

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

2-2014

Barometric phone sensors: More hype than hope!

Kartik MURALIDHARAN

Singapore Management University, kartikm.2010@smu.edu.sg

Azeem Javed KHAN

Oriental Institute Of Management

Archan MISRA

Singapore Management University, archanm@smu.edu.sg

Rajesh Krishna BALAN

Singapore Management University, rajesh@smu.edu.sg

Sharad AGARWAL

Microsoft Research

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Software Engineering Commons](#)

Citation

MURALIDHARAN, Kartik; KHAN, Azeem Javed; MISRA, Archan; BALAN, Rajesh Krishna; and AGARWAL, Sharad. Barometric phone sensors: More hype than hope!. (2014). *HotMobile '14: Proceedings of the 15th Workshop on Mobile Computing Systems and Applications, Santa Barbara, CA, February 26-27*. 12:1-6. Available at: https://ink.library.smu.edu.sg/sis_research/2061

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.

Barometric Phone Sensors – More Hype Than Hope!

Kartik Muralidharan[†], Azeem Javed Khan[§], Archan Misra[‡],
Rajesh Krishna Balan[†], Sharad Agarwal[‡]

[†]Singapore Management University, [§]Oriental Institute of Management, [‡]Microsoft Research
[†]{kartikm.2010,archanm,rajesh}@smu.edu.sg, [§]azeem@oes.ac.in, [‡]sagarwal@microsoft.com

ABSTRACT

The inclusion of the barometer sensor in smartphones signaled an opportunity for aiding indoor localization efforts. In this paper, we therefore investigate a possible use of the barometer sensor for detecting vertically oriented activities. We start by showing the accuracies of various commodity measurement devices and the challenges they bring forth. We then show how to use the barometer values to build a predictor that can detect floor changes and the mode (elevator, escalator, or stairs) used to change floors with nearly 100% accuracy. We validate these properties with data collected using 3 different measurement devices from 7 different buildings. Our investigation reveals that while the barometer sensor has potential, there is still a lot left to be desired.

1. INTRODUCTION

There are a number of papers in ACM MobiSys and HotMobile that have profited from the plethora of sensors on smartphones, including the gyroscope, accelerometer, compass and microphone. Many have employed these sensors for location and contextual awareness. Imagine our excitement when we noticed a *new* sensor appearing on smartphones – the barometer – that has not been tapped by our community. This sensor has appeared on “top shelf” devices, including the Samsung Galaxy S4, Google Nexus 4 and 10.

Naturally, our first instinct was to use this sensor to solve one of the key challenges in mobile computing – indoor location. Since the barometer sensor measures ambient atmospheric pressure, and since that pressure varies primarily with height, our objective is to determine which floor of an indoor building the user is on.

In practice, this is a challenging problem. Atmospheric pressure outdoors varies throughout the course of a day, subject to natural weather phenomena. Some large indoor buildings are pressurized, either for heating and cooling efficiency or for sanitary reasons. While in some indoor location scenarios, an error of a few meters may be tolerable, in this situation an error of even 2 meters might indicate the wrong floor. Given these real-world challenges, can we still determine which floor of a building a user is on?

For this paper, we conducted extensive measurements using several devices in 7 buildings across Singapore across different times

of the day and days of the week. As we anticipated, there is *absolute pressure variation* observed by a phone that remains stationary on a floor. Worse, we observed *inter- and intra- phone model variations* in absolute pressure readings that were large enough that two phones next to each other of the same model give pressure readings that are off by the equivalent of over one building floor. We subsequently focus on *relative pressure differences*, the intuition being that a change in pressure equivalent to a floor height indicates that the phone has changed floors. Unfortunately, we also encounter *sensor drift*, where the readings from a barometer sensor vary not necessarily because the actual pressure has changed. Despite these challenges, we have made the following findings:

- 1 Absolute pressure readings are unreliable indicators of floor-level information, and there are several subtle pitfalls that practitioners must consciously tackle.
- 2 The pressure difference between two floors of a building shows significantly less variability across time and phone model. Using readings from a single device, we can determine that *i*) the user has changed floors, *ii*) and how many floors have changed (or the change in absolute height) with almost perfect accuracy.
- 3 The three common modes for changing floors (escalators, stairs and elevators) can be distinguished by looking at the rate of pressure change. These rates are distinctly higher than the typical drift we observe on barometer sensors.
- 4 The barometer is significantly more robust than an accelerometer at detecting vertical activities. However, it consumes similar amounts of energy, even though on paper, the barometer is a very cheap sensor.

Unfortunately, there are key challenges that prevent us from using the barometer sensor on phones for determining which floor a user is on. Since the absolute value reported by different phones varies, we cannot simply use sensors embedded in the building to provide reference values to compare against and then determine what height or floor a user is on. Many large buildings have multiple entrances that may be on different floors, hence it is not trivial to identify which floor a user has entered on and then simply count up or down each time the pressure changes. Finally, while the barometer sensor could be used in conjunction with an infrastructure-based

This research is supported by the Singapore National Research Foundation under its IDM Futures Funding Initiative and administered by the Interactive Digital Media Programme Office, Media Development Authority and supported in part by the Singapore Ministry of Education Academic Research Fund Tier 2 under the research grant MOE2011-T2-1-001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the granting agency or Singapore Management University.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM HotMobile'14, February 26–27, 2014, Santa Barbara, CA, USA.
Copyright 2014 ACM 978-1-4503-2742-8 ...\$15.00.

indoor location system, it provides limited value since it only provides how many floors have changed when the user moves.

If not for indoor location, what can the barometer sensor be used for? The sensor was originally introduced to provide faster GPS location fixes by providing an initial altitude estimate for the GPS location calculations. However, we have run experiments with comparable devices, ones that have the barometer and ones that do not, and we see no improvement in GPS lock speed or accuracy (average speed of 39.0s for a cold start GPS lock for a Nexus 4 with a barometer compared to 20.3s for a Lumia 720 without). We therefore encourage the mobile community to find solutions, that overcome the limitations we encountered in using barometer sensor for indoor location, as well as identify other compelling use cases for this new sensor.

2. MOTIVATION

The use of a barometer for identifying height (or floor) change information is motivated by the fundamental property that atmospheric pressure drops with an increase in altitude. It is well-known that this relation between barometric pressure and altitude is affected not just by the temperature, but also by various environmental phenomena, such as weather patterns and humidity. For example, during hurricanes or temperate depressions, the pressure readings will obviously drop. Our focus is *not* on studying or investigating this fundamental relationship, but on *ascertaining the properties of pressure variations in indoor buildings, and on the interplay between such variations and the measurement accuracy of the phone-embedded barometer sensor.*

We expect the use of the barometer for indoor height/floor estimation to be a non-obvious exercise, principally because indoor environments have several distinct artifacts that we do not observe outdoors. In particular, we can envision the following artifacts:

- Buildings are often pressurized and climate controlled (more specifically, in the context of Singapore, *air-conditioned*). As a consequence, we expect the humidity and temperature indoors to be quite distinct from that outdoors. Also, the pressure-gradient indoors may not follow a simple relationship, as the building pressure on different floors may be regulated by different air-conditioning units or controllers.
- The floor heights of buildings are typically in the range 2.5-6.0 meters, which may be well within the range of measurement error of the smartphone-embedded barometer sensor. So, while a 20 meter variation in height estimation may be inconsequential for GPS, it may translate into an error of 8 or more floors inside a building!
- The floor heights of buildings are not only different, but are non-uniform (even within the same building). In particular, our empirical studies showed that the heights of lower-level floors (notably the entrance lobby) are often larger, and even the heights of otherwise homogeneous floors (e.g. the 4th and 5th floors of a campus building) show unexpected construction-specific artefacts.

3. DEVICE-SPECIFIC CHARACTERISTICS

We used three different devices for recording our data; 1) a Samsung Galaxy S III smartphone (*S3*), 2) a Google Galaxy Nexus smartphone (*NX*), and 3) an external USB weather board (*UB*) [3]. For the Android platform devices (*S3* and *NX*), we wrote a small application that queried the Android API for the barometer values (at 15 Hz for *NX* and 25 Hz for *S3*). For the weather board (*UB*), we used a laptop running a terminal program that captured the *UB*'s barometer output at 1 Hz for 120 seconds. In the case of the *UB*




Figure 1: 12 Hour pressure trend on a single floor




Figure 2: Temporal Variation in height as computed from pressure

device, we collected at least 120 samples (or 2 minutes of data) for every result. Unless otherwise mentioned, every result for the Android devices uses individual ‘data points’ that are an average of at least 1,000 consecutive samples.

The *NX* and *UB* devices used Bosch barometer sensor [2] while the *S3* used a STM sensor [1]. The Bosch sensors were rated for an RMS measurement error of 0.5 meters (corresponding to an RMS pressure error of 0.06 hPa), while the STM sensors were rated for an RMS measurement error of 0.65 meters (corresponding to an RMS pressure error of 0.08 hPa).

3.1 Device Impact on Accuracy

Figure 1 shows the change in pressure as reported by four devices (2 *S3*'s and 2 *NX*'s) for a 12-hour duration. All four devices were left stationary on a table, next to one other, in a closed room in our campus building. Figure 2 takes a closer look at this change in pressure on the *NX* device for a shorter duration of 15 minutes. There are several observations we make:

- Pressure does not remain constant for a given floor throughout the day in an HVAC environment. To understand why, we observe the variation in temperature and the corresponding variation in pressure [Figure 3]¹. The figure suggests that this variation in pressure is perhaps a result of the variation in temperature.
- The absolute pressure reported (at any given time) across different device models as well as that reported within the *same* device model are different. This difference is significant and can correspond to the pressure difference across a floor pair. For example, the average pressure difference between the two *NX* devices was 1.2 hPa which is equivalent

¹For this observation we collected data using a Samsung Galaxy S4 device that has a built-in barometer and temperature sensor.




Figure 3: Temporal variation of Pressure and Temperature.

to height difference of ~ 10 meters. However, while the absolute pressure reported is different, the pressure trend (change in pressure) across all devices is similar.

- Even when stationary, the pressure drift on a device can report a 1.4 meter change in height after a short duration of 15 minutes. Thus change in pressure need not always be associated with a floor change.

These observations have the following implications:

- Identifying the floor level based on pressure alone will not work.
- While the short-term (within tens of seconds) measurement errors are reasonably low, the medium-term variations are not. Thus to identify floor change, change in pressure as well as the rate at which pressure changes matter.

In particular, we shall see the practical implications for these observations, both on the accuracy of fingerprinting-based strategies for obtaining the number of floors changed (in Section 6) and on the reliability of pressure-change as a criteria for detecting the onset of vertical activities (in Section 5). We note that while the first issue can be resolved by calibrating the sensor, practical deployment will require additional infrastructure to assist the calibration. In this paper we therefore explore the sensor without the need for any additional setup.

4. TIME, LOCATION, & DEVICE EFFECTS

In this section, we show that even with the measurement drifts reported earlier, the barometer’s results are quite robust even across different types of buildings, different times of the day, and importantly, different devices. We collected data from 7 different buildings, using the 3 devices explained in Section 3. The details of these buildings are shown in Table 1.

4.1 Effect of Time, Loc., & Building Type




Figure 4: Robustness of Absolute Pressure Values

As stated in Section 3.1, the absolute values reported by the measurement devices have inherent noise in them. Hence, comparing

Level	PLA	SIS	SOE	SMU	CIT	HDB	BUG
B1 to 1	0.38	0.77	0.84		0.72		
1 to 2	0.49	0.53	0.49		0.60	0.36	0.67
2 to 3	0.49	0.49	0.52	0.89	0.60	0.32	0.67
3 to 4	0.49	0.43	0.37	0.72	0.63	0.28	0.70
4 to 5	0.47	0.47	0.47	0.61	0.65	0.29	0.72
5 to 6	0.50			0.76	0.62	0.30	
6 to 7	0.50					0.32	

Empty values indicates that that building did not have those levels. For *SMU* there is no level 2

Table 2: Pressure Diff. Across Levels and Buildings

the absolute values directly would result in quite a high error. In addition, as shown in Figure 4 (where we took three pressure readings on three different levels of *SIS* (levels 3, 4, and 5) at 3 different times across 3 different days), time of day and day of the week effects can also significantly change the absolute pressure values. However, what if we compared the relative values between different levels in a building? Would the difference between the absolute values help to reduce the error caused by measurement noise and time/day effects?




Figure 5: Robustness of Pressure Diff. Across Building Types

Figure 5 shows the results of this experiment. For this experiment, we computed the average pressure difference between consecutive floor pairs for multiple building types across 3 different days and 3 different times. The results showed that the maximum error in the pressure differences (for the same building and levels across the multiple days and times) was at most 0.2 hPa. This translates to an error of about 1.6 meters. Hence, even with the various sources of error, as long as the distance between floors exceeded 1.6 meters, the barometer could still be a useful sensor to detect floor changes and activities related to that.

Figure 5 also demonstrates that the type of building, the time of day, and the day of the week does not really affect the usefulness of the barometer. However, it was most effective if the inter-floor height of the building was at least 0.2 hPa or 1.6 meters. To understand if this requirement was a restriction in practice, we computed the inter-floor pressure differences for our entire set. This is shown in Table 2.

From the table, we see that most of the inter-level differences are much greater than 0.2 hPa. However, there are some cases (such as all the levels in *HDB*), where the pressure difference is quite close to the required minimum level. In these cases, the probability of having erroneous results (e.g., we predict a floor change that never happened) is higher. Indeed, as shown later, our floor change and vertical activity prediction accuracy is lower than 100% in part because of pressure differences that are close to the 0.2 hPa minimum requirement.

4.2 Pressure Difference is Device Independent

So far we’ve seen that the pressure difference between different levels in a building are relatively independent of time and location. However, for a practical solution this property needs to hold across

Name	Code	Purpose	Size (sq.m)	Floors	Location	Type
University Building	SIS	Academic Building	9,064	6	Singapore	Short, Narrow
University Admin Building	SMU	Office Building	52,637	13	Singapore	Tall, Narrow
University Building	SOE	Academic Building	17,250	5	Singapore	Short, Narrow
Residential Housing	HDB	Apartment Building	21,375	18	Singapore	Tall, Open Air
Plaza Singapura	PLA	Shopping Mall	67,500	9	Singapore	Tall, Wide
Bugis Junction	BUG	Shopping Mall	38,907	5	Singapore	Short, Wide
City Square	CIT	Shopping Mall	32,516	6	Singapore	Narrow

All the buildings, except for *HDB*, were air conditioned and not exposed to the outside air pressure.

Table 1: Buildings Used for Our Experiments




Figure 6: Absolute Pressure versus Pressure Difference

the different measurement devices as well. Figure 6 plots the absolute pressure reported by the *NX* and *S3* devices, as we moved several times (up and down) between a floor pair, as well as the pressure difference between the floors. Clearly, while the absolute pressure reported by the two devices differ by a magnitude of around 2hPa the pressure difference are relatively in-sync.

Overall, our measurements show that pressure differences are quite robust even across difference measurement devices, different types of buildings, different times of the day, different days of the week, and different locations. The main requirement is that the pressure difference between the two floors needs to be at least 0.2 hPa (or 1.6 meters). In the next two sections, we show how to use these pressure differences to build two extremely accurate (close to 100% accuracy) and robust services — namely, a floor change predictor that can also determine the exact number of floors changed, and a predictor that can determine whether an escalator, elevator, or stairs was used to change floors.

5. DETECTING MODE OF TRANSPORT

We hypothesise that the different mode of transports differ principally in two parameters: a) rate of change of height (or equivalently, *rate of change of pressure* P), given by $\frac{\Delta P}{\Delta T}$, and b) total time duration T needed to move between floors.

However, the mappings between $\frac{\Delta P}{\Delta T}$ and the modes of transport is not straightforward. For example, for elevators, a) the rate of pressure change (ΔP) varies depending on the 'number of floors' changed (the period for which an elevator accelerates and decelerates depends on the number of floors changed), and b) the average rate of change of pressure, in most buildings, is higher while ascending than descending, suggesting that elevators travel faster while ascending. There are also similarly confounding effects for escalators and stairs usage.

To navigate these effects in a tractable way, we experimented with four different classification strategies, each of which uses a different combination of 3 distinct features: i) ΔP (the pressure change between the final and initial floors), ii) $\frac{\Delta P}{\Delta T}$ and iii) the transition time T . We used the J48 classifier in Weka, with 10-fold cross validation, to compute the accuracy of our various classifiers.

Mode of Transport	dp	dp,T	(dp/dt)	(dp/dt),T
Escalator/Elevator/Stairs	69	99.67	98.3	100
Escalator	79	99	100	100
Elevator	28	100	95	100
Stairs	100	100	100	100

For each mode of transport, we collected 100 samples with 50 samples going up and 50 going down. The samples were collected in 6 buildings by 2 different people using 3 devices.

Table 3: % accuracy for detecting vertical mode of transport

In Table 3, we present the effectiveness of the barometer at detecting the mode of transport used to change floor levels. When we use both the pressure difference, the time difference, and the total time taken ($(dp/dt)/T$), we achieve 100% accuracy at detecting the mode of transport. However, this is under the assumption that the time taken to change floors using the elevator, escalator, and stairs is different from each other. It also assumes that the floors are at least 1.6 meters apart. This translates to about 5-7 seconds of movement on a typical escalator, stairs, and often less for an elevator.

6. DETECTING FLOOR CHANGES

In this section, we describe the effectiveness of the barometer in determining that the measurement device has changed floors. We show this in two steps; 1) how effective is the barometer at just detecting that a floor change has happened, and 2) how effective is the barometer at detecting the exact number of floors changed?

6.1 Detecting That The User Has Changed Floors

To determine the accuracy of the barometer at determining that a device has changed floors, we collected 220 samples. This data included 110 samples that involved the user changing floors and the rest when the user did not. The floor transitioning data also had a 50-50 split between going up and coming down. In both cases we recorded the change in pressure for the duration of the activity.

We then used Weka to perform a 10-fold cross validation and obtained a floor change accuracy, using just the difference in pressure values (dp), of 99.54%. I.e., we can use the barometer to very accurately determine if the measurement device has changed floors. However, this again assumes that the floor change is at least 1.6 meters.

6.2 Identifying the Exact Floor

Up to this point, the barometer can be used to very accurately determine that a floor change has occurred. However, can we also determine the number of floors that have been changed? For example, can we determine that the measurement device has gone up 2 floors, or gone down 3 floors? This information could be incredibly useful in any kind of indoor localization tool.

6.2.1 Without Any Additional Information

We investigate the accuracy levels achievable without any additional information (beyond the pressure and time differences ob-




Figure 7: Accuracy At Determining No. of Floors Changed

tained when changing floors). To do this, we collected data from every possible floor transitions of 1-5 levels (up and down) from all our buildings, using all the three measurement devices. For each possible floor transition, we collected 11 samples going up and 11 samples going down for a total of $22 * 5 = 110$ samples. We entered all our measurements into Weka and built a suitable predictor that could be used to accurately determine the exact number of floors changed.

Figure 7 shows the accuracy of determining the exact number of floors changed using just pressure and time differences. We observe that the accuracy of determining just a single floor change is very high (close to 100% using the $(dp/dt)/T$ predictor). However, the accuracy drops significantly for two and three floor changes. This is because the time taken to change these number of levels can overlap (for example, an elevator can go up three floors just as fast as another elevator can go up two floors).

However, if we are willing to accept up to a one floor error, we can achieve a 100% accuracy using the $(dp/dt)/T (+1/-1)$ predictor. This predictor uses the direction of motion to accurately predict either the exact number of floors changed or the number of floors changed + 1. For example, if you went up 3 floors, it will return either 3 or 4 floors up (and no other value). If you are went down 2 floors, it will return either 2 or 3 floors down (and no other values).

6.2.2 With Basic Fingerprinting Information




Figure 8: No. of Floors Changed Accuracy with Fingerprinting

Finally, we investigate the accuracy of determining the number of floors changed if we are able to complement our data with a little bit of fingerprint information. In particular, what if we knew the height of the various floors or in other terms the relative pressure difference between the various floors of the building the device is currently in. If this information was available, we would not need to measure the pressure values while changing floors and then using that to compute the difference in pressure and difference in time needed by our previous predictor. Instead, after the floor change was done, we could just measure the absolute pressure seen at the new floor and immediately compute the number of floors changed (using the difference between the current floor and the level 1 pressure as an index into a pre-computed “relative pressure map” table).

Figure 8 shows the accuracy possible under these assumptions. Each floor pair in the figure was the average of at least 1,000 sam-

Device	S3		Nexus	
Frequency (Hz) Acc/Bar	100/25	15/5	125/15	15/5
Measured Power Consumption (mW) After 10 mins of Use				
Base	28.3		65.06	
Accelerometer	719.05	594.97	748.67	662.29
Barometer	582.37	574.72	662.62	621.66
% Improvement	23.74	3.48	12.99	6.54

Table 4: Power Consumed by the Accelerometer & Barometer

ples. We observe that we can obtain 96.5% under these condition with just one measurement taken on the new floor by the device. As mentioned earlier in Section 3.1, the measurement device noise can be quite high and this affects our accuracy. To achieve better accuracy, we should use more samples to smooth out the noise. If we are willing to take 30 samples, for example, the accuracy becomes almost 100% (99.94%).

7. BETTER THAN AN ACCELEROMETER?

Prior work has used the accelerometer to perform some of the functions described earlier [13]. We therefore compare the barometer with the accelerometer for these functions, as well as the power consumed, to evaluate the true utility of the barometer. As before, we use the S3 and set the 3-axis linear accelerometer to a sampling frequency of 15 Hz. We performed the experiments under two different conditions. One, when the phone was kept flat in the right hand for the entire duration of the recording, and the second when the phone was perturbed for the entire duration of the recording. There were two types of perturbations. In the first type, the phone was being used to play a game. In the second type, the user was attending to a phone call. The perturbed data was collected in S/S from six stairs, six escalators and six elevators for the floor combinations of 1, 2 and 4 floor changes (both up and down). The un-perturbed data contained 100 different sets of measurements for using the stairs (across all the buildings in our dataset). Similarly, we had 100 different measurements for the escalators and elevators in our un-perturbed data. When taking the stairs, escalators or elevators, the sensor values were recorded just before stepping onto the transport and stopped just after stepping off the transport. The un-perturbed data was labeled and used to create a training set for the perturbed data using the J48 classifier in Weka.

Figure 9 shows the results of this study. Figure 9a) shows the results for detecting the mode of vertical transport (elevators, escalators, stairs) while the measurement device was left stationary (Normal), used to play a game, and used to make a call. Figure 9b) shows the results for detecting that the measurement device has changed a floor under the same 3 conditions. We observe that when the measurement device was left in a normal position, the accuracy of both solutions was almost identical and quite high. However, when the measurement device was perturbed, the accuracy of the accelerometer-based solution dropped significantly while the accuracy of the barometer-based solution remained almost constant (and close to 100% accurate).

Finally, the power consumption of each sensor measured over a 10 minute time period using the Monsoon power measurement device showed the barometer (on both devices) consumes between 3 to 23% less power than the accelerometer (Table 4). This suggests that the barometer, in practice, is more energy efficient than the accelerometer — but only just.

8. RELATED WORK

Barometers have recently started appearing on smartphones and tablets. Given the novelty of these sensors, we have found little prior work that has used them for indoor location.




Figure 9: Robustness Of The Accelerometer Versus The Barometer

Brush [4] unsuccessfully attempted to use a barometer to help users find where they parked their car. Lester [8] focused on activity recognition, including going up or down stairs and elevators. While their data collection included barometric pressure, it was not used to infer where a user is.

Varshavsky [12] used a GSM fingerprinting based system to infer the current floor of a user, but with lower accuracy than our system and with higher fingerprinting overhead. Ojeda [10] deployed a dead reckoning system to capture floor level transitions. However, that system is limited to stairs and cannot capture floor changes via an elevator or escalator.

Johannsson [6] used vertical accelerometer information over time to determine the number of floors traveled. While the paper reports high accuracy in classifying floor level transition it is unclear if the results hold for multiple buildings. Similarly Ye [13] used a phone’s accelerometer to record the time taken for traveling across different floors via an elevator and the step count in the case of stairs. The system was shown to achieve a high accuracy, but as we observe, the accelerometer sensor is susceptible to any sort of perturbation which can result in lower accuracies.

A large body of literature proposes indoor location techniques which can be used to identify the physical location of users with varying levels of accuracy. This could in theory be translated to detection of floor changes. Some techniques rely on custom hardware or radios [5, 7, 9, 11]. Custom radios can also offer TOA (time of arrival), TDOA (time difference of arrival) or AOA (angle of arrival) information for location. These approaches have hardware adoption challenges.

9. CONCLUSIONS

We have provided what we believe to be the first exhaustive study on the properties of mobile-embedded barometers across a number of buildings with heterogeneous characteristics. Our results show that while absolute pressure readings have significant time-of-day variations, the difference in pressure across different floor pairs is remarkably consistent and steady for any given building. As a consequence, we are able to use pressure difference as a useful fingerprint to detect the exact number of floors changed with almost 100% accuracy. Additionally, pressure-based features (such as the change in pressure) enable us to classify vertical activities (such as taking escalators, stairs or elevators) with high accuracy. The barometer is highly robust to changes in the phone’s on-body placement and orientation, making it a significantly more robust sensor than the accelerometer for real-life vertical activity detection.

Unfortunately, we also conclude that it is difficult to use the barometer to determine the actual floor that a user is on. Knowing how many floors a user has changed and what modality was used is not a particularly useful piece of context. While the promise of

the barometer was that it would aid GPS location, in practice we find that it does not help there either. In summary, we advise our colleagues in the mobile computing community to be aware of the limitations of this sensor when being considered as part of their system.

10. REFERENCES

- [1] LPS331AP MEMS Pressure sensor. <http://www.st.com/st-web-ui/static/active/en/resource/technical/document/datasheet/DM00036196.pdf>.
- [2] MEMS Pressure Sensors Data sheets. http://www.bosch-sensortec.com/en/homepage/products_3/environmental_sensors_1/bmp180_1/bmp180.
- [3] USB Weather Board. <https://www.sparkfun.com/products/8311?>
- [4] Brush, A. B., Karlson, A. K., Scott, J., Sarin, R., Jacobs, A., Bond, B., Murillo, O., Hunt, G., Sinclair, M., Hammil, K., and Levi, S. User experiences with activity-based navigation on mobile devices. *MobileHCI*, 2010.
- [5] Chang, H., Tian, J., Lai, T., Chu, H., and Huang, P. Spinning beacons for precise indoor localization. *SenSys*, 2008.
- [6] Johannsson, H., Kaess, M., Fallon, M., and Leonard, J. Temporally scalable visual SLAM using a reduced pose graph. *RSS Workshop on Long-term Operation of Autonomous Robotic Systems in Changing Environments*, 2012.
- [7] Kusy, B., Maroti, M., Balogh, G., Volgyesi, P., Sallai, J., Nadas, A., Ledeczi, A., and Meertens, L. Node density independent localization. *IPSN/SPOTS*, 2006.
- [8] Lester, J., Choudhury, T., and Borriello, G. A practical approach to recognizing physical activities. *Pervasive*, 2006.
- [9] Muthukrishnan, K., Dulman, S., and Langendoen, K. Towards a rapidly deployable positioning system for emergency responders. *International Conference on Ubiquitous, Positioning, Indoor Navigation and Location-Based Service*, 2010.
- [10] Ojeda, L. and Borenstein, J. Personal dead-reckoning system for GPS-denied environments. *Safety, Security and Rescue Robotics, 2007. SSRR 2007. IEEE International Workshop on*, pages 1–6, sept. 2007.
- [11] Priyantha, N. B., Chakraborty, A., and Balakrishnan, H. The Cricket Location-Support system. *MobiCom*, 2000.
- [12] Varshavsky, A., LaMarca, A., Hightower, J., and de Lara, E. The SkyLoc floor localization system. *Percom*, 2007.
- [13] Ye, H., Gu, T., Zhu, X., Xu, J., Tao, X., Lu, J., and Jin, N. FTrack: Infrastructure-free floor localization via mobile phone sensing. *Percom*, 2012.