanjib Sur, Teng Wei, and Xinyu Zhang. 2014. Autodirective Audio Capturing Through a Synchronized Smartphone Array (*MobiSys*). <http://doi.acm.org/10.1145/2594368.2594380>
- [42] Souvik Sen, Božidar Radunovic, Romit Roy Choudhury, and Tom Minka. 2012. You Are Facing the Mona Lisa: Spot Localization Using PHY Layer Information (*MobiSys*). <http://doi.acm.org/10.1145/2307636.2307654>
- [43] Souvik Sen, Dongho Kim, Stephane Laroche, Kyu-Han Kim, and Jeongkeun Lee. 2015. Bringing CUPID Indoor Positioning System to Practice (*WWW*). <https://doi.org/10.1145/2736277.2741686>
- [44] Stephen P. Tarzia, Peter A. Dinda, Robert P. Dick, and Gokhan Memik. 2011. Indoor Localization Without Infrastructure Using the Acoustic Background Spectrum (*MobiSys*).
- [45] Stuart A. Golden and Steve S. Bateman. 2007. Sensor Measurements for Wi-Fi Location with Emphasis on Time-of-Arrival Ranging (*TMC*). <http://dx.doi.org/10.1109/TMC.2007.1002>
- [46] Swarun Kumar, Stephanie Gil, Dina Katabi, and Daniela Rus. 2014. Accurate Indoor Localization with Zero Start-up Cost (*MobiCom*). <http://doi.acm.org/10.1145/2639108.2639142>
- [47] Teng Wei and Xinyu Zhang. 2015. mTrack: High-Precision Passive Tracking Using Millimeter Wave Radios (*MobiCom*). <http://doi.acm.org/10.1145/2789168.2790113>
- [48] Viktor Erdelyi, Trung Kien Le, Bobby Bhattacharjee, Peter Druschel, and Nobutaka Ono. 2018. Sonoloc: Scalable positioning of commodity mobile devices (*MobiSys*).
- [49] Wei Wang, Alex X. Liu, Muhammad Shahzad, Kang Ling, and Sanglu Ling. 2015. Understanding and Modeling of WiFi Signal Based Human Activity Recognition (*MobiCom*). <http://doi.acm.org/10.1145/2789168.2790093>
- [50] Yanzi Zhu, Yibo Zhu, Ben Y. Zhao, and Haito Zheng. 2015. Reusing 60GHz Radios for Mobile Radar Imaging (*MobiCom*). <http://doi.acm.org/10.1145/2789168.2790112>
- [51] Yogita Chapre, Prasant Mohapatra, Sanjay Jha, and Aruna Seneviratne. 2013. Received signal strength indicator and its analysis in a typical WLAN system (short paper) (*LCN*).
- [52] Yuanfang Chen, Noel Crespi, Lin Lv, Mingchu Li, Antonio M. Ortiz, and Lei Shu. 2013. Locating Using Prior Information: Wireless Indoor Localization Algorithm (*SIGCOMM*). <http://doi.acm.org/10.1145/2486001.2491688>
- [53] F. Zafari, A. Gkelias, and K. Leung. 2017. A Survey of Indoor Localization Systems and Technologies. *ArXiv e-prints* (2017).
- [54] Zengbin Zhang, David Chu, Xiaomeng Chen, and Thomas Moscibroda. 2012. SwordFight: Enabling a New Class of Phone-to-phone Action Games on Commodity Phones (*MobiSys*). <http://doi.acm.org/10.1145/2307636.2307638>
- [55] Zheng Yang, Chenshu Wu, and Yunhao Liu. 2012. Locating in Fingerprint Space: Wireless Indoor Localization with Little Human Intervention (*MobiCom*). <http://doi.acm.org/10.1145/2348543.2348578>