Polaris: Getting Accurate Indoor Orientations for Mobile Devices Using Ubiquitous Visual Patterns on Ceilings

[1]Zheng Sun, [1]Aveek Purohit, [2]Shijia Pan, [1]Frank Mokaya, [3]Raja Bose, and [1]Pei Zhang
[1]Electrical and Computer Engineering, Carnegie Mellon University
[2]Computer Science and Technology, University of Science and Technology of China
[3]Nokia Research Center, Palo Alto
{zhengs, apurohit, peizhang, fmokaya}@cmu.edu, {shijia.pan}@gmail.com, {raja.bose}@nokia.com

ABSTRACT

Ubiquitous computing applications commonly use digital compass sensors to obtain orientation of a device relative to the magnetic north of the earth. However, these compass readings are always prone to significant errors in indoor environments due to presence of metallic objects in close proximity. Such errors can adversely affect the performance and quality of user experience of the applications utilizing digital compass sensors.

In this paper, we propose Polaris, a novel approach to provide reliable orientation information for mobile devices in indoor environments. Polaris achieves this by aggregating pictures of the ceiling of an indoor environment and applies computer vision based pattern matching techniques to utilize them as orientation references for correcting digital compass readings. To show the feasibility of the Polaris system, we implemented the Polaris system on mobile devices, and field tested the system in multiple office buildings. Our results show that Polaris achieves 4.5° average orientation accuracy, which is about 3.5 times better than what can be achieved through sole use of raw digital compass readings.

Categories and Subject Descriptors

C.3 [**Special-purpose and application-based systems**]: Signal processing systems

General Terms

Algorithms, Design, Experimentation

Keywords

Orientation, digital compass, ceiling pictures

1. INTRODUCTION

Digital compass equipped mobile devices have become fairly common now and are playing increasingly important roles in ubiquitous computing application domains, such as localization [1], activity recognition [4], photographing [12], and gaming [11].

A digital compass sensor provides the orientation of the device relative to the magnetic north of the earth. However, when used within indoor environments they suffer from significant errors, due to the existence of metallic objects in close proximity.

To compensate for compass errors, previous work has explored different approaches, such as filtering [1], averaging [14], or integrating compass readings with gyroscopes using Kalman Filters [8]. However, these approaches usually assume that the magnetic interference is low, and that a 'correct' initial compass reading can be obtained, otherwise compass errors will accumulate quickly over time. Hence, these approaches perform poorly in conditions where there is high magnetic interference.

In this paper, we propose Polaris, a system that provides reliable orientation information for mobile devices within indoor environments. Our approach is based on an observation that indoor environments, such as classrooms, offices, and supermarkets, are highly likely to have regular rectangular or square peripheries visible on their ceilings. These objects include ceiling beams, panels, tube lamps, gas pipes, electricity wires, and ventilation fans. The edges of these objects are usually straight, and are parallel or perpendicular to the orientation of the buildings. Figure 1 shows some examples of such patterns. Polaris uses these straight lines as orientation references for mobile devices to correct errors in magnetic compass measurements. Therefore, in indoor environments having such ceiling patterns, when a user wants to orient her mobile device, she only needs to take a picture of the ceiling, and the orientation can then be inferred by incorporating both the visual patterns on the ceiling and the raw magnetic compass measurements.

As a preliminary effort, we implemented the Polaris system using iPhones and HTC phones, and tested the system in multiple office buildings. Our experimental results show that Polaris achieves 4.5° average orientation accuracy, which is about 3.5 times better than what can be achieved using raw compass readings.

The key contributions of this paper are as follows:

1. We propose the novel idea of using ceiling features such as parallel and perpendicular straight line edges of ceiling-mounted objects, to correct magnetic compass measurements for mobile devices.

2. We implemented the Polaris system using commercial off-the-shelf mobile phones, and show that Polaris achieves significantly better accuracy as compared to raw compass readings.

The rest of this paper is organized as follows. Section 2 motivates our work and Section 3 gives a system overview of Polaris. We describe the algorithms in Section 4 and provide evaluation results in Section 5. Related work is shown in Section 6. Finally, Section 7 describes future work and Section 8 concludes the paper.

HotMobile'12 February 28–29, 2012, San Diego, CA, USA.

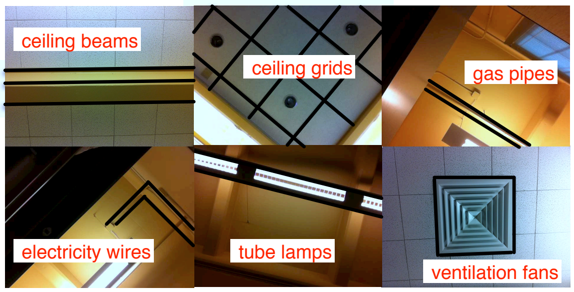

Figure 1: Examples of objects with parallel and perpendicular lines on ceilings in indoor environments that can be leveraged by Polaris. Their peripheries are emphasized using bold black lines.

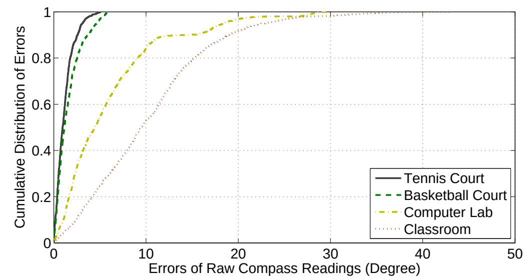

Figure 2: The cumulative distribution of compass errors under the four cases. The errors in the two indoor cases (classroom and computer lab) are much bigger than those in the two outdoor cases (basketball and tennis courts) due to indoor magnetic interferences.

2. COMPASS ERRORS IN INDOOR ENVIRONMENTS

Even though digital compasses are commonly used in ubiquitous computing applications, the limitations imposed on their accuracy due to the existence of localized magnetic interference leads to poor quality of performance and user experience. Previous work has pointed out that digital compass errors up to 23° can be observed in indoor environments even given a correctly calibrated compass, especially in situations where the compass sensor is placed in close proximity to electromagnetic sources such as computers, metal cubicles, and high-voltage power lines [10]. Gusenbauer et al. [6] have pointed out that the primary source of error accumulation in their self-contained indoor positioning system was the heading error, which was caused by the exclusive use of an electronic magnetic compass for heading determination. Similarly, Youssef et al. [14] also talked about their experiences where noisy compass readings led to rapid accumulation of errors in experimental results over time.

In order to quantitatively investigate the effect of magnetic interference to digital compasses indoors, we collected raw compass readings in two different types of indoor locations namely, classrooms and computer labs. We also collected readings in two different types of outdoor locations namely, basketball courts and tennis courts. In each case, we placed an HTC G7 Desire smart phone at four different locations, each oriented to four directions that were perpendicular to each other. For each location and each direction, 50 samples were collected, resulting in a total of 200 samples per location. To measure ground truths for the two indoor cases, we collected compass readings outside of the building, facing towards the same directions as we did inside the building.

Figure 2 shows the cumulative distribution of the compass errors. The median error in the two outdoor cases was 1.6° for the basketball courts and 1.2° for tennis courts, whereas median error in the two indoor cases was 16.8° for the classrooms and 7.5° for the computer labs. These results indicate that the compass errors in indoor environments are $5\sim15$ times higher than outdoors. We also noticed that, the desks on which the phone was placed in classrooms had iron supports and legs, and the air conditioners in the rooms were a powerful source of magnetic fields. Hence, the classroom compass readings generated much higher errors in measurement of orientation. Furthermore, the maximal error observed in the classroom case was 42.9°.

These results indicate that digital compasses are significantly affected by magnetic fields indoors and are unable to unilaterally provide reliable orientation information.

3. SYSTEM OVERVIEW

To address the inaccuracy of compass measurements, Polaris uses ceiling patterns, i.e. parallel and perpendicular straight lines visible on the ceiling, as orientation references for mobile devices. This is based on two fundamental assumptions. The first assumption is that parallel and perpendicular straight lines commonly exist on ceilings. This assumption is intuitively reasonable in that when objects, such as grid panels of *dropped ceilings*, tube lamps or electricity wires, are hung on the ceilings, they are usually oriented to be parallel or perpendicular to the ceiling edges for aesthetic considerations such as visual neatness and consistency.

The second assumption is that if such straight line patterns exist on the ceiling, they should be parallel (or perpendicular) to each other over the entire building, indicating that the buildings are rectangular in shape. This assumption may not be true for some particular buildings, such as the Pentagon. However, architects A. F. Bemis and M. J. T. Kruger have carried out surveys of buildings and found out the majority (83% and 98% in their respective surveys) of modern buildings are predominantly rectangular [13].

Based on these two key assumptions, we design the Polaris system as described below. The system architecture is shown in Figure 3. Polaris leverages crowdsourcing to collect ceiling images and raw compass readings associated with the location and orientation from which each ceiling image was captured. When people who choose to participate in the crowdsourcing activity are in their indoor environments, they will be prompted to contribute ceiling pictures to Polaris. This is aided by our observation that unobstructed pictures of the ceiling are easily obtainable using the front camera on mobile phones during the user's normal interaction with the device. These pictures, along with raw compass readings recorded when the pictures are taken, are then paired, time stamped and sent to a back-end server where they are aggregated.

The server collects such *ceiling picture - compass reading* pairs from multiple mobile devices in the building over a period of time, and creates a mapping relation between general directions of ceiling patterns (i.e. parallel or perpendicular straight lines) and the magnetic north. Since the raw compass readings are taken throughout the entire building at different time, the effect of any localized magnetic interference is expected to be minimized.

After creating the general mapping relation for a building, a person can orient her mobile phone by simply taking a picture of the ceiling. The ceiling patterns contained in the picture will then be processed by incorporating the mapping relation, and finally the phone orientation can be determined.

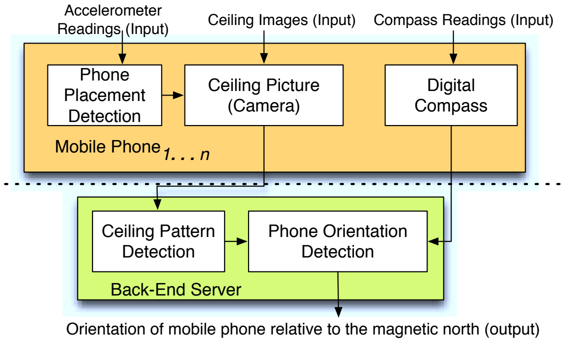

Figure 3: The system architecture of the Polaris system.

The Polaris system has the following major components:

1. Aggregating ceiling pictures through crowdsourcing. Polaris aggregates data from different users in the same building at different times. For enhanced usability, it is desirable that the aggregation process be automatic with minimal user participation. Moreover, the pose of the mobile phone in the user's hand can distort the observed orientation of ceiling patterns. Polaris addresses this by automatically detecting the suitable stationary and horizontal phone placement when the user wishes to contribute pictures.

2. Extracting effective ceiling patterns. Ceiling pictures usually contain complex and diverse objects, such as ceiling beams, electricity wires, fans, or fluorescent lamps with different shapes. Variation in distances to light sources can also result in brightness changes in the captured pictures, which make ceiling pattern detections challenging. Polaris addresses these issues using histogram equalization and multiple edge detection techniques to robustly detect straight lines.

3. Accurate estimation of mobile device orientation. Due to the existence of *both* parallel and perpendicular lines on the ceiling, Polaris cannot differentiate ceiling patterns that are 90° rotational-symmetric, thereby causing directional ambiguities. To get correct directions, Polaris leverages the raw compass readings on the mobile phones as references, and eliminates these ambiguities.

The following sections will describe these techniques in details.

4. ALGORITHM DESIGN

Polaris leverages 'crowdsourcing' to obtain information regarding ceiling patterns and their corresponding directions. People working in the same office building are invited to contribute ceiling pictures to the system.

4.1 Aggregating Ceiling Pictures

To minimize the involvement from users, ceiling pictures have to be taken as automatically as possible, without sacrificing picture quality. Toward this end, one important thing to ensure is that the phones are held horizontally by the users, facing up to the ceiling, such that the orientations of the straight lines would not be affected by possible perspective-distortions of pictures.

To detect horizontal and stationary placements, Polaris keeps tracking 3D accelerations a_X, a_Y, and a_Z using accelerometers in the mobile phones to detect the horizontal phone placement. Specifically, if

$$\begin{cases} \text{Abs}(\text{Mean}(a_X)) \leq \epsilon_{mean}, \text{ and } \text{Var}(a_X) \leq \epsilon_{var} \\ \text{Abs}(\text{Mean}(a_Y)) \leq \epsilon_{mean}, \text{ and } \text{Var}(a_Y) \leq \epsilon_{var} \\ \text{Abs}(\text{Mean}(a_Z) - g) \leq \epsilon_{mean}, \text{ and } \text{Var}(a_Z) \leq \epsilon_{var} \end{cases}, \quad (1)$$

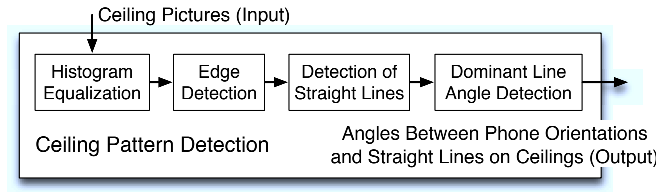

Figure 4: The processing pipeline of the extraction of ceiling patterns.

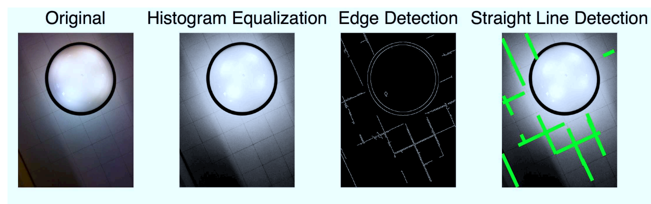

Figure 5: Pattern extraction of a ceiling picture that contains a fluorescent lamp. The Hough transform is capable of avoiding the round edges of the lamp and correctly detecting straight lines of the ceiling grid. After the Hough transform, the detected straight lines are shown using bold bright lines on the rightmost plot.

holds for 2 seconds, where g represents the gravity of Earth, Polaris takes one picture of the ceiling, and transmits it back to the server. The ϵ_{mean} and ϵ_{var} are parameters used to tolerate minor arm vibrations of the user. In experiments, we empirically set them to 0.03g and 0.01g, respectively.

We note that, image processing techniques could also be used to transform a skewed ceiling view to a horizontal one given the accelerometer readings of the phone. However, in the current implementation, we focus on the placement detection approach only.

4.2 Extracting Effective Ceiling Patterns

Ceiling pictures taken by the users usually contain complex and diverse objects, such as ceiling beams, panels, lamps, and wires. These objects can sometimes significantly change the brightness and contrast of the pictures and make the detection of straight lines challenging, especially when a fluorescent lamp is in the picture.

To address this issue, the server first converts the pictures to grayscale images, and then equalizes the histograms of the images to reduce brightness and contrast variations [15]. Then, edges are detected from the pictures using the 'Canny' and the 'Sobel' edge detection algorithms [3]. During our experiments, we found that the 'Canny' algorithm worked well for detecting thinner lines, whereas the 'Sobel' algorithm had much better robustness against pictures that had low brightness. Finally, the standard Hough transform detects straight lines in the pictures [2]. Occasionally, ceiling pictures can contain objects with round or irregular edges, such as lamps. However, we found that the Hough transform was robust enough to avoid such outliers, as illustrated in Figure 5. The processing pipeline of extracting the ceiling patterns is shown in Figure 4.

4.3 Accurate Estimation of Phone Orientations

Finding the Two Directional Axes of A Building. Along with the ceiling pictures, the back-end server also receives raw compass readings from the phones. These compass readings indicate the phone orientation relative to the magnetic north that the compass measures when the picture is being taken by the user, which is denoted as α_P. The Hough transform, as mentioned in Section 4.2, detects straight lines in every ceiling picture, and also provides the orientations of the lines relative to the phone. Assuming that most ceiling lines are parallel or perpendicular, the server tests two hypotheses to find dominant line orientations: First, if more than $a\%$ lines in one image have a dominant angle with $\pm b$ tolerance, the server considers the image to only have parallel lines, and refers this angle as the ceiling pattern orientation relative to the phone, which is denoted as $\alpha_{C(P)}$. Otherwise, the server considers the image to have perpendicular lines and conducts a linear search for two angles that are $\pm 90^\circ$ apart and have the most lines reside on, and consider one of the two angles as $\alpha_{C(P)}$. In the implementation, we empirically set a and b as 90% and 1°, as discussed in Section 5.1.

Through the Cartesian coordinate transforms, a candidate orientation of the building relative to the magnetic north can be derived as

$$\hat{\alpha}_C = \mathrm{Mod}(\alpha_{C(P)} + \alpha_P, 360^\circ), \tag{2}$$

where Mod is the operation that computes the remainder of division.

Through crowdsourcing, Polaris is able to aggregate compass readings and ceiling pictures from all over the building. Finally, Polaris combines all compass readings and determines candidate orientations of the building as

$$\begin{aligned}\alpha_C &= \mathrm{atan2}\left(\frac{\sum_{k=1}^{K}\sin(\hat{\alpha}_{C,k})}{\sum_{k=1}^{K}\cos(\hat{\alpha}_{C,k})}\right)\\ &= \mathrm{atan2}\left(\frac{\sum_{k=1}^{K}\sin(\mathrm{Mod}(\alpha_{C(P),k}+\alpha_{P,k},360^\circ))}{\sum_{k=1}^{K}\cos(\mathrm{Mod}(\alpha_{C(P),k}+\alpha_{P,k},360^\circ))}\right),\end{aligned} \tag{3}$$

where K is the total number of ceiling pictures aggregated in the same building.

Due to the existence of perpendicular lines on the ceiling, Polaris cannot differentiate straight lines in a 90°-rotated ceiling picture. Therefore, the actual orientation of the building could be either α_C, $\alpha_C \pm 90^\circ$, or $\alpha_C \pm 180^\circ$, as shown in Figure 6.

Estimating Phone Orientations Using Polaris. To solve the ambiguity among directional axes, when a user wants to determine her phone's orientation, Polaris uses the raw compass reading in the phone as directional references. First, the user takes a picture of the ceiling. Then, straight lines will be detected using edge detection techniques and the Hough transform, as described in Section 4.2, along with the orientations of the lines relative to the phone, i.e. $\alpha_{C(P)}$. Finally, the phone orientation is derived as

$$\alpha_P = \mathrm{Mod}(\alpha_C - \alpha_{C(P)}, 360^\circ), \tag{4}$$

where α_C is the orientation of the building that is determined through crowdsourcing in Equation 3. Again, the ambiguity problem will arise. For example, when Polaris gives a value as 30°NE, the actual phone orientation could be either 30°NE, 120°SE, 210°SW, or 300°NW.

To eliminate the ambiguities, Polaris refers to the raw compass reading on the phone, and chooses the orientation value that is the closest to the raw compass reading as the final phone orientation estimate. As in the example above, if the digital compass gave 47°NE, Polaris would choose 30°NE as the final orientation estimate. We would like to note that the readings given by the digital compass may be prone to errors caused by magnetic interferences. However, as long as the errors of the compass is less than $\pm 45^\circ$,

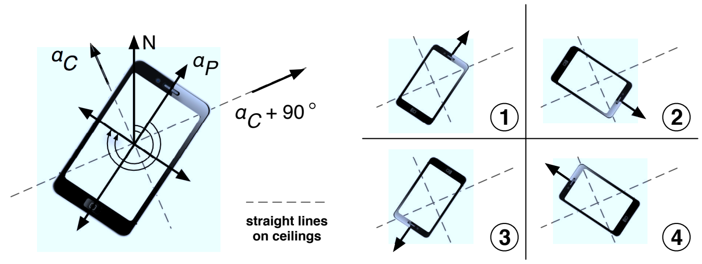

Figure 6: Left: Relations between the orientation of the ceiling patterns α_C and that of the phone α_P, both relative to the magnetic north. Right: The four ambiguous phone orientations that arise due to the existence of perpendicular lines on the ceiling and the phone's incapability of differentiating pictures with 90° rotations.

Polaris is able to choose the correct orientation and solves the ambiguity problem.

5. EVALUATION RESULTS

As a proof of concept, we implemented the Polaris system using iPhones (iOS 4.3.5) and HTC G7 Desires (Android 2.3.2, with built-in AK8973 3-axis electronic compasses), and performed real experiments in several office buildings. The server-end processing was implemented using the Image Processing Toolboxes in MATLAB. The compass readings collected on the iPhones were from iPhone's built-in compass app, and those on the HTC phones were from the Android API, with tilt-compensations by default. In both cases, the compass readings were raw data, without any filtering or adjustments from gyroscopes. The ceiling pictures were taken using the default camera program on the phones, and were transferred to the server through WiFi connections.

5.1 Performance of Straight Lines Detection

To get accurate orientation information, Polaris leverages ceiling patterns, i.e. straight lines on ceilings, as orientation references. We conducted experiments in three different buildings to evaluate the performance of straight line detection on ceilings and the accurate calculation of the angles between the lines relative to the phone.

The ceiling pictures was taken using an iPhone in the Main Building and Building 19 of CMU Silicon Valley campus in Mountain View, California, and using an HTC G7 Desire phone in the Computer Lab Building in the University of Science and Technology of China in Hefei, China. In the Computer Lab Building and Building 19, we did the experiment in 10 different rooms, each having pictures taken at three and five different locations, respectively. In the Main Building, since the building features a huge public cubicle area, we took pictures at 16 different locations in the cubicle area.

We examined the detected straight lines vs. real patterns marked using human eyes in each ceiling picture. Furthermore, if straight lines are detected, the difference between the detected orientation of these lines and the ground truth is manually computed using a protractor. If a detected orientation is within $\pm 1^\circ$ of the manual measurement, we consider it as a successful detection, and further quantize its orientation difference. Table 1 shows the results. We found that in average the Polaris system could correctly detect straight lines in most pictures, with a 88.5% successful detection rate. When the brightness was moderate and the contrast was high, the detection rate could achieve 100%.

Table 1: Experiment results of detecting straight lines on ceilings and orientation estimations relative to the phones.

Locations	Straight Line Detection Rates	Averaged Errors
Computer Lab Building	25/30 (83.3%)	0.5°
Main Building	16/16 (100.0%)	0.4°
Building 19	44/50 (88%)	0.2°

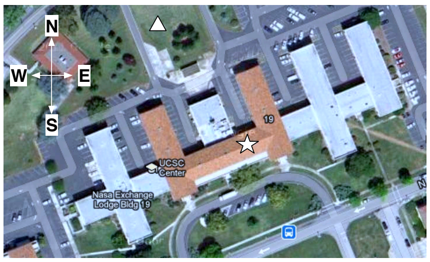

Figure 7: The satellite view of Building 19, the building in which our experiment was conducted. The star shows Building 19, and the triangle shows the lawn outside the building on which we took ground truth orientations for the experiment.

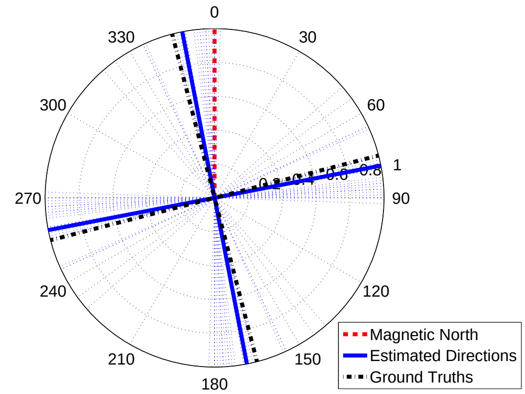

Figure 8: The directional axes estimated using the collected building orientations through crowdsourcing (shown in dotted light blue lines) of Building 19. They are compared with the ground truths. The angular error of the estimated directional axes is about 3.5°.

5.2 Performance of Orientation Determination

Based on the pattern detection techniques, we evaluated the performance of orientation determination of Polaris. We used the 50 pictures taken in Building 19 to build the general mapping relation between orientations of straight lines and the magnetic north. Since the pictures were taken at different rooms facing different directions, we used them to simulate the process of picture aggregations through crowdsourcing.

To generate ground truths, we measured the actual orientation of Building 19 (as shown in Figure 7) relative to the magnetic north on a lawn outside of the building, to minimize effects from magnetic interference. Figure 8 shows the candidate directional axes of Building 19 estimated using the 50 ceiling pictures. Since both the positive and negative differences exist between the estimated directions and the ground truths, the errors after averaging the candidates tend to diminish. The ultimate error between the estimated directional axes and the ground truths was only 3.5°. This result indicates that using ceiling pictures and compass readings aggregated at different places in a building to estimate the orientation of the building is possible and accurate.

To evaluate the performance of orientation determination, a student took an iPhone and walked inside Building 19, following a square-shaped route. Each edge of the route contained ten steps, with about 0.7m intervals, toward the same direction. After taking each step, the student took a ceiling picture, and recorded the compass readings using the iPhone. After the experiment, the orientations of the 40 steps on the route were estimated using Polaris. Figure 9 shows the estimated orientations and the raw compass readings and Figure 10 shows the cumulative distributions of the orientation errors. The median error of using the raw compass readings was 15.5°, whereas that of using Polaris was only 4.5°, about a 3.5X improvement. Furthermore, the standard deviations and the 95th percentile errors were also reduced from 9.24° and 35.5° to 1.21°, and 5.5°, respectively, about 7.6X and 6.5X better. These results indicate that Polaris significantly improves the accuracy of orientation determination.

We would like to note that, of the 4.5° median error, 3.5° is a static offset that is contributed from errors in detection of building orientation. If the accuracy of determining building orientations can be increased, the accuracy of Polaris is expected to improve, as we will discuss in Section 7.

6. RELATED WORK

Digital compasses are predominantly used in ubiquitous computing research for determining device orientation. Their applications range from localization [1, 10, 14], activity recognition [4], and computer-human interaction [5].

Along with the widespread use, however, digital compasses have been widely known to provide erroneous readings in indoor environments due to the existence of metallic objects and magnetic fields [1, 7, 14]. There are multiple techniques that have been proposed to compensate for compass errors. The first approach is through the use of additional sensors, especially the techniques that combine compasses with gyroscopes through the Kalman filtering [10]. However, since the gyroscope does not measure the absolute orientation, it rely on the initial value of digital compasses, which may suffer from systematic offsets. Although some authors come up with the idea to combine even more sensors, such as the GPS or vision-based recognition software, to calibrate with known landmarks, the overhead and redundancy is still concerning [10].

The second approach is through averaging of sensor readings. For example, in GAC [14], the authors average multiple compass readings to estimate device orientation over time. However, if the location of the compass does not change over time, such as a user holding her mobile phone while sitting in her cubicle, the magnetic interference from the metallic cubicle cannot be eliminated by only using averaging. Moreover, when magnetic anomaly is detected, such as in the iPhone, existing methods often require frequent recalibration if the user is mobile.

By contrast, Polaris avoids these problems associated with indoor environments by using existing and invariant patterns on the celling. Since ceiling patterns are universal and unrelated to magnetic fields, Polaris can provide accurate orientations for mobile devices even under severe magnetic interferences.

7. DISCUSSIONS AND FUTURE WORK

Polaris leverages crowdsourcing to aggregate ceiling pictures as crowdsourcing, as evaluated in Section 5.2. In future work, we will leverage online map services to eliminate the need of crowdsourcing. As the accuracy through online maps is 0.1°, we expect to improve the accuracy of Polaris to be within 1°.

Due to the use of image capturing and simple line detection techniques in Polaris, there is a slight increase in power consumption as compared to simply using raw compass readings from the mobile device. We intend to do power profiling as part of our future work to determine the impact Polaris has on power consumption across different mobile devices. Moreover, as battery capabilities and power management on these mobile devices improve, this concern is expected to diminish in the future.

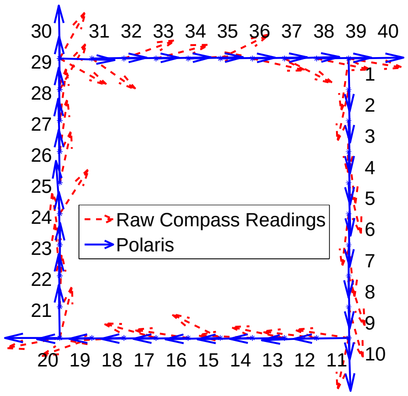

Figure 9: The estimated orientations of footsteps when the student walked following a square-shaped route in Building 19. The raw compass readings that were taken at each step are also shown as comparisons. It can be seen that the orientations estimated using Polaris is significantly more accurate than those using raw compass readings.

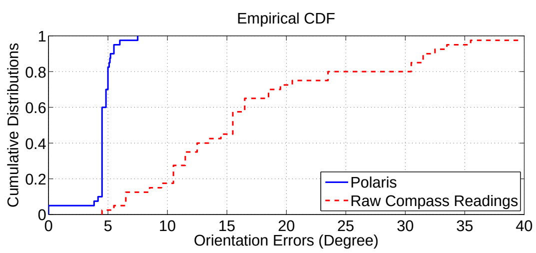

Figure 10: The cumulative distributions of the orientation errors when using raw compass readings vs. Polaris. The median, the standard deviation, and the 95th percentile errors of using Polaris are 3.5X, 7.6X, and 6.5X better than those of using the raw compass readings, respectively, indicating a significant performance improvement.

well as raw compass readings in the building. This means that the system needs a sufficient amount of data before starting to provide orientations to the users. In some certain circumstances, this process may take days or a couple of weeks to accomplish.

As the widespread use of online map services, such as the Google Maps and the Bing Maps, obtaining building orientations can be automatically done using image processing techniques to the maps of the buildings. Our preliminary experiments on determining building orientations using extracted straight-line perimeters of buildings on Google Maps' *map views* achieved 0.1° accuracy in the orientations of the detected lines. However, we found detecting straight lines became much more difficult on *satellite views*, primarily due to the existence of cars, trees, and road marks on the map. There are related work of robust techniques to detect building perimeters from satellite views [9], but that is beyond the scope of this work.

It should be noted that within Polaris' current 4.5° orientation error, 3.5° was caused by the erroneous building orientation through

8. CONCLUSIONS

In this paper, we propose Polaris, a novel approach to provide reliable orientation information for mobile devices in indoor environments. By using computer vision techniques to robustly extract orientation information from patterns visible on ceilings, our approach is not affected by magnetic interferences which are commonly present within indoor environments. As a preliminary work, we tested Polaris in multiple office buildings, and demonstrated that Polaris achieved 4.5° average orientation accuracy, which is about 3.5 times better than what can be achieved by simply using raw compass readings. As part of our ongoing work, we are working on leveraging online map services in the process of determining building orientations, and also incorporating Polaris in indoor localization and navigation systems.

9. REFERENCES

[1] M. Azizyan, I. Constandache, and R. Roy Choudhury. SurroundSense: Mobile Phone Localization via Ambience Fingerprinting. In *Proceedings of the 15th Annual International Conference on Mobile Computing and Networking*, pages 261–272. ACM, 2009.

[2] D. Ballard. Generalizing the Hough Transform to Detect Arbitrary Shapes. *Pattern Recognition*, 13(2):111–122, 1981.

[3] J. Canny. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698, Nov. 1986.

[4] T. Choudhury, G. Borriello, and S. Consolvo. The Mobile Sensing Platform: An Embedded Activity Recognition System. *IEEE Pervasive*, pages 32–41, 2008.

[5] G. Essl and M. Rohs. SHAMUS: A Sensor-Based Integrated Mobile Phone Instrument. In *Proceedings of the International Computer Music Conference ICMC (2007)*, volume 40, page 50. Citeseer, 2007.

[6] D. Gusenbauer, C. Isert, and J. Krosche. Self-Contained Indoor Positioning on Off-The-Shelf Mobile Devices. In *Proceedings of the 2010 International Conference on Indoor Positioning and Indoor Navigation*, number September, pages 1–9. IEEE, 2010.

[7] B. Hoff and R. Azuma. Autocalibration of An Electronic Compass In An Outdoor Augmented Reality System. *Proceedings IEEE and ACM International Symposium on Augmented Reality (ISAR 2000)*, 2000(October):159–164, 2000.

[8] M. Hoshino, Y. Gunji, S. Oho, and K. Takano. A Kalman Filter to Estimate Direction for Automotive Navigation. In *Proceedings of the 1996 IEEE/SICE/RSJ International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pages 145–150. IEEE, 1996.

[9] X. Jin and C. H. Davis. Automated Building Extraction from High-Resolution Satellite Imagery in Urban Areas Using Structural, Contextual, and Spectral Information. *EURASIP Journal on Advances in Signal Processing*, 2005(14):2196–2206, Jan. 2005.

[10] T. King, S. Kopf, T. Haenselmann, C. Lubberger, and W. Effelsberg. Compass: A Probabilistic Indoor Positioning System Based on 802.11 And Digital Compasses. In *Proceedings of the 1st International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*, number September, pages 34–40. ACM, 2006.

[11] W. Piekarski and B. Thomas. ARQuake: The Outdoor Augmented Reality Gaming System. *Communications of the ACM*, 45(1):36–38, 2002.

[12] D. Spinellis. Position-Annotated Photographs: A Geotemporal Web. *IEEE Pervasive Computing*, 2(2):72–79, Apr. 2003.

[13] P. Steadman. Why Are Most Buildings Rectangular? *Architectural Research Quarterly*, 10(02):119–130, June 2006.

[14] M. Youssef, M. Yosef, and M. El-Derini. GAC : Energy-Efficient Hybrid GPS-accelerometer-compass GSM Localization. In *Proceedings of the IEEE Globecom 2010*, 2010.

[15] Y.T.Kim. Contrast Enhancement Using Brightness Preserving Bi-histogram Equalization. *IEEE Transactions on Consumer Electronics*, 43(1):1–8, 1997.