6-2016

Indoor location error-detection via crowdsourced multi-dimensional mobile data

Savina SINGLA
PEC University of Technology, India

Archan MISRA
Singapore Management University, archanm@smu.edu.sg

Citation

SINGLA, Savina and MISRA, Archan. Indoor location error-detection via crowdsourced multi-dimensional mobile data. (2016). *MobiData '16: Proceedings of the First Workshop on Mobile Data, Singapore, June 30*. 19-24.

Available at: https://ink.library.smu.edu.sg/sis_research/3638

Indoor Location Error-Detection via Crowdsourced Multi-Dimensional Mobile Data

Savina Singla*
PEC University of Technology
Chandigarh, India

Archan Misra
Singapore Management University
Singapore

ABSTRACT

We explore the use of multi-dimensional mobile sensing data as a means of identifying errors in one or more of those data streams. More specifically, we look at the possibility of identifying indoor locations with likely incorrect/stale Wi-Fi fingerprints, by using concurrent readings from Wi-Fi and barometer sensors from a collection of mobile devices. Our key contribution is a novel two-step process: (i) using longitudinal, crowd-sourced readings of (possibly incorrect) Wi-Fi location estimates to statistically estimate the barometer calibration offset of individual mobile devices, and (ii) then, using such offset-corrected barometer readings from devices (that are supposedly collocated) to identify likely errors in indoor localization. We evaluate this approach using data collected from 104 devices collected on the SMU campus over a period of 61 days: our results show that (i) 49% of the devices had barometer offsets that result in errors in floor-level estimation, and (iii) 46% of the Wi-Fi location estimates were potentially incorrect. By identifying specific locations with unusually high fraction of incorrect location estimates, we attempt to more accurately pinpoint the areas that need re-fingerprinting.

1. INTRODUCTION

Mobile sensing data, collected from a large pool of visitors, are being increasingly used to perform movement and behavioral analytics (e.g., group detection [11] or prediction of shopper movement [3]) at various indoor public venues, such as museums, shopping malls & college campuses. The underlying indoor location estimates are mostly based on Wi-Fi fingerprinting techniques (e.g., RADAR) and are known to be quite noisy for two distinct reasons: (a) Wi-Fi RSSI measurements at any given location fluctuate significantly, due to factors such as changes in the indoor layout or fluctuations in crowd density, and (b) the underlying fingerprints themselves become *stale*, as fingerprinting itself requires significant human effort and is not something that can be performed routinely.

*Work performed while visiting Singapore Management University.

MobiData'16, June 30 2016, Singapore, Singapore
© 2016 ACM. ISBN 978-1-4503-4327-5/16/06. . . $15.00
DOI: http://dx.doi.org/10.1145/2935755.2935762

As part of a broader investigation into strategies for *detecting* and *cleaning* errors in large mobile sensor datasets, we address the important question: *can measurements obtained from additional mobile sensors help in identifying likely errors in such Wi-Fi based indoor location estimation?* More specifically, we look at the possible use of concurrently-gathered barometer sensor data (used to estimate floor-level location in [7]) as a means of identifying likely errors in the Wi-Fi RSSI-based location estimates.

The key idea is fairly simple: **when multiple devices share a common location (via Wi-Fi based localization), we expect that their barometer readings will be identical (within error bounds) as well.** If neighboring devices, however, show large deviations in the barometer readings (e.g., if the readings suggest that the devices are on different floors of a building), then one can suspect that the location estimates are faulty. Note that this approach of "crowdsourced corroboration" is distinct from alternative error correction techniques that fuse Wi-Fi and additional sensor data from the same individual device (e.g., smoothing Wi-Fi location estimates with inertial sensor data).

Challenges: Practical realization of this idea, however, has several challenges: (a) *Barometer Offset Errors:* Barometric sensors on smartphones are known to have potentially-high calibration errors– e.g., (offsets of 15 meters reported in [13]). Moreover, this error can be *device-specific*–different units of the same model could have very different errors. Before barometer-based floor estimates are used for corroboration, we must first correct for errors in the barometer readings themselves, without required explicit per-device calibration; (b) *Latency and Error in Wi-Fi location estimates:* A server-side based localization system (as implemented in SMU's LiveLabs testbed [6]) provides more universal, OS-independent coverage. However, commercial implementations of server-side solutions often exhibit update latencies of 3-4 minutes, and achieve accuracies around 6-8 meters. The error detection process must accommodate such latency and error characteristics; (c) *Time-varying errors:* Both barometer and Wi-Fi RSSI measurements vary with changes in the ambient environment–e.g., even in indoor pressurized buildings, true pressure values change due to changes in HVAC operating set-points, while Wi-Fi readings vary with fluctuations in crowd levels. The error detection technique must be robust to such underlying unknown fluctuations.

Key Contributions: We develop a two-step error detection process, consisting of an initial crowdsourcing-based estimation of device-specific barometer offsets, followed by use of such offset-compensated values to identify anomalous Wi-Fi readings. We utilize a reasonably large crowd-sourced dataset, consisting of readings from 104 users on the LiveLabs testbed over a period of 60 days (with 1 hour of barometer reading per day), to make the following contributions:

- *Barometer Offset Estimation with No Additional Sensing:* We propose and evaluate a barometer-offset computation approach, that relies only on "likely to be reliable" readings from devices observed to be concurrently located in close proximity. Our analysis reveals the existence of significant offsets: about 71% of the participant devices had errors of more than 2.5 hPa (corresponding to errors greater than a floor's height). Moreover, such errors were both model and device-dependent: the Samsung Galaxy S5 had average offset of 1.04 hPa, with some specific handsets requiring offsets as high as 3.5 hPa ($\approx 90 meters$)!

- *Wi-Fi Error Detection using Corrected Barometer Data:* We then identify likely incorrect Wi-Fi readings as those segments of location estimates where the Wi-Fi and barometer-based floor level estimates diverge by more than 12 feet. We show that barometer calibration matters: while raw readings would indicate errors in $\approx$ 88% of the Wi-Fi data, the calibrated readings reduce the possible error rate to $\approx$ 46% (given the high latency of Wi-Fi reports, the true error rate is likely lower).

- *Help Pinpoint Locations for Re-fingerprinting:* We finally look at the occurences of anomalous Wi-Fi readings on a per-location basis, and thereby identify the locations/regions that need re-fingerprinting (have a disproportionately high occurrence of such anomalies). Our approach allows us to use a threshold-based approach to priortize such locations–our studies show that over 40% of locations need a fresh set of Wi-Fi fingerprints after 60 days.

2. RELATED WORK

Wi-Fi Fingerprinting & Maintenance: In indoor environments, Wi-Fi fingerprinting approaches, which require the development of a database that contains $(location, [RSSI\ vector])$ readings for each landmark, are the most popular. The key drawback of this approach is the need for extensive manual effort to build this fingerprint database; approaches such as EZ [1] have utilized crowdsourced measurements to progressively develop and update fingerprints, but require active participation by visitors to the venue. The maintenance of accurate fingerprints is, however, key to such indoor location: by evaluating multiple approaches, Farshad et al [2] concluded that fingerprinting itself is as importance as the choice of the localization algorithm.

Barometer-based Context Sensing: Muralidharan et al [7] investigated the use of smartphone-embedded barometers for indoor floor-level localization–their investigations showed that the absolute barometer values could be different for different devices, but the *change in readings* due to floor-level transitions was essentially stable and device-independent. Liu et al [4] applied a similar strategy, using the barometer reading on the ground floor of a building as a reference data point, to directly compute the phone's height indoors by applying Kalman filtering on the sensor stream. In contrast, Xia et al [12] utilize a set of static per-floor barometer sensors to provide a real-time reference value. Majethia et al [5] further showed that the barometer static offset is not constant, but affected by additional context factors, such as the phone's temperature. Barometer sensing has also been used to estimate other outdoor context, such as the detection of transportation mode [10]. In this work, we do not rely on either continuous barometer measurements or use of reference values, but instead estimate device-specific offsets based on location-tagged measurements from a variety of smartphones.

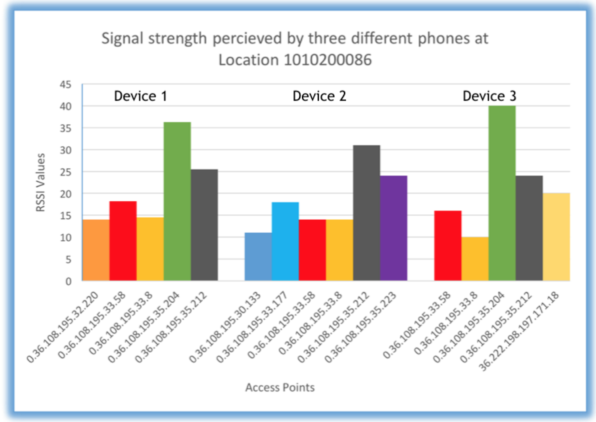

Figure 1: AP reports of RSSI values of 3 collocated devices

Crowdsourced Sensor Calibration: Our work is similar to that of SBC [13], which uses a large body of crowdsourced mobile sensor data to compute barometer drifts. SBC however uses accelerometer traces to identify devices that concurrently perform a specified set of collocated vertical (e.g., taking a lift) and horizontal (e.g., taking a bus) activities, and then computes barometer offsets in a transitive manner. Large-sized crowdsourced datasets have been used for modeling other parameters of mobile devices–e.g., the Constella system [8] used such crowdsourcing data to build models of smartphone energy consumption. In contrast to such work, we focus on the joint use of two noisy sensor streams (barometer & Wi-Fi) to isolate the likely erroneous segments of context based on one of those sensor streams (Wi-Fi based location).

3. ARCHITECTURE AND STUDY DETAILS

In this section we discuss about our overall approach for detecting erroneous Wi-Fi sensing data, and also outline the dataset used in our studies.

3.1 Dataset Details

The results reported in this study (approved by SMU's Institutional Review Board) are based on data from 160 participants in SMU's LiveLabs testbed, which spans 5 academic buildings across the university's downtown campus. This testbed includes a server-side Wi-Fi based localization system, that provides both the raw RSSI readings (for each mobile device, as received at multiple listening APs) and the computed location of each device, based on a customized version of the classical RADAR algorithm. The server-side location system reports a new location reading for each Wi-Fi enabled device on campus once every 3-4 minutes. As the location accuracy is observed to be approx. 6-8 meters, the system reports the location updates at *section-level* granularity (each section is roughly 15m wide). Moreover, each of these 160 participants had a LiveLabs data collector application installed on their Android phone (with their consent)–this collector collected various sensor readings (including barometer, accelerometer, etc.) and uploaded the data to a backend server. Note that each user used their own personal phone-accordingly, the dataset contains over 24 distinct device models (across manufacturers such as Samsung, Google (Nexus), Xiaomi, HTC & Asus). To prevent unnecessary energy drainage, this collector was activated only for one hour (12pm-1pm) daily, over a period of 61 days. Out of these 160 devices, only 104 were observed to have a reasonable number of *useful* reports (i.e., readings where its location is collocated with other devices); we restrict further analysis to this set of 104 devices.

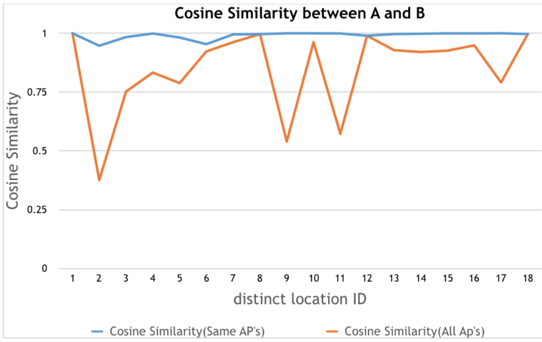

Figure 2: Cosine Similarity of RSSI for 2 collocated clients

3.2 Inaccurate Wi-Fi Data

Given the data, we first inspected the Wi-Fi measurement data reported by multiple concurrently-collocated mobile devices. We discovered that there were indeed *many cases* where the RSSI readings of nearby devices, as reported by one or more APs, diverged significantly. As an example, Figure 1 plots the RSSI values reported for three clients that were truly in the same room simultaneously–each bar/color corresponds to a different reporting AP. We notice that not only are the values for different clients reported by a single AP quite different, even the relative ordering of signal strengths across APs is not consistent (in fact, different devices are heard by a different set of APs).

While absolute values for different devices can indeed be different (due to different transmission power and antenna gain parameters), we expect that their relative values would be similar, at different locations. However, this is not the case either: Figure 2 plots the *Cosine Similarity* (between the RSSI vectors reported by the APs) for two devices, at each of 18 distinct locations where they were concurrently located. Ideally, this similarity should be 1. We see that while the similarity value is indeed close to 1 when we only consider (at each location) only those 'common APs' (that are contained in the reporting set of both devices), this similarity drops significantly when we consider the entire AP vector. Our results illustrate that the *server-side Wi-Fi localization process is indeed noisy, and that some form of error filtering is thus necessary*.

3.3 *High-Level Architecture of our system*

To tackle this problem, we implement an analytics pipeline whose high-level logic (illustrated in Figure 3) is as follows:

1. We first analyze the barometer data (assuming that the computed location tags are correct) to determine the specific offset for each distinct mobile device–the details of this are presented in Section 4.

2. We then use the appropriately calibrated barometer data to compute, for each Wi-Fi data stream, the corresponding floor-level location. We then flag as *anomalous* those location values where the barometer-supplied floor-level estimate of the collocated devices diverge significantly from that obtained via Wi-Fi.

3. We finally use the relative frequency of such anomalous location estimates, computed separately at section-level granularity, to identify the sections that have disproportionately high instances of anomalous locations, and flag them for re-fingerprinting.

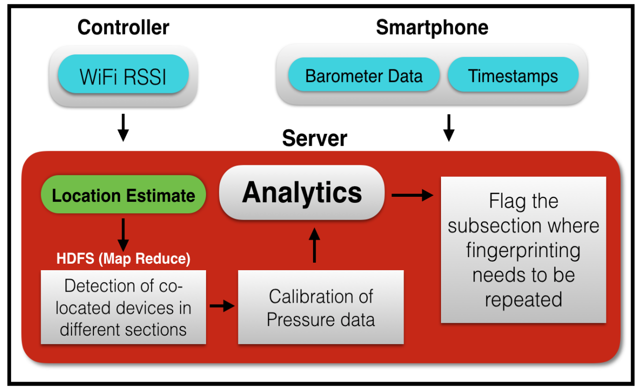

Figure 3: Architecture of the system

4. PER-DEVICE BAROMETER OFFSET

Our barometric calibration strategy works on the basic assumption that the readings for collocated devices should be identical, *after one has compensated for their individual calibration offsets*.

4.1 Pressure to Height Relationship

Equation 1 relates the pressure reading p (in millibars) to the altitude, under normal temperatures [9]:

$$\textbf{Height} = \textbf{1-(p/1013.25)}^{\textbf{0.190284}} * \textbf{145366.45} \tag{1}$$

Roughly speaking, a difference of 1 hPa corresponds to a height difference of 8.3 meters (27.3 feet).

4.2 The Reality of Device-Specific Offsets

We first performed some controlled micro-studies to observe the divergence between different devices. Four devices (two Samsung Galaxy S3s, 1 Samsung Galaxy S4 and one Nexus) were kept at the same place, right next to one another, and collected the pressure data multiple times, across two different days. Figure 4(a) and (b) shows the mean of pressure readings (hPa) recorded by four different devices for two days respectively. Figure 4(c) and (d), shows the difference(delta) between each pair for both the days. We see that, in terms of absolute readings, even phones of the same model show a significant relative difference (the two S3s differ by almost 6 hPa or approx. 150 feet!). However, the relative difference remains constant between the pairs, irrespective of any causality factors like air-conditioning, climate etc., indicating that this difference is due to a *calibration offset*.

We performed additional micro-studies to confirm that the on-body location of the phone did not impact the barometer readings. As part of these studies, 3 different phone models were placed side by side in 3 different locations (held in the hand, placed inside the pant pocket and carried in a backpack). Figure 5 demonstrates that the pressure readings reported by the barometer are independent of how and where the user keeps his phone. Accordingly, we can utilize the crowdsourced data without having to worry about the phone's placement artefact.

4.3 Finding the Device Offset Δ_B

To compute the unknown device-specific offsets, we first identify a collection of devices that are reported (by the Wi-Fi location system) to be located at the same section for contiguous period of at least 20 minutes. We deliberately use a duration of 20 minutes to filter out scenarios where the devices may be moving, not just across different floors but even within a floor (even with a floor, there are differences between sections that are climate-controlled

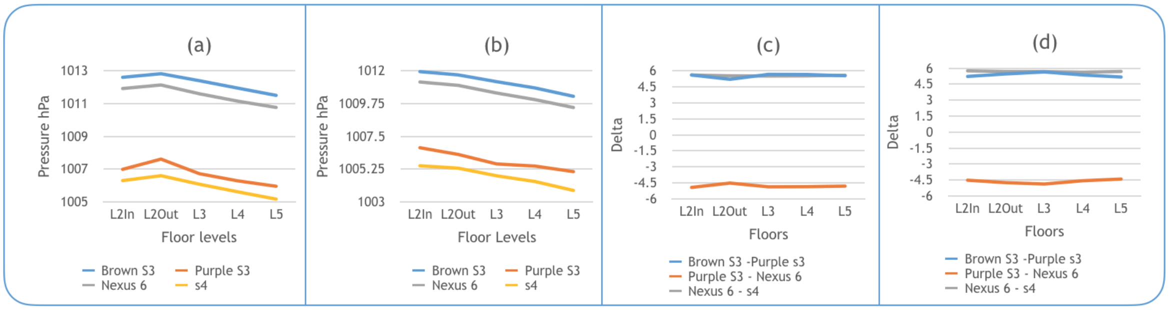

Figure 4: (a)Comparison of different phone model barometer readings Day 1, (b) Comparison of different phone model barometer readings Day 2, (c) Deltas between different phones day 1 and (d) Deltas between different models day 2

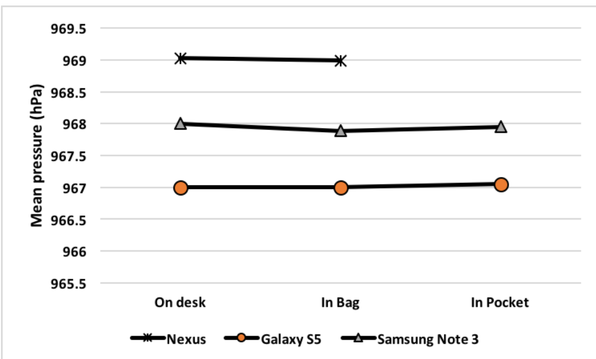

Figure 5: Barometric readings are independent of phone placement.

```
for (each subsection):
        colocatedDevices[] <— group all the devices located
        for (each device in colocatedDevices[])
                pressureSum <— pressureSum + deviceBarometerReading
        find meanPressure
        for (each device in colocatedDevices[])
                localOffSet <— devicePressure - meanPressure
for (each device):
        localOffsetList[] <— localOffsets at eachLocation
        deltaValue <— median(localOffsetList[])
for (each device):
        calibratedData <— rawData - deltaValue
```

Figure 6: Algorithm to find $\Delta_B(.)$

and sections that are not enclosed). Over the entire study, there are many such collections formed, each with a potentially distinct set of devices, and corresponding to a different (section, time) tuple.

For each such collection, we compute the "collection-average" (μ) of the barometer readings reported from the constituent devices, and then compute the *collection-specific offset*, for each device, as the difference of the device's mean readings from this collection-average. As each device will be part of multiple such collections, we will eventually obtain a set of collection-specific offsets for this device. We then obtain our estimated *device-specific offset*, denoted by $\Delta_B(.)$, as the *median* of this set (the median helps eliminate the outliers). Figure 6 provides the pseudocode for this offset computation process. A negative Δ_B indicates that its pressure readings are lower than that of other comparable devices (the reverse holds for positive Δ_B values).

We use an example scenario to further illustrate this approach. Assume a scenario where we have a group formed of four devices in section S_1: $\{D_1,D_2,D_3,D_4\}$, three devices in section S_2: $\{D_1,D_3,D_5\}$ and six devices in section $S_3\{D_1,D_2,D_3,D_6,D_7,D_8\}$. To proceed further we find the mean value of each section (collection): $\mu(S_1) = \mu(D_1) + \mu(D_2) + \mu(D_3)+ \mu(D_4)$ / 4 where $\mu(D_n)$ is the average of the pressure readings reported by device n over that interval. We similarly calculate $\mu(S_2)$ and $\mu(S_3)$. Now the

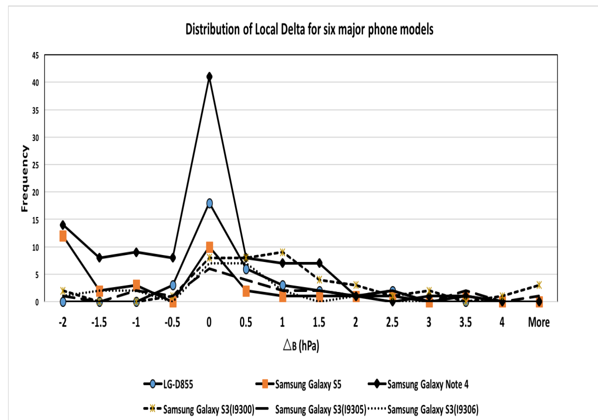

Figure 7: Δ_B Distribution for Different Phone Models

collection-specific offset $\{O_1,O_2,O_3, O_4\}$ for all the four devices at Subsection S_1 comes out to be ($\mu(D_1)$ - $\mu(S_1)$), ($\mu(D_2)$ - $\mu(S_1)$), ($\mu(D_3)$ - $\mu(S_1)$) and ($\mu(D_4)$ - $\mu(S_1)$) respectively. In the same manner local offsets of all the devices are calculated in the other sections.

Suppose the device D_1 is seen in five such different "collections", then we will obtain 5 different local offsets for device D_1 say, $O_{(1,1)}$, $O_{(1,2)}$, $O_{(1,3)}$, $O_{(1,4)}$, $O_{(1,5)}$. For D_1, its *device-specific offset*, $\Delta_B(D_1)$ is then obtained by finding the median of this set.

4.4 Experimental Results & Insights

We applied the algorithm described above to compute the value of Δ_B for each of the 104 devices. Figure 7 plots the distribution of Δ_B for different handsets, for 6 commonly occuring phone models. Clearly, we see that the calibration offset is definitely (a) *model-dependent*: some models have a smaller range of Δ_B values, implying that the barometer readings vary within tighter bounds. and also (b) *device-dependent*: For many models, the range is, in fact, quite wide, often spanning $\approx$ 6 hPa (equivalent to a variation of over 50 meters!).

Figure 8 then zooms in on the different offset values for multiple devices of key selected popular models (around 20 different handsets of Galaxy Note 3 and Note 4, almost 8 different handsets of Galaxy S3 and LG D855). We see that the offset values are themselves not uniform. Moreover, we also noticed that, for the same device, its "collection-specific offset" was not constant and could vary by $\approx$ 0.5 hPa. This could be due to two reasons: (i) the location estimate was incorrect (the high Wi-Fi reporting latency implies that a user could have temporarily visited another floor and returned within 3-4 minutes), and thus the barometer readings are being erroneously compared with devices on another floor, or (ii)

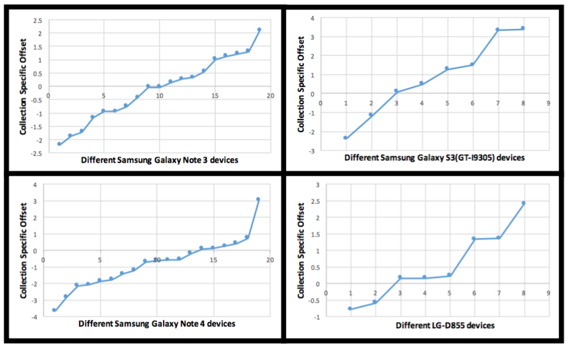

Figure 8: "Collection-specific Offsets" (4 Models)

the offset itself may have changed, due to dynamic conditions such as the device's operating temperature. More importantly, these results suggest that the offsets are not simply transitive, and thus validate our strategy of using statistical outlier filtering to compute the most-representative offset value.

We next verified that our offset computation was indeed accurate–i.e., it could truly result in identical readings across multiple collocated devices. To test this, we performed a control study with 3 of the handsets (belonging to our lab) placed next to one another. Figure 9 shows both the raw (uncalibrated) and offset-adjusted readings of the three devices. We see that our calibration process is quite effective, as the application of the computed offset results in barometer readings that truly become device-independent.

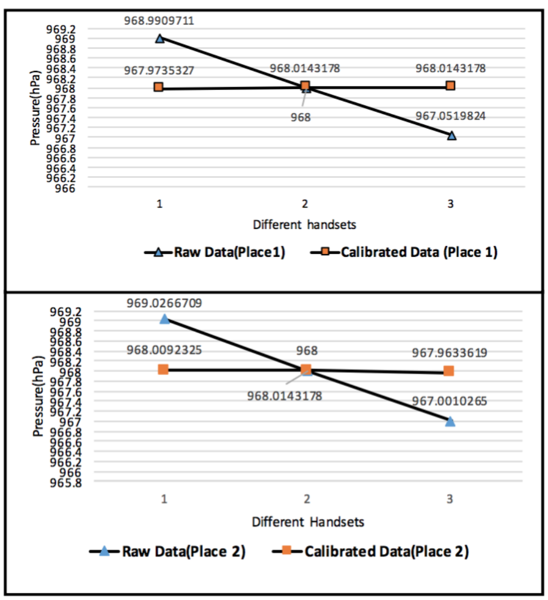

Figure 9: Pressure Data of different phones before and after calibration in the controlled environment.

Overall, our results showed that the computed Δ_B value was higher than 2.4 hPa (i.e., approximately the height of a single floor) for 71% of the devices, indicating the critical importance of proper calibration prior to the use of such barometer sensor data for detecting Wi-Fi localization errors.

5. INCORRECT WI-FI LOCATION

Having estimated the per-device offset Δ_B, we can then 'correct' the barometer readings for each device by subtracting Δ_B from the original raw readings. We now describe how we use this corrected, calibrated barometer readings to identify (a) the likely Wi-Fi location errors and (b) the sections that need fingerprinting.

5.1 Pressure difference-based Wi-Fi Anomalies

First, we observe that the height of a single level of a building on the SIS campus varies between 10-12 feet. Given our earlier observation of 1hPa pressure difference being equivalent to ≈ 27 feet, we can see that a pressure changes of ≈ 0.45hPa corresponds to a single floor change. Accordingly, we assume that the 'corrected barometer' readings of different devices must lie within a range of 0.45 hPa, if they are truly collocated (on the same floor). However, if their readings diverge by more than 0.45 hPa, it is likely that one or more of the location estimates are erroneous.

Detecting Erroneous Wi-Fi Measurements: Based on these insights, we detect the likely time segments with location estimation errors as follows. For each "collection" (i.e., a specific section observed over a 20 minute interval), we compute the number of "anomalous Wi-Fi readings"–i.e., one where the Wi-Fi reported location (floor level) differs from the barometer-based floor level estimate (after applying the $\Delta_B(.)$ offset) by more than 0.45 hPa. We then compute the total fraction of such readings (for each section).

Identify Locations that need Re-fingerprinting: To identify the locations (sections) that need fingerprinting, we proceed in analogous fashion. For each section, we consider all the "collections" corresponding to that section (i.e., all the time-intervals where we observed a set of collocated mobile devices within that section). We then compute the total percentage of "anomalous Wi-Fi readings" for that section. If this fraction exceeds a configurable threshold γ, we declare that section as having an unacceptably high rate of inconsistent Wi-Fi readings, and thus flag it for re-fingerprinting.

5.2 Results

We first plot (in Figure 10) the percentage of anomalous Wi-Fi readings (those where the barometer-based floor estimate differs by more than one level from the Wi-Fi based estimate), at several representative sections. Figure 10 plots this percentage based on both (i) the raw barometer readings and (ii) the offset-corrected barometer readings. We notice that, after offset correction, almost 46% of subsections (in contrast to a much higher number of 88% when the raw readings are used) have anomaly rates higher than 50%. Moreover, 71% of all the devices exhibited at least one discrepancy between barometer vs. Wi-Fi based location estimates. *These results clearly demonstrate the importance of barometer-offset correction in reducing false-positives and false-negatives during Wi-Fi location error detection.*

We then provide a sensitivity analysis on how the choice of γ affects the number of sections needing re-fingerprinting. Figure 11 illustrates this tradeoff: clearly, a higher value of γ allows for a larger tolerance to inferred Wi-Fi location estimation errors, and thus reduces the number of locations identified for re-fingerprinting. Note that, because of our high Wi-Fi update latency, our discrepancy estimates are higher than reality, as we cannot capture situations where a user visits another locations and returns within 3-4 mins. We notice that, if $\gamma = 70\%$, then $\approx 26\%$ of the sections are identified as candidates for re-fingerprinting.

Note that these error percentages (which can appear to be quite high) apply only to the instantaneous Wi-Fi based location estimates. In actual location tracking systems, additional path-based smoothing techniques are typically used to reduce such instantaneous errors.

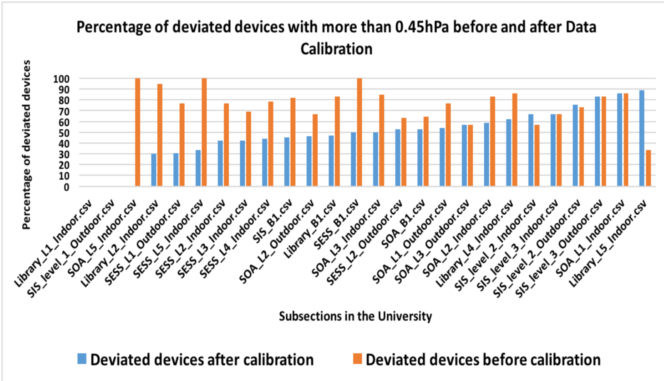

Figure 10: Fraction of Wi-Fi locations inconsistent with concurrent barometer-based floor estimates

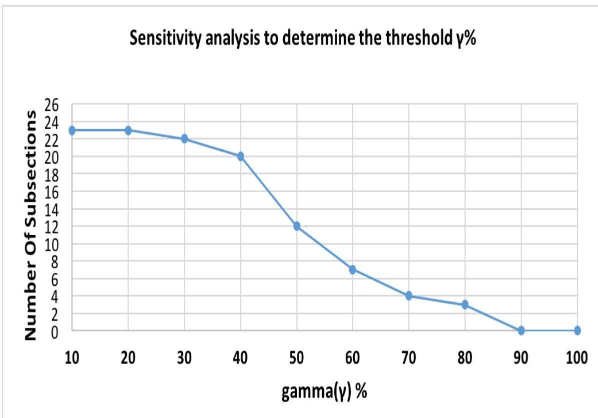

Figure 11: Number of sections where $\geq \gamma\%$ anomalous readings are observed.

6. CONCLUSION

In this paper, we have devised a way to combine crowd-sourced, possibly erroneous, Wi-Fi RSSI and barometer-based pressure data, to correct for the errors in both sensing modalities. Using longitudional observations of barometer readings from multiple supposedly-collocated devices, we figure out the device-specific calibration offset in the barometric sensor. This turned out to be a vital step: we found that the calibration errors could be quite large ($\approx$ 50 meters) and that significant variations existed both across and within device models. Using this offset-adjusted barometer data, we were then able to identify anomalous Wi-Fi location estimates more precisely: while raw barometer readings identified 88% of Wi-Fi estimates as erroneous, the number dropped to 46% when offset-adjusted data was used. This error correction process becomes an important intermediary step in the automatic identification of indoor locations where fingerprints need to be updated.

On the broader scale, our work illustrates how careful processing of large, multi-dimensional mobile sensing datasets can help eliminate or reduce errors in each distinct sensing channel. In ongoing work, we shall apply a similar approach to further refine and isolate errors, by taking crowd-sourced inertial data into consideration as well.

7. ACKNOWLEDGEMENT

This research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its IDM Futures Funding Initiative. Savina Singla's work is also conducted as a part of the project "CARTS- Communication Assisted Road Transportation System" supported by ITRA, Media Lab Asia.

8. REFERENCES

[1] Krishna Chintalapudi, Anand Padmanabha Iyer, and Venkata N. Padmanabhan. Indoor localization without the pain. In *Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking*, MobiCom '10, 2010.

[2] Azade Farshad, Jiwei Li, Mahesh K Marina, and Francisco J Garcia. A microscopic look at wifi fingerprinting for indoor mobile phone localization in diverse environments. In *Indoor Positioning and Indoor Navigation (IPIN), 2013 International Conference on*, pages 1–10. IEEE, 2013.

[3] SangJeong Lee, Chulhong Min, Chungkuk Yoo, and Junehwa Song. Understanding customer malling behavior in an urban shopping mall using smartphones. In *Proceedings of the 2013 ACM Conference on Pervasive and Ubiquitous Computing Adjunct Publication*, UbiComp '13 Adjunct, 2013.

[4] Guangwen Liu, Masayuki Iwai, Yoshito Tobe, Dunstan Matekenya, Khan Muhammad Asif Hossain, Masaki Ito, and Kaoru Sezaki. Beyond horizontal location context: Measuring elevation using smartphone's barometer. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, UbiComp '14 Adjunct, pages 459–468, 2014.

[5] Rahul Majethia, Varun Mishra, Prasad Pathak, Divya Lohani, Debopam Acharya, and Seema Sehrawat. Contextual sensitivity of the ambient temperature sensor in smartphones. In *Communication Systems and Networks (COMSNETS), 2015 7th International Conference on*, pages 1–8. IEEE, 2015.

[6] Archan Misra and Rajesh Krishna Balan. Livelabs: Initial reflections on building a large-scale mobile behavioral experimentation testbed. *SIGMOBILE Mob. Comput. Commun. Rev.*, 17(4), December 2013.

[7] Kartik Muralidharan, Azeem Javed Khan, Archan Misra, Rajesh Krishna Balan, and Sharad Agarwal. Barometric phone sensors: More hype than hope! In *Proceedings of the 15th Workshop on Mobile Computing Systems and Applications*, page 12. ACM, 2014.

[8] Ella Peltonen, Eemil Lagerspetz, Petteri Nurmi, and Sasu Tarkoma. Constella. *Pervasive Mob. Comput.*, 26(C):71–90, February 2016.

[9] Pressure altitude. https://en.wikipedia.org/wiki/Pressure_altitude. Accessed: 2016-04-09.

[10] Kartik Sankaran, Minhui Zhu, Xiang Fa Guo, Akkihebbal L Ananda, Mun Choon Chan, and Li-Shiuan Peh. Using mobile phone barometer for low-power transportation context detection. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems*, pages 191–205. ACM, 2014.

[11] Rijurekha Sen, Youngki Lee, Kasthuri Jayarajah, Archan Misra, and Rajesh Krishna Balan. Grumon: Fast and accurate group monitoring for heterogeneous urban spaces. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems*, SenSys '14, 2014.

[12] Hao Xia, Xiaogang Wang, Yanyou Qiao, Jun Jian, and Yuanfei Chang. Using multiple barometers to detect the floor location of smart phones with built-in barometric sensors for indoor positioning. *Sensors*, 15(4), 2015.

[13] Haibo Ye, Tao Gu, Xianping Tao, and Jian Lu. Sbc: Scalable smartphone barometer calibration through crowdsourcing. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, MOBIQUITOUS '14, 2014.