IODetector: A Generic Service for Indoor/Outdoor Detection

MO LI, PENGFEI ZHOU, YUANQING ZHENG, and ZHENJIANG LI,
Nanyang Technological University
GUOBIN SHEN, Microsoft Research Asia

The location and context switching, especially the indoor/outdoor switching, provides essential and primitive information for upper-layer mobile applications. In this article, we present IODetector: a lightweight sensing service that runs on the mobile phone and detects the indoor/outdoor environment in a fast, accurate, and efficient manner. Constrained by the energy budget, IODetector primarily leverages lightweight sensing resources, such as light sensors, magnetism sensors, and cell tower signals. For universal applicability, IODetector assumes no prior knowledge (e.g., fingerprints) of the environment and uses only on-board sensors common to mainstream mobile phones. Being a generic and lightweight service component, IODetector greatly benefits many location-based and context-aware applications. We prototype the IODetector on Android mobile phones and evaluate the system comprehensively with data collected from 34 traces that include 133 different places during a 6-week period, employing different phone models. We further perform a case study where we make use of IODetector to instantly infer the GPS availability and localization accuracy in different indoor/outdoor environments.

Categories and Subject Descriptors: C.2.4 [**Computer Communication Networks**]: Distributed Systems—*Distributed applications*; C.3.3 [**Special-Purpose and Application-based Systems**]: Real-Time and Embedded Systems

General Terms: Design, Implementation, Measurement

Additional Key Words and Phrases: Indoor and outdoor detection, mobile phones, GPS availability

ACM Reference Format:
Mo Li, Pengfei Zhou, Yuanqing Zheng, Zhenjiang Li, and Guobin Shen. 2014. IODetector: A generic service for indoor/outdoor detection. ACM Trans. Sensor Netw. 11, 2, Article 28 (December 2014), 29 pages.
DOI: http://dx.doi.org/10.1145/2659466

1. INTRODUCTION

Current mobile phones are becoming important platforms that serve the ubiquitous sensing and communication needs of people [Lane et al. 2010]. The sensing and communication modules on mobile phones are usually developed to provide location and context-aware services. However, they may have different availabilities, energy, and accuracy profiles in different environments. An effective indoor/outdoor detection scheme can provide primitive environment information for a variety of mobile applications and thus potentially improve their performance. For example, in location-based applications, people usually source GPS for an accurate location reference when they are in the outdoor environment. In contrast, GPS performs poorly without line-of-sight paths to satellites when mobile devices are inside buildings [Bahl and Padmanabhan 2000; Thiagarajan et al. 2009]. In mobile data services, mobile phones normally observe more

Author Addresses: M. Li, P. Zhou, Y. Zheng, and Z. Li, School of Computer Engineering, 50 Nanyang Avenue, Singapore 639798; G. Shen, Building 2, No. 5 Dan Ling Street, Haidian District, Beijing, P. R. China, 100080.

DOI: http://dx.doi.org/10.1145/2659466

WiFi access points (APs) with strong signals inside buildings, whereas it is unlikely to have good WiFi connections in outdoor environments. Therefore, knowing whether the environment is indoors or outdoors can help to make smarter decisions regarding whether to turn on GPS or to perform AP scanning. In the context and activity recognition applications, the knowledge of the surrounding indoor/outdoor environment potentially leads to more accurate recognition. Although many applications may benefit from accurate and prompt indoor/outdoor information, the research study toward generic indoor/outdoor detection surprisingly lacks. Many location-related works simply assume that a clear preknowledge on the indoor/outdoor environment has been known, but such an assumption hardly holds in practice. The unavailability or performance degradation of GPS is sometimes used to infer the indoor/outdoor environment, yet such an approach suffers from low accuracy, high energy consumption, and long response time.

In this article, we present the Indoor/Outdoor Detector (IODetector): a generic and lightweight service for the indoor/outdoor detection in mobile applications. Constrained by the energy budget on mobile phones, we primarily make use of three lightweight sensing resources: light sensor, cellular module, and magnetism sensor. Through a 6-week experiment, we observe that the light intensity, the cell tower signal, and the intensity of magnetic field all individually exhibit distinct patterns in the indoor and outdoor environments. Those patterns turn out to be viable for an accurate classification of the ambient environments. More precisely, light signals exhibit distinct patterns when they are captured inside and outside buildings. The reason behind is that the natural and manmade light sources contain inherent difference by nature. The received signal strength (RSS) from a cell tower by a mobile phone changes dramatically from the outdoor to indoor environments as the dividing walls block the line-of-sight paths between the mobile phone and the cell tower. The intensity of magnetic field varies significantly across different places inside buildings due to the ambient electric appliances and steel structures but remains much less fluctuated across an outdoor environment. Motivated by those facts and observations, we target at achieving the indoor/outdoor detection by exploiting the three sensing resources.

Translating such an idea into a practical indoor/outdoor detection service entails a wide range of challenges, as the three aforementioned sensing resources show distinct pros and cons. The ambient light intensity may vary over time and is potentially influenced by various factors (e.g., people movement, phone pose, and cover of sight). The absolute cell tower signal strength may vary significantly at different places and across different mobile phone models, making it difficult to confidently set a uniform rule for the indoor/outdoor classification. The magnetometer readings are error prone without careful calibrations. We develop practical solutions to cope with the preceding challenges in IODetector. In particular, we extract unique identifiable indoor lighting features to detect the indoor/outdoor environment and leverage particular light intensity patterns to improve the detection accuracy (Section 3.2). We exploit the abrupt period of the cell tower signal strength rather than its absolute value to distinguish the indoor/outdoor context. We track the cellular signals from multiple visible cell towers so as to enhance the robustness of the indoor/outdoor detection (Section 3.3). We take advantage of the magnetic disturbance inside buildings and make use of the movement status from accelerometers to ensure the detection performance (Section 3.4).

We constructively combine the three sensing components and develop an extensible indoor/outdoor detection framework. By taking other ambient sensing readings and evaluating the confidence levels of three sensing units, we intellectually aggregate their detection results and guarantee optimized reliance on those sensing units. The developed IODetector then works as an underlying service module that can be invoked by upper-layer applications to provide instant indoor/outdoor information (Section 3.5).

We implement and evaluate IODetector with an Android prototype using different mobile phone models. We test IODetector in 34 traces including 133 different sites in our campus and city areas, and demonstrate quite encouraging results. Since IODetector only relies on lightweight sensors, the low energy cost allows continuous tracking of indoor/outdoor state transitions. In particular, we perform a case study and show that we can utilize IODetector to cheaply and accurately infer the current availability and accuracy of the GPS module for mobile phones.

The rest of this article is organized as follow. In Section 2, we detail the background of and motivation for this work. We describe the technical solutions of IODetector in Section 3. We present the evaluation results in Section 4 and review related works in Section 5. Finally, we conclude the article in Section 6.

2. BACKGROUND AND MOTIVATION

Indoor/outdoor detection can provide essential and primitive information for upper-layer mobile applications. For example, before turning on GPS, one may first check whether it is outside a building to ensure the GPS performance. As another example, before searching for WiFi APs, one may check whether it is inside or near a building and adapt the scanning strategy accordingly. The indoor/outdoor information is also highly useful for cameras whose energy consumption and processing time depend on the ambient environments [LiKamWa et al. 2013]. Primitive indoor/outdoor information helps them achieve energy efficiency and higher performance. Many cellular service providers [SingTel 2013; China Telecom 2012] usually receive complaints from their customers about the poor communication performance. A deeper understanding of the users' indoor/outdoor environment definitely helps them diagnose the communication problems and improve their service performance. There are also many research works on human mobility modeling from both the users and cellular providers [Isaacman et al. 2012; Liu et al. 2013], and the indoor/outdoor detection is the most significant input for those models. Many other applications, including automatic image annotation [Qin et al. 2011], context and activity recognition [Keally et al. 2011], and indoor localization [Chung et al. 2011], may also rely on the indoor/outdoor knowledge for a proper working scheme. If the detection overhead (depending on the application profile) is sufficiently small, most location and context-aware applications will greatly benefit from such indoor/outdoor detection.

While practically useful, the problem of indoor/outdoor detection has not been thoroughly studied yet. Existing localization and tracking applications may indirectly infer the ambient environment with the availability and accuracy of the GPS signal. It is well known that localization and tracking systems perform poorly in the indoor environment as the line-of-sight paths to GPS satellites are blocked. The unavailability of GPS signals and the the decreasing number of the visible satellites can thus infer the indoor environment [Ravindranath et al. 2011]. Typical GPS modules, however, draw substantial amount of energy and take minutes to warm up and conduct the GPS satellite scanning on mobile phones [Thiagarajan et al. 2009]. As a result, detecting indoor/outdoor environments solely with GPS can be slow and inefficient. There are some other works relying on dedicated devices to assist the ambient environment detection. The deployment cost of such infrastructure-based approaches significantly limits the flexibility and scalability for general-purpose detection [Smith et al. 2004]. On the other hand, some recent works study the problem of logical localization by sensing the surrounding environment [Lu et al. 2009; Azizyan et al. 2009; Liu et al. 2010]. By painstakingly fingerprinting ambient signals (e.g., sound, floor color, user movement), the mobile phones can learn the ambient environment through an intensive site survey. A central server is normally needed to store such ambient fingerprints and answer queries from users. Such an approach is unlikely to be generalized to deal

Environment	**Outdoor**	**Semioutdoor**	**Indoor**
Definition	Outside a building	Near a building	Inside a building
Example	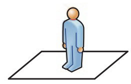	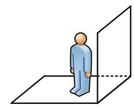	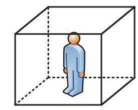
Scene	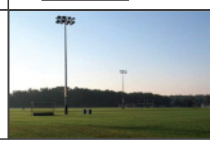		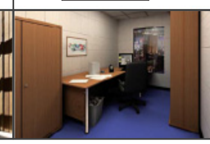

Fig. 1. Three indoor/outdoor environment types and the representative scenes. In reality, specific applications may need finer or rougher classification based on their requirements, and they can make parameterized classification using the techniques introduced in this article.

with universal indoor/outdoor detection. Many works in image processing and pattern recognition study the problem of the indoor/outdoor image classification and automatic image tagging [Payne and Singh 2005; Szummer and Picard 1998; Qin et al. 2011]. Such approaches cannot directly be applied to our problem, as they require explicit, manual input from users.

In this work, we propose IODetector, a lightweight indoor/outdoor detection framework that independently runs on each mobile phone and provides generic service to upper-layer applications. As a basic component that might frequently be invoked by many applications on energy-constrained mobile phones, IODetector needs to meet several stringent design requirements:

—*High accuracy*. As a generic framework that many other applications would potentially rely on, IODetector should accurately detect the indoor/outdoor environment.
—*Prompt response*. IODetector should promptly distinguish the indoor/outdoor environment. An outdated detection result may be less valuable for many instantaneous applications.
—*Energy efficiency*. Being a generic service running on mobile phones with constrained energy budgets, IODetector should be energy efficient and use only inexpensive sensing resources.
—*Universal applicability*. IODetector should avoid relying on a priori knowledge (or site survey), special sensors, or explicit user feedback to ensure wide applicability.

Before we present the design of IODetector in detail, we formally define the indoor/outdoor environment types studied in this work. To provide fine-grained context information for upper-layer applications, we classify the environment into three categories: outdoor (outside a building), semioutdoor (close to a building or a semiopen building), and indoor (inside a building). Figure 1 illustrates representative scenes for the three environment types. The reason of introducing the category of semioutdoor is mainly due to potential application needs. For instance, the GPS may not necessarily perform well even if it is outdoors, as the number of visible line-of-sight satellites might be insufficient in many semiopen environments. In such cases, we may not prefer to launch the GPS component. On the contrary, the situation could become different for other types of applications. One typical example is that mobile phones normally can find WiFi APs in indoor environments. Yet in most semioutdoor environments, mobile phones may still detect a number of APs with good connections. Additionally, some rooms with large windows could be treated as being semioutdoors,

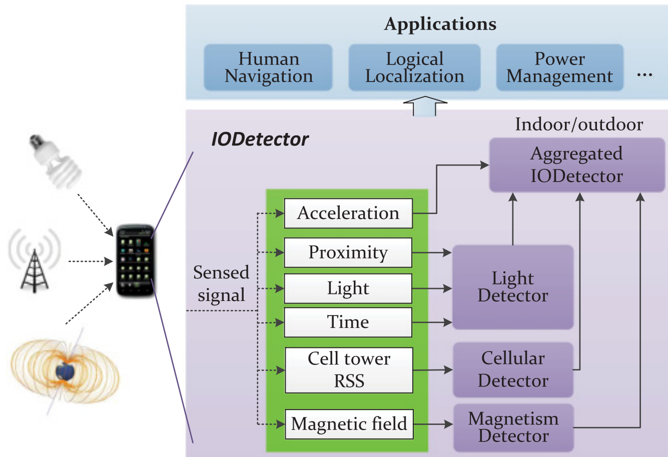

Fig. 2. System architecture of IODetector. There are three major sub-detectors and some assistant sensors. We use accelerometer to detect user mobility. Each sub-detector works independently and their detection results are aggregated to make the final detection decision.

since it is possible to receive good GPS signal in such environment, although it can be less accurate as in fully outdoor environments. Thus, the designed IODetector does not simply output a binary result (i.e., indoor or outdoor) for upper-layer applications. Instead, it provides finer-grained classification on the indoor/outdoor scenes and thus better meets different application needs.

3. SYSTEM DESIGN

In this section, we first introduce the system architecture and design details for each component in IODetector. Then we specify how to aggregate the outputs obtained from each component to construct a comprehensive and effective indoor/outdoor detector.

3.1. System Overview

Figure 2 illustrates the system architecture of IODetector. To meet stringent design requirements, IODetector utilizes a series of lightweight sensors for the indoor/outdoor detection. IODetector primarily makes use of three lightweight detectors: the light detector, cellular detector, and magnetism detector. The *light detector* adopts light sensors to capture ambient light signals to determine the surrounding environment type. It also utilizes two other lightweight sensors, the proximity sensor and the system time clock, to assist the detection. The *cellular detector* detects the attenuation of cellular signals caused by obstacles (e.g., walls). It normally indicates the entrance/exit of the device to/from an indoor environment. The *magnetism detector* exploits the dramatic disturbance of the magnetic field inside or in the vicinity of buildings during the movement of the mobile phone. It thus can distinguish the indoor/semioutdoor environments from the outdoor environment.

Note that each component of IODetector shows unique advantages and disadvantages in different environmental contexts. They process the sensor data and report the *respective* and *partial* detection results. IODetector then aggregates those results and generates a final decision, which is provided to upper-layer applications through a service interface.

There are many other sensing resources available on today's mobile phones, such as gyroscope, WiFi, camera, and microphone. Compared with the sensing resources used in our system, they are relatively not suitable to be used for indoor/outdoor detection due to the following considerations. First, they are not distinguishable for different environment types. For example, there is no common observation that the outdoor ambient noises are different from indoor ambient noises. There could be some difference for specific sites but not for common sites. One way to use such kind of signals is to war drive and label the sites using the signals as fingerprints. However, such a method incurs substantial operation overhead and is not scalable. Second, some of them are not stable and accurate enough to be used. For example, in our system, we make use of cellular signal over other wireless signals (e.g., WiFi) mainly due to such a consideration. Cell tower signal is available with no additional energy cost since mobile phones have to maintain connectivity to cell towers for basic communication, and cellular networks have almost universal coverage, both outdoors and indoors. But for those high-frequency band signals such as 2.4GHz WiFi signal, because of the short wavelength, they may severely suffer from the shielding effect of surrounding objects or even the human body itself [Zhang et al. 2011b], which will bring too much noise into the detection system. On the contrary, the cell tower signal of much longer wavelength can easily diffract around these objects. The WiFi signals also suffer from poor availability in many outdoor areas. Third, some of them are not energy efficient. For example, the GPS lock is potentially useful to detect indoor/outdoor environments, but its energy consumption is high. To achieve a generic service for indoor/outdoor detection, we avoid using energy-expensive sensing resources.

In the rest of this section, we will describe the design details of each component. To reveal the signal features with different environments, we empirically study the patterns of light signal, cell tower signal, and magnetism signal in different environments for 2 weeks. All of the signals are collected in 31 different environments under different weather conditions, including sunny, cloudy, and rainy days, and at different times of the day. The studied sites include indoor offices, homes, stores, outdoor campuses, some downtown areas, and so forth. For each site, we collect light signal six times, magnetism signal four times, and cell tower signal four times on average with different sampling rates. The light signal is collected with different orientations of the light sensor, and the cell tower signal is collected when the user walks from outdoors to indoors and vice versa.

3.2. Light Detector

In outdoor and semioutdoor environments, the sun is the primary light source in the daytime. In the indoor environment, however, we normally rely on artificial light sources (e.g., fluorescent lamps).

3.2.1. Light Intensity. Our primary observation for the light detector is that the light intensity inside buildings is typically much lower than that in either the outdoor or semioutdoor environment even on cloudy or rainy days. The major reason for this is that the intensity of sunlight within the visible spectrum is normally much higher than that from ordinary lighting lamps. In addition, light sensors can also detect the light in the invisible spectrums (e.g., infrared and ultraviolet). As a result, even when the brightness of sunlight and artificial light looks similar, the luminous flux from sunlight is much higher than that from artificial light sources during the daytime. Therefore, the indoor environment can be accurately distinguished from the outdoor/semioutdoor environment by using the observed light intensity.

To verify the preceding statement, we conduct a set of experiments. We measure the light intensities in different environment types under different weather conditions. In Figure 3(a), we plot the light sensor readings from three different types of mobile phones (HTC Desire S, HTC Sensation G14, and Samsung Galaxy S2 i9100). Many current Android platforms, however, only provides coarsely quantized light sensor readings for

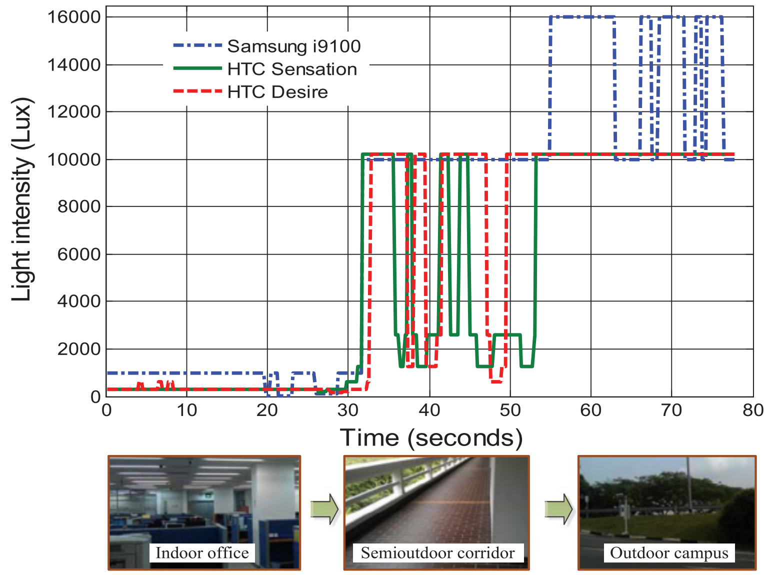

(a) Phone light sensor readings in different scenes with three types of mobile phones. The light sensor readings are discrete but clearly different for different environments.

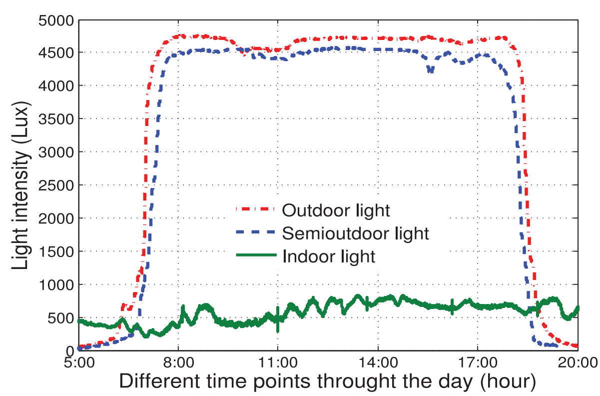

(b) Light variation throughout the whole day (from 5:00AM to 20:00PM). The measurements are made in three different environments simultaneously.

Fig. 3. Light signal measurements.

upper-layer applications. For instance, the Samsung Galaxy S2 i9100 only provides five quantized levels (10, 100, 1000, 10,000, and 160,000), and the light intensity will be rounded to the closest quantized level. From Figure 3(a), we can see that readings of the light intensity from all three mobile phones are discrete and coarse. Yet the readings still show clear and consistent transition behaviors in the experiments. When the user moves outside of the office at the 30-second point, the light sensor readings collected on all the three mobile phones increase significantly.

To further investigate the effectiveness of utilizing the light intensity to distinguish the indoor environment, we further collect light intensities in three different environments using three TelosB motes in a cloudy and rainy day. We modify TinyOS code to

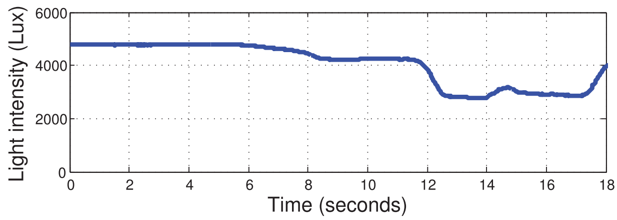

Fig. 4. Outdoor light intensity during the rotation. The sensor is uniformly rotated from facing the sun (the 1st second) to back to the sun (the 16th second).

directly read the voltage on the light sensor S1087-01. The sensitivity range of the light sensor is from 300nm to 1,200nm with a full coverage of the visible light spectrum and a partial coverage of the infrared and ultraviolet spectrum. We note that the results and observations obtained in this experiment can also be used to improve system performance. In this experiment, the sampling rate of the light sensor is set to be one sample per second. From Figure 3(b), the light intensities in both the outdoor and semioutdoor scenarios are above 2,000Lux and much higher than that in the indoor environment in the daytime (from 8:00AM to 17:00PM). We also find that during the night (from 20:00PM to 5:00AM), the outdoor light intensity is much smaller than indoor light intensity. In addition, the light intensities in the indoor and outdoor environments are both relatively stable. This observation is consistent with the indoor lighting standards and measurements [Wikipedia 2013b], which show that in most cases, the indoor light intensity is within the interval from 100Lux to 1,000Lux.

In practical scenarios, the mobile phone does not necessarily face the sun and the phone may be dynamically rotated. To examine the robustness of our method, we record the detected light intensity when rotating a TelosB mote in Figure 4. The light sensor initially faces the sun and is gradually rotated until being toward an opposite direction. Figure 4 shows that even when the light sensor is back to the sun (from the 13th sec to 17th sec), the light intensity is relative high as well (e.g., around 3,000Lux). Compared with the light intensity observed in the indoor scenario as shown in Figure 3(b), we can still distinguish them easily. Therefore, the detection of the light intensity is robust to the mobile phone dynamics.

3.2.2. Light Flicker. We have demonstrated that the light intensity can serve as a good feature of the indoor environment in the daytime. One problem, however, still remains. If the detected light intensity is low, we still cannot confidently determine whether it is an indoor or outdoor environment at night. As shown in Figure 3(b), the outdoor light intensity might become comparable to or even lower than the light intensity from artificial light sources at night. In this case, solely using the light intensity is not adequate.

To tackle this issue, we propose utilizing more light features to jointly make the classification decision. Our proposed solution is based on the following facts. Fluorescent lamps are the most widely available light sources in many indoor environments because of the low manufacturing cost and the high energy efficiency. Fluorescent lamps, powered by the alternating current (AC) power, emit the light with periodic patterns, which we call *light flicker* in this article. Suppose that the frequency of AC is f. After rectifying, the frequency of fluorescent light intensity will be $2f$. Although the periodical varying light intensity from fluorescent lamps is almost imperceptible to eyes, such distinctive periodic patterns can be captured by light sensors [Li et al. 2012]. The sunlight, on the contrary, does not exhibit such a periodic pattern.

To verify the effectiveness of the method mentioned earlier, we measure the frequency of fluorescent light flicker using light sensors on TelosB mote in an office building. The AC frequency is 50Hz, and we sample the light sensor with a 4kHz sampling rate

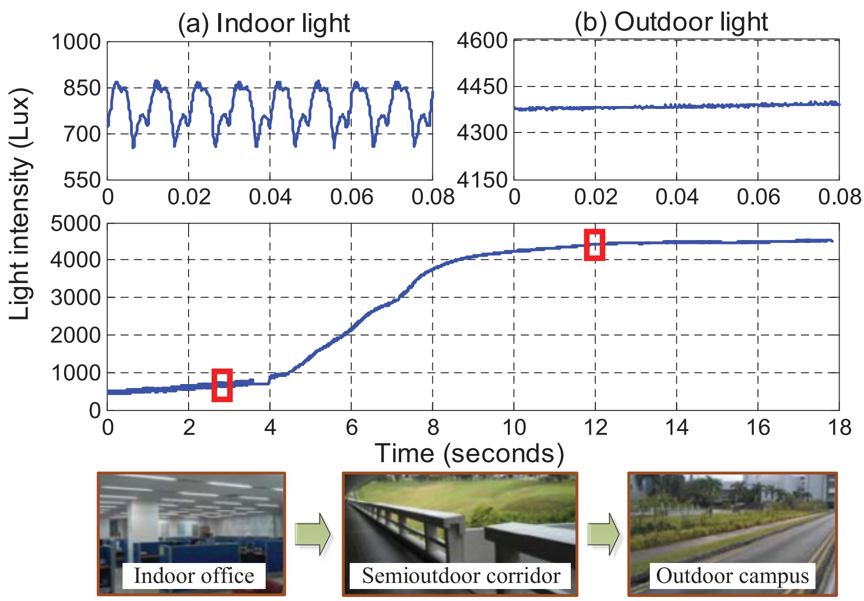

Fig. 5. Light intensity exhibits a periodic pattern in the indoor environment as shown in (a), whereas the periodic pattern is hardly observed outside buildings during the daytime as shown in (b). The overall trend of light intensity variation is plotted in the bottom part of the figure.

that is sufficiently high to capture the periodic pattern of the fluorescent light. In the experiments, we see a clear peak at 100Hz frequency band after a fast Fourier transform (FFT) in the indoor environment, whereas in the outdoor environment, such a peak is hardly observable in the daytime. In Figure 5 (bottom), we plot the light intensity along a trajectory from an indoor office to an outdoor field passing through a semioutdoor corridor. Figure 5(a) plots the finer variation of indoor light signal at the 3rd second. We can observe the distinct periodical light intensity due to the alternating pattern in the indoor office environment. When leaving the office and entering the semioutdoor corridor (e.g., the 12th second), we find that the periodicity of the light signal disappears immediately (Figure 5(b)). We also examine the detection performance when the light from fluorescent lamps is mixed with the sunlight. It usually occurs when the mobile phone is close to the window. From our experiment, we observe that the sunlight contributes a direct component to the final result. The periodical pattern from the fluorescent lamps nevertheless remains to appear, and we can still detect its existence using FFT. The observed periodic patterns in one period may exhibit multiple peaks (Figure 5(a)). This is because the observed result is a mix of the light from multiple lamps, and each lamp may have different phases. In this example, there are three dominant phases, as shown in Figure 5(a). The electric characteristics of the bulb starter probably cause the light phase difference. Although all of the fluorescent lamps may be powered by the same AC power, the bulb starter of each lamp performs differently, which results in phase difference. As a matter of fact, the bulb starter of a particular fluorescent lamp performs differently when we turn on the lamp at different two times.

We also examine the robustness of the light flicker detection with dynamics. We record the light intensity when rotating a TelosB mote in the indoor environment. Figure 6 plots the light intensity in the indoor environment. The light sensor initially faces to fluorescent lamps (at the 1st second) and is gradually rotated until being

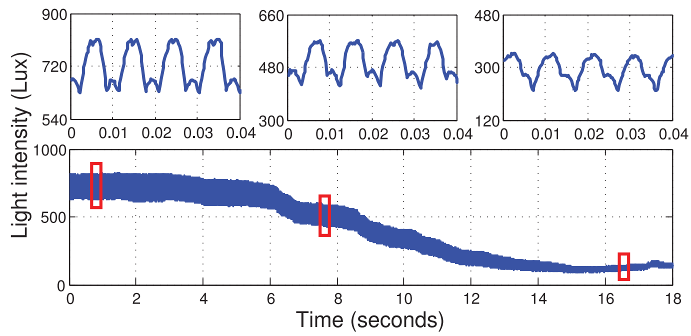

Fig. 6. Indoor light flicker during the rotation. The sensor is uniformly rotated from facing the lamp (the 1st second) to back to the lamp (the 16th second). The top three graphs show the consistent periodical patterns during the rotation.

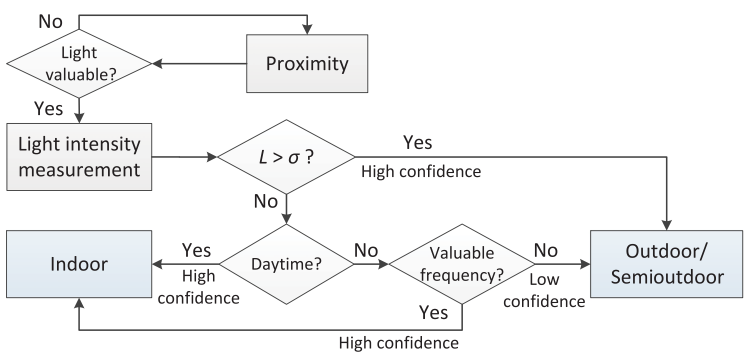

Fig. 7. Detection flow of the light detector.

turned back to the lamps (at the 16th second). From Figure 6 (bottom), we observe that the overall light intensity decreases. Yet when we take a fine look at any specific portion (the 1st, 8th, and 16th seconds) of the collected light signals (as shown in the top three graphs), we still see distinct periodic patterns at 100Hz due to the alternating intensity of the fluorescent lamps. Such periodic component can be easily captured in the frequency domain at the 100Hz band after FFT.

3.2.3. Detection Process in Light Detector. Since mobile phones may be placed in pockets or bags, the light sensors may not be always available. We use the proximity sensor, which is usually embedded at the same position as the light sensor on mobile phones, to detect the presence of nearby objects that may block the light sensor. We associate a confidence level $C_L \in [0, 1]$ for the detection result. Different light signals will lead to different detection confidence levels.

Figure 7 summarizes the work flow of the light detector component. We denote L as the detected light intensity. The light detector first queries the proximity sensor to check whether the light sensor is currently available. If the light sensor is available, the light intensity L is then compared with a threshold σ. If $L > \sigma$, the light detector confirms an outdoor/semioutdoor environment detection with a high level confidence $C_L = 1$; if $L \leq \sigma$, it needs to further differentiate whether it is an indoor environment or an outdoor/semioutdoor environment at night. To this end, the light detector refers to the system clock. If the clock indicates daytime, the detector infers the environment to be indoors with a high confidence. If not, the light detector turns to the frequency domain detection using FFT. After FFT, if the light flicker frequency f_{light} is within the interval $[f_{lower}, f_{upper}]$ (the selection of f_{lower} and f_{upper} will be discussed soon), it indicates an indoor environment with a high confidence level $C_L = 1$; otherwise, the mobile phone is in an outdoor/semioutdoor environment with a confidence level $C_L = \frac{\sigma - L}{\sigma}$.

We now study how to properly configure each parameter involved. From Figure 3(b), the sunlight intensity in the daytime is distinguishable from that of indoor lights. According to our empirical study, we set the threshold σ to 2,000Lux. The frequency of the electrical system varies by countries; nevertheless, most electric power operates at either at 50 or 60 Hz [Rowe et al. 2009]. Thus, the fluorescent light intensity frequency is normally 100 or 120 Hz, which simplifies the frequency domain detection. We therefore choose a reasonable frequency interval $[f_{lower} = 95, f_{upper} = 125]$ to enhance detection robustness and reduce the computation overhead of FFT. Say that $L = 1{,}000 < \sigma$ and $f_{light} = 100 \in [95, 125]$, then the light detector will confirm the ambient environment as the indoor environment with the confidence level $C_L = 1$.

Although the readings of light sensors on many current mobile phone models are coarsely quantized to upper applications, we see the trend that more and more mobile phones grant upper-layer applications finer access to those lower-level sensors, such as with Samsung Google Nexus One. We believe that most mobile phone models will support such a functionality in the near future. The limitation of the light detector is that the light signal is not always available. In addition, we cannot confidently distinguish the outdoor and semioutdoor environments by merely using light sensors.

3.3. Cellular Detector

Mobile phones maintain connections to nearby cell towers to support the primary functionality—that is, the telephone call. The marginal energy consumption of collecting cellular RSS is thus negligible. Previous works utilize the information about visible cell towers and their signal strength for localization and tracking [Thiagarajan et al. 2010]. Such approaches, however, suffer from low accuracy due to various factors. One primary issue is the dividing wall effect, which refers to the fact that the dividing wall significantly blocks the cellular signal and hence leads to dramatic signal strength drop when people get into indoor environments. Unlike the localization works where the dividing wall effect is undesired, in this work, we embrace and exploit the cellular RSS variation for indoor/outdoor detection.

3.3.1. Associated Cell Tower Signal Strength. We aim to find the correlation between the cellular signal variation and the surrounding environment transitions. We first measure the cellular RSS in several representative places, such as offices and homes (indoor), corridors and paths in the vicinity of building (semioutdoor), and plaza and football field (outdoor). We find that the absolute value of the cellular RSS provides limited information for the detection. It varies across different places, times, and phone models. In contrast, the RSS variation within a short period of time normally indicates the context transition. In our experiments, we observe a significant variation of the cellular RSS when the ambient environment changes. For instance, when the user walks into an office building from outside, the cellular RSS significantly drops due to the dividing walls that block the line-of-sight paths to cell towers. Therefore, we exploit the abrupt variation of the cellular signal strength rather than its absolute value to distinguish the indoor/outdoor context that is invariant across different places and phone models.

To guarantee the communication quality, a mobile phone usually connects to the cell tower with the strongest RSS. Figure 8(a) shows the RSS value from the connected cell tower when the user walks out to the corridor and then back to the office. The user walks outside at the 30th second. We can see that the RSS rises by approximately

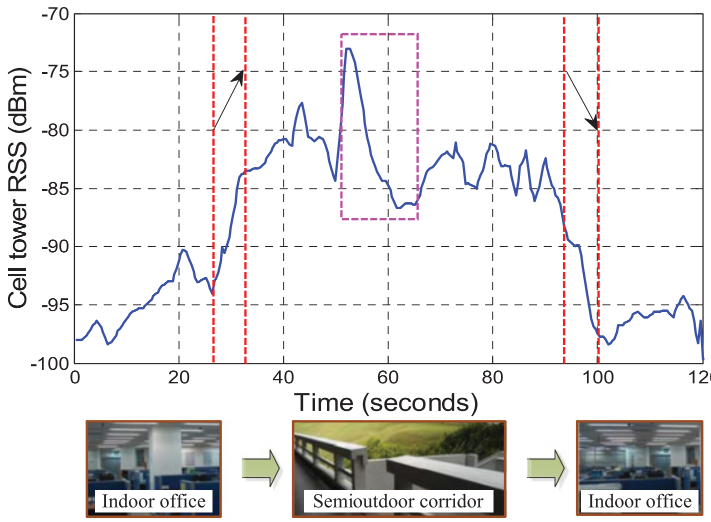

(a) Single cell tower's RSS variation in environmental change. The RSS jumps from indoor to outdoor and drops from outdoor to indoor.

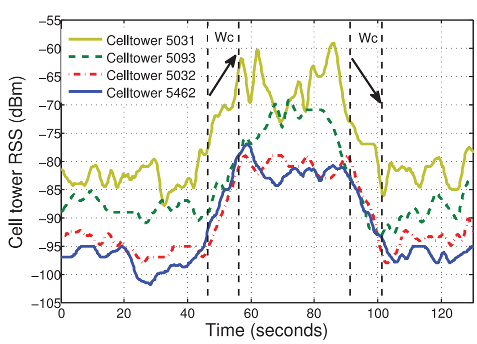

(b) Multiple cell towers' RSS variation in environmental change. The RSS of all cell towers exhibits the same trend in the environment change.

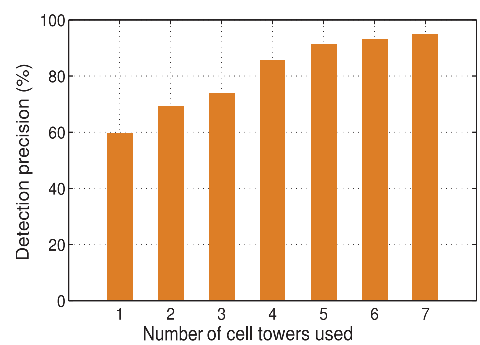

(c) Detection accuracy with the varied number of cell towers. Generally, the accuracy increases as the number of cell towers increases.

Fig. 8. Cell tower signal strength variation for indoor/outdoor detection.

15dB. Then at about the 90th second, the user comes back to the office and the RSS drops back within 10 seconds. Such sharp cellular RSS variation can be used to detect the ambient environment changes. On the other hand, since the antenna gain may vary across different mobile phone models, it is hard to accurately map different RSS values to different environments. Adopting the RSS variation can avoid the detection error that would arise if the absolute RSS value were used, especially when applied on diversified devices and environments. In short, our cellular detector is independent of mobile phone models and environments, which ensures universal applicability.

However, we notice that using RSS information of the single associated cell tower suffers from two inherent limitations. First, mobile phones may handover from one cell tower to another. Such a handover normally introduces a significant cellular RSS variation. In this case, the RSS variation may not necessarily imply an indoor/outdoor transition. Second, due to the corner effect [Tripathi et al. 1998], the cellular RSS may dramatically change in the semioutdoor environment. For example, in Figure 8(a), the RSS suddenly drops by about 15dB at 50 seconds when the user turns around at a corner. The corner effect usually happens in the semioutdoor environment due to the change of the line-of-sight to cell towers.

3.3.2. Visible Cell Tower Signal Strength. A mobile phone is normally within the coverage of multiple cell towers. Instead of using the single associated cell tower, we take a full advantage of all visible cell towers to improve the detection accuracy [Zhou et al. 2012a]. In particular, we measure the signal strengths of all of cell towers and track their RSS variation. Thereby, we naturally solve the inherent handover problem since the cell tower that the phone may connect to is also among the observed cell towers. In addition, with a rich set of cell towers, we can mitigate the problem of the corner effect. Actually, since the evident corner effect usually indicates a semioutdoor environment, we can exploit such a property to refine the detection.

We denote the RSS from cell tower i at time t as $R_i(t)$, $1 \leq i \leq n$. We track the RSS variation within a time interval Δ and denote the variation of cell tower i as $V_i(t) = R_i(t + \Delta) - R_i(t)$. We refer $N_+(t)$ as the number of cell towers whose RSS increases more than ν—that is, $N_+(t) = |\{i|V_i(t) \geq \nu, 0 \leq i \leq n\}|$; we also denote by $N_-(t)$ the number of cell towers whose RSS decreases more than ν—that is, $N_-(t) = |\{i|V_i(t) \leq -\nu, 0 \leq i \leq n\}|$. In some cases, we will also see that $N_+(t) + N_-(t) < n$, since the RSS of many cell towers remains quite stable and the differences do not exceed ν. We define $N_0(t) = n - N_+(t) - N_-(t)$ to represent the stability of cell tower RSS. In our experiments, we set $\Delta = 10$ seconds and $\nu = 15$dB.

Intuitively, if a user moves from indoors to outdoors, the RSS of cell towers will increase, and vice versa. In addition, the more cell towers whose RSS exhibits the same trend, the more confident the detection will be. We correspond the detection results with different confidence levels C_C. Say that we find $N_0(t) = 1$, $N_+(t) = 1$, $N_-(t) = 4$, and $n = 6$, then the cellular detector will confirm the ambient environment as the indoor environment with confidence level $C_C = N_-(t)/n = 0.67$. The cellular detector will also report the confidence level for the semioutdoor/outdoor environment as $N_+(t)/n = 0.17$.

Figure 8(b) illustrates the RSS of multiple cell towers when the user walks out to the corridor (at the 45th second) and then returns to the office (at the 90th second). In Figure 8(b), we see that the RSS of the four cell towers rapidly climbs up, which implies that the user has moved from the indoor environment to the outside. At the 90th second, the RSS of the four cell towers drops sharply, which means that the user gets back to the indoor office. During the period from the 60th to 70th second, the RSS of the associated cell tower varies significantly, whereas other cell towers remain relatively stable. In this case, the majority rule helps filter out bursts and reduces detection errors.

We note that the visible cell towers are not necessarily from the same GSM network operator. A phone may detect cellular signals from multiple GSM networks, which ensures a sufficient number of visible cell towers. In our experiment, mobile phones typically see four to six cell towers at one time. Figure 8(c) plots the detection precision of the cellular detector with the varying number of cell towers. We find that the detection precision increases as the number of visible cell tower increases, and it is satisfactory when the number of cell towers is more than four. Since mobile phones need to maintain connections to cell towers, the energy consumption of the cellular detector is almost negligible. The major limitation is that the cellular detector may perform poorly without a sufficient number of visible cell towers in some cases.

3.4. Magnetism Detector

Many steel structures and electric appliances disturb the geomagnetic field and generate the electromagnetic fields in the indoor environment [Rowe et al. 2009]. The

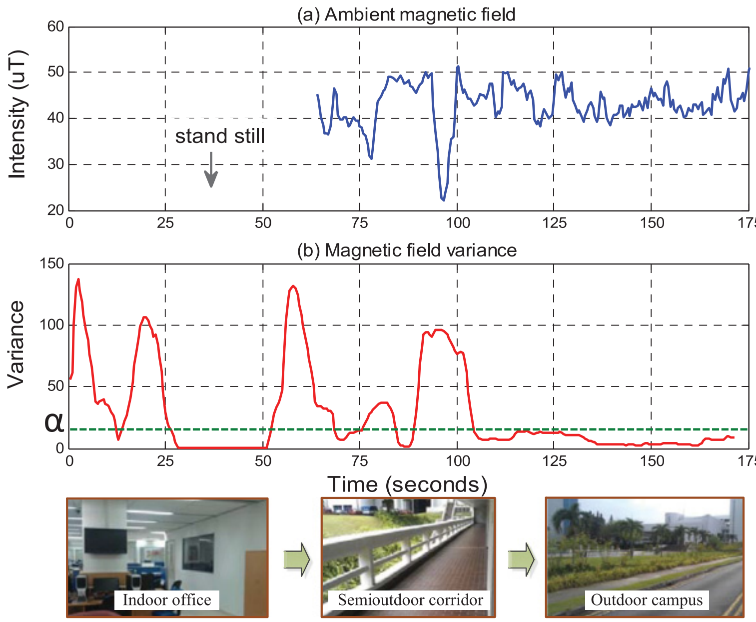

Fig. 9. The variation of magnetic field intensity. The magnetism signal varies significantly when the user moves to indoor environments but keeps relatively stable when the user gets to outdoor environments.

disturbance of the earth's magnetic field inside buildings can be utilized as fingerprints for the indoor localization [Chung et al. 2011]. However, such a localization approach requires a labor-intensive fingerprinting and cannot be applied to the indoor/outdoor detection directly. In this section, we seek to explore useful characteristics of magnetic fields in different ambient environments that may help to enhance the indoor/outdoor detection.

The magnetic field exhibits distinct patterns in indoor/outdoor environments. In the indoor environment, the earth's geomagnetic field varies at different positions due to the disturbance of steel structures and electric appliances inside buildings. For instance, the intensity of the magnetic field near the equator and near the pole varies from 0.25 to 0.65 G (i.e., 25 to 65μT). In comparison, a strong refrigerator magnet has a field of around 100G (two orders of magnitude higher) [Wikipedia 2013a]. Therefore, the intensity of magnetic fields shows a high variance across different places near and inside buildings than that in the open space.

Figure 9 plots the magnetic field intensity and its variance in an example scenario in which a user walks outside of the office, passing through a corridor. In particular, the user walks from the 1st second to 25th second, stops walking from the 25th second to 50th second inside the building, and then walks along the corridor from the 50th second to 100th second. In the end, the user walks along the road. In Figure 9(a), we find that the intensity of magnetic field in the indoor environment varies dramatically. Figure 9(b) plots the variance averaged over τ seconds to filter out noises. We find that the variance is very high when the user moves (from the 1st second to 25th second). When the user is walking through the corridor, the magnetic field intensity also shows significant variance. In contrary, after the user goess outside after the 100th second, the variance drops significantly. Therefore, by choosing a suitable threshold α, we could distinguish the indoor/semioutdoor from the outdoor environment.

Table I. Summary of the Three Subdetectors

Detector	Accuracy	Latency	Efficacy	Availability
Light	Medium	Fast	Indoor vs. Outdoor/Semioutdoor	Without blockage
Cellular	High	Slow	Indoor vs. Outdoor/Semioutdoor	Sufficient cell tower coverage
Magnetism	High	Slow	Outdoor vs. Indoor/Semioutdoor	When moving

We vary the threshold α from 0 to 40 with step length 2 and statistically analyze the detection accuracy using the collected data described in Section 3.1. If the threshold is small, most indoor/semioutdoor environments will be correctly classified, whereas many outdoor environments will be wrongly detected as the indoor/semioutdoor environment. On the other hand, if the threshold is too large, most outdoor environments will be correctly classified, but we will miss the detection of many indoor/semioutdoor environments. Therefore, we select an empirical threshold 18 to achieve a balance. In our implementation, we first refer to the accelerometer to detect whether the mobile phone is moving. If so, the magnetism detector samples the magnetism sensor and uses the variance averaged over $\tau = 10$ seconds to detect the environment. When the user stops walking (from the 25th second to the 50th second), the variance becomes very small. When the user is moving, the magnetism detector confirms the detection of an indoor/semioutdoor environment if the field variance is larger than α; otherwise, it reports an outdoor environment. Since a larger τ yields a higher detection robustness, we set the confidence level of the magnetism detector $C_M = \tau/10$.

3.5. Aggregated IODetector

Each of the three subdetectors shows unique advantages and disadvantages as summarized in Table I. They best fit different scenarios. For instance, the light detector can rapidly detect the ambient environment. The light detector, however, requires the mobile phone to be exposed in the space. If the phone is inside a pocket or bags, the light detector cannot provide accurate detection results. The cellular detector needs sufficient cell tower coverage to confidently detect the ambient context. The detection response is also slower. The magnetism detector is only available when the user is moving around such that the magnetic disturbance inside buildings can be exploited. We refer to the three individual detectors as subdetectors and integrate them to output an arbitrated decision.

At first, we directly aggregate the instant detection results of all three subdetectors. We let each subdetector report a detection profile-that is, a triplet of confidence levels for the three possible environment types and sum the confidence levels from all three subdetectors. The environment type with the highest summed confidence level will be output as the final detection result. Such a combination makes a stateless decision—in other words, the detection output is solely determined by the current environment status and the instant sensor readings. We refer to it as the stateless IODetector in the following.

Figure 10(a) shows the aggregation processing of the stateless IODetector. We denote the detection profile from the three subdetectors as $[D_L(t), C_L(t)]$ (light), $[D_C(t), C_C(t)]$ (cellular), and $[D_M(t), C_M(t)]$ (magnetism), where D is the output detection result from each subdetector and C is the set of associated confidence levels for the three possible environment types. As described in Section 3, each individual subdetector outputs the possible environment types and associated confidence levels for them. For example, each detection profile of the light detector can be denoted as $[D_L, C_L] = \{(indoor, C_{L,indoor}), (semioutdoor, C_{L,semi-outdoor}), (outdoor, C_{L,outdoor})\}$. For each possible environment type, we sum the confidence levels from the three subdetectors and obtain the

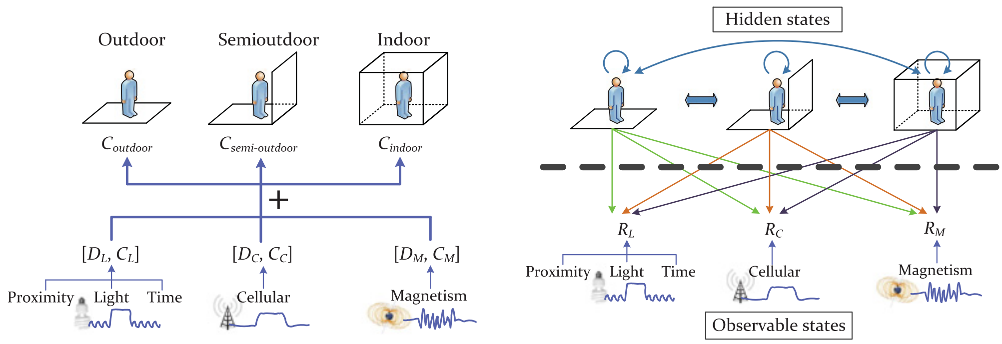

(a) Stateless IODetector. (b) Stateful IODetector.

Fig. 10. Aggregated stateless and stateful IODetectors.

triplet of overall confidence levels $C_E \in \{C_{indoor}, C_{semi-outdoor}, C_{outdoor}\}$. The environment type with the highest overall confidence level will be reported as the final detection result.

The stateless IODetector provides us with instant detection results. Users can activate IODetector on an as-needed basis. Thus, the significant out-of-the-box functionality ensures the energy efficiency of the stateless IODetector. In our experiments, however, we find that the current environment state of human beings is usually related to the previous state. For example, during the movement from indoors to outdoors, the user has a good chance of experiencing the semioutdoor environment. The stateless IODetector does not consider previous states and thus may suffer from noises. In the following, we alternatively consider a stateful integration of the three subdetectors, which makes decisions on top of both current and previous observations.

To do so, we let all subdetectors perform continuous detection and return sequential results. Figure 10(b) sketches an illustrative example of the stateful IODetector. We make use of the hidden Markov model (HMM) [Thiagarajan et al. 2009] to integrate the subdetectors. The HMM models a Markov process with underlying hidden states. Every hidden state emits observable states with particular conditional probability distribution called the *emission probability distribution*. The HMM traverses the states, and the transitions among the hidden states are governed by the transition probabilities. With the HMM, we estimate the most likely sequence of hidden states that may produce the sequence of observable states. We use the first-order HMM in which the current environment state is only affected by the immediate previous state. We denote the hidden state at time t as $H(t) \in \{indoor, semioutdoor, outdoor\}$ and the observed results from the three subdetectors as $R_L(t)$(light), $R_C(t)$(cellular), and $R_M(t)$(magnetism), where R is the output environment type with the highest confidence level from each individual subdetector. For example, the detection result from the light detector is $R_L \in \{indoor, semioutdoor/outdoor\}$. IODetector incorporates the detection results from all sub-detectors and treats them as the observable state $B(t) = [R_L(t), R_C(t), R_M(t)]$. IODetector will thus infer the most likely hidden state $H(t)$ from the previous hidden state $H(t-1)$ and the current observable state $B(t)$. The transition and emission probabilities determine the inference result.

Transition probability. We set the transition probability based on the following observation: (1) for each environment state, there is a substantial probability that the user will remain in his previous environment; and (2) there is constrain on the user's moving speed, so the transition from the previous environment's state to the current environment's state is constrained as well.

Table II. Transition Probability Settings

To		Indoor	Semioutdoor	Outdoor
From	Indoor	0.85	0.15	0
	Semioutdoor	0.33	0.4	0.27
	Outdoor	0	0.23	0.77

Table III. Emission Probability Settings

Detector	Observable State	Indoor	Semioutdoor	Outdoor
Light detector	Indoor	0.9	0.11	0.11
	Semi/outdoor	0.1	0.89	0.89
Cellular detector	Indoor	0.82	0.16	0.16
	Semi/outdoor	0.18	0.84	0.84
Magnetism detector	Semi/indoor	0.88	0.88	0.17
	Outdoor	0.12	0.12	0.83

We thus determine the transition probabilities based on the observations and the characteristics of IODetector. Since the detection period of IODetector is set to 10 seconds, when a user is previously indoors, the current environment state is probably indoors and might be semioutdoors but is not likely outdoors because the user unlikely moves directly from indoors to a fully outdoor environment. It is similar when a user is outdoors. When the user is semioutdoors, however, he could be able to directly move indoors or outdoors, or he may stay semioutdoors.

We denote the transition probability from environment H_1 to H_2 (elaborated as I:*indoor*, O:*outdoor*, and S:*semioutdoor* in the following) as $T(H_1, H_2)$. We make statistical analysis from more than 30 independent real user trajectories that cover many sites across all three environment types. We calculate the transition probabilities (summarized in Table II) based on the environment transitions experienced in those trajectories. The transition interval is set to 10 seconds. Such settings are according the environment type distributions on those trajectories that for most of the time users are either indoors or outdoors.

Emission probability. The emission probability $E(B, H)$ is the likelihood that an observable state B is observed in H environment. We set the emission probability according to the training data as described Section 3.1. Table III shows the emission probability of each hidden state (*indoor*, *semioutdoor*, and *outdoor*) to each observable state in detail.

In each detection window, the probability of each environment state $p(H) = \{p(indoor), p(semioutdoor), p(outdoor)\}$ is first calculated based on the emission probability and the detection results of the three subdetectors. Then, $p(H)$ is further updated according to the transition probability from the previous environment state. Finally, the stateful IODetector outputs the environment state with the highest probability. Thus, the transition probability does not accumulate across multiple states. For the stateful IODetector, we do not keep all the sensors on. We use the accelerometer as a trigger. Only when the accelerometer detects the user movement does IODetector activate the sensors and start to infer the new environment state from the HMM. When the user is stationary, the user environment state is deemed unchanged and all sensors are deactivated, as is is the HMM processing.

The HMM-based stateful IODetector provides sequential detection results from the sensor readings. For easy system implementation, we also provide a simplified design alternative for the stateful IODetector. The environment state of current state is denoted as s_t, and the previous environment state is denoted as s_0. For each detection,

Table IV. Experimental Sites

Environment Type	Representative Places	Total
Outdoor	12 campus sites, 11 downtown areas, on roofs of 10 buildings, 6 buses	39
Semioutdoor	15 campus sites, 12 downtown areas, 16 open windows	43
Indoor	10 office rooms, 18 stores, 6 restaurants, 8 classrooms, 7 underground rapid train stations, 2 shopping malls	51

the probability of each environment type $p(s_i)$ is calculated as

$$p(s_i) = p_t(s_i|s_0) * (p(s_i|O_l) + p(s_i|O_c) + p(s_i|O_m)),$$

where the transition probability $p_t(s_i|s_0) = 0$ for transition between indoors and outdoors directly and $p_t(s_i|s_0) = 1$ for the rest. $p(s_i|O_l)$, $p(s_i|O_c)$ and $p(s_i|O_m)$ are the probability based on the observations of the light detector, cellular detector, and magnetism detector, respectively. They can be adopted from the emission probability.

The stateful IODetector further explores the sequential observations and provides stateful detection results, which are robust to noisy measurements [Thiagarajan et al. 2009]. Its detection accuracy, which we show in Section 4.2.2, is better than the stateless IODetector. However, the stateful IODetector may consume extra energy since it has to perform continuous detection. We show the energy consumption of IODetector Sections 4.2.2 and 4.3.3. Users can choose either the stateless or stateful IODetector, whichever is more suitable for the application scenarios.

4. EVALUATION

We implement a prototype system on the Android platform with different types of mobile phones. We collect sensor data at 34 traces including 133 different sites over a 6-week period of experiments. The following details the experiment methodology and the results.

4.1. Experimental Methodology

Mobile phones. We implement IODetector on the Android platform and test its performance using four different types of mobile phones (Samsung Galaxy S2 i9100, HTC Desire S, HTC Sensation G14, and Samsung Google Nexus One). All types of mobile phones are equipped with light sensors, proximity sensors, magnetism sensors, accelerometers, and so forth. Since the light sensor on three types of mobile phones (Galaxy S2 i9100, HTC Desire S, and Sensation G14) does not provide continuous sensor readings, we only use the light intensity difference for the light detector [Zhou et al. 2012b]. The Nexus One provides us with continuous and real sensor readings, and we make full implementation of the light detector using both intensity and flicker difference on Nexus One. As IODetector is independent of platforms, we believe that the proposed indoor/outdoor detection method can be simply implanted to other mobile computing platforms, such as Apple iOS and Windows Phone.

Experiment environment. We experiment with 34 different walking traces and collect sensor readings from 39 outdoor segments (covering football fields, downtown squares, etc.), 43 semioutdoor segments (covering corridors and paths near buildings), and 51 indoor segments (including offices and shopping malls) mainly in campus and city areas (Table IV) during the period 5:00 to 22:00 in 6 weeks with different weather conditions. The users walk along these traces, and the mobile phones perform continuous detection for the experimental sites along the traces. These sites are different from the environments where we collected prior data and developed the IODetector philosophy.

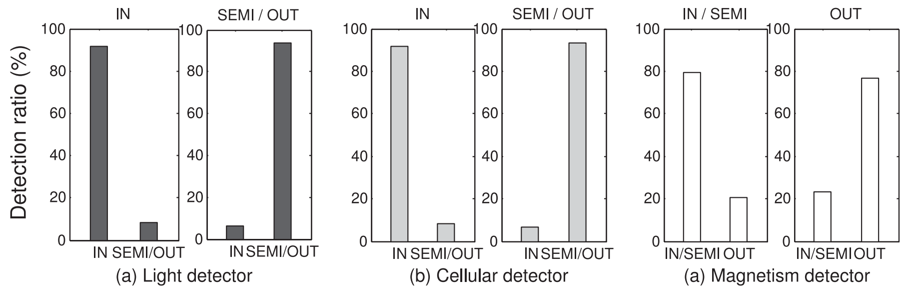

Fig. 11. Detection performance of three subdetectors. For each subdetector, we report the performance under two kinds of scenarios. For example, we test the light detector in indoor environments and outdoor/semioutdoor environments, respectively.

4.2. System Performance

In this section, we show the detection performance of the three individual subdetectors as well as the aggregated IODetector. We also compare the performance of the stateless and stateful IODetectors.

4.2.1. Performance of Subdetectors. One may query the three different detectors independently and select an arbitrary one in practice. To evaluate the contribution of each detector (i.e., light detector, cellular detector, and magnetism detector), we examine the detection performance independently in Figure 11. Each detector reports the environment type with the highest confidence level after the local computation.

The light detector is available when there are clear paths between mobile phones and ambient light sources. Figure 11(a) depicts the detection performance of the light detector. We find that the light detector can effectively distinguish the indoor environment from the semioutdoor/outdoor environment. In Figure 11(a), when mobile phones are in the indoor environment, the detection accuracy is around 90%. When the phones are in the semioutdoor/outdoor environment, the detection accuracy is around 92%. Figure 11(b) shows the detection performance of the cellular detector that classifies the indoor environment from the semioutdoor/outdoor environment. We obtain quite a close performance of the cellular detector compared with that of the light detector. Our experiments mainly cover the campus and city areas where most sites are covered by at least five cell towers. In such experiment settings, the cell tower–based detection performs with 87% accuracy.

We note that both the light detector and cellular detector can effectively classify the indoor environment from the semioutdoor/outdoor environment. On the other hand, the magnetism detector can enhance the detection capability of IODetector in classifying the semioutdoor and outdoor environment. Figure 11(c) plots the performance of the magnetism detector. The magnetism detector can successfully distinguish the indoor and semioutdoor environments from the outdoor environment with an accuracy around 78%.

4.2.2. Performance of Aggregated IODetector. As described in Section 3.5, there are two approaches to constructively combine the results from the three subdetectors.

Detection accuracy. In Figure 12, we show the detection accuracy of both stateless and stateful IODetectors. We report the average detection results, including detection *precision* and *recall* [Junker et al. 1999], in different scenarios. The overall detection accuracy of the stateless IODetector is about 85%. When the three subdetectors are aggregated as the stateful IODetector, there is improvement of detection accuracy for all different types of indoor/outdoor environments, but not much. According to the

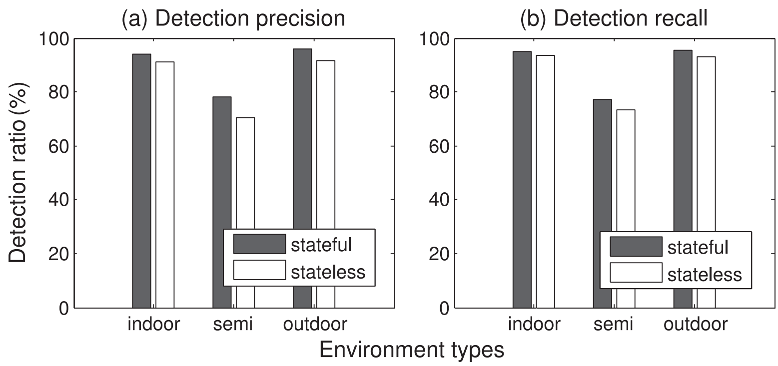

Fig. 12. Detection precision and recall of stateless and stateful IODetectors.

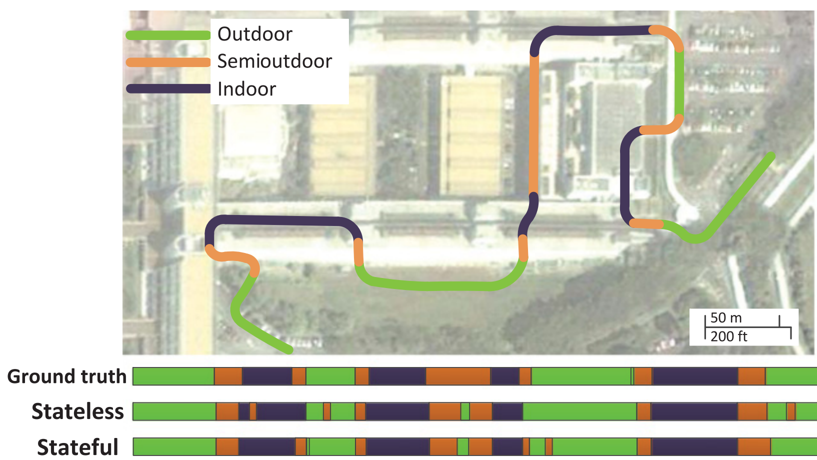

Fig. 13. An experiment trace in the university campus. The user uniformly walks along the path in 20 minutes, and the detection results are stored on the user's mobile phone. The ground truth data is manually collected.

experiment results, for both stateless and stateful IODetectors, the detection precision and the recall for the semioutdoor detection are much lower than the detection results for the other two environment types. Compared with the other two types, the semioutdoor environment is defined flexibly as "near a building," as shown in Figure 1. Different applications may have different definitions for it and can parameterize such environment type based on their specific requirements; the detection accuracy could then be improved accordingly. Compared with less than 82% detection accuracy of individual detectors, in the aggregated IODetector both the precision and the recall are consistently above 92% for the indoor and outdoor environment type. The experiment results suggest that IODetector accurately classifies the indoor/outdoor environments for most cases. For the stateful IODetector, with the optimization of the HMM parameters, the detection accuracy could be further improved.

In Figure 13, we show one of the walking traces that we experiment with at the NTU campus. The experiment was done on a rainy day. The detection results from stateless and stateful IODetectors can be seen in the bottom portion of Figure 13. The detection results of both IODetectors are accurate. When we look at their detection results separately, there are some differences. In some segments, the stateless IODetector

Table V. Battery Duration for Different Sensor Settings (in Hours)

Sensors	Samsung i9100	HTC Desire	HTC Sensation
No sensor	18.3	15.4	18.1
Magnetism 2Hz	18.0	14.9	17.8
Light 400Hz+FFT	17.8	15.0	17.5
Cell tower 2Hz	18.1	15.1	17.9
Stateful IODetector	17.4	14.3	17.1

suffers from misdetection of some semioutdoor environments, which are usually in the trace between indoor and outdoor environments. In some segments, although the ambient environment does not change, the detection result of the stateless IODetector may vary. The detection result of the stateful IODetector is relatively more stable due to the effect of the HMM. Considering the previous state, the HMM filters out some noise and avoids the misdetection of semioutdoor environments during user movements. However, the stateful IODetector may give inaccurate results for frequent environment changes, as it reacts insensitively to the sudden change of environment types and there are extra energy consumptions for the stateful IODetector due to its continuous operation.

Detection latency. The detection latency of IODetector is bounded by the time consumed by three subdetectors. The light detector is fast, sampling at 400Hz, which is sufficient to capture the alternating light intensity. We set the same detection window length of 10 seconds for both the cellular detector and magnetism detector. Considering that three detectors can run in parallel, it typically takes 10 seconds to warm up and then starts reporting detection results. After that, IODetector can keep tracking the indoor/outdoor transitions according to the application requirements.

System overhead. We measure the energy consumption of continuously sampling light sensor, magnetism sensor, and cellular signals. Table V shows the measured battery lifetime when the mobile phones continuously sample different sensors. The mobile phone's battery is fully charged before the measurement. During the measurement, we set the mobile phone screen brightness to the minimum brightness (but not turned off) while completely discharging the phone's battery. The experiments were done with the screen turned on, and we can expect the phone battery lifetime to be significantly longer if the screen were turned off. During the measurement, the mobile phone is in idle mode with cellular network connection. The setting "No sensor" means that the sensors (exclude the idle cellular module) used in our system are not triggered by other applications. In Table V, we find that the battery durations for sampling magnetism sensor at 2Hz and sampling light sensor at 400Hz with the FFT are quite close to the battery duration without sampling any sensors. Sampling the cellular signal consumes little extra battery power as well. Thus, although the stateful IODetector needs to perform continuous detection, the low energy consumption makes it affordable for the users.

4.3. Case Study: Inferring GPS Availability

In this section, we conduct a case study and demonstrate how IODetector can be used to provide indicative information on GPS availability. Today, many smartphones are equipped with commodity GPS modules that provide localization and navigation services for mobile applications. Traditional works [Thiagarajan et al. 2009] study how to adaptively use GPS/GSM/WiFi signals for energy-efficient localization or tracking. Such approaches, however, either assume the preknowledge of the ambient environment or infer it passively with high overhead and low efficiency. Serving as a generic and lightweight service, IODetector can be used to provide cheap and instant triggers for switching on/off the GPS component to achieve energy efficiency. Many research works [Burigat and Chittaro 2011; Zhang et al. 2011a] aim to improve GPS accuracy and are orthogonal to IODetector. IODetector tells the applications when GPS satellites are visible, and the applications can operate accordingly.

4.3.1. Indoor/Outdoor-Dependent Performance. For accurate localization, GPS normally needs clear line-of-sight paths to more than four GPS satellites. In the outdoor environment, with clear paths to a sufficient number of satellites, commodity GPS modules can achieve high localization accuracy within 20m. In the shadow of tall buildings, the line-of-sight paths to some satellites would be blocked and the mobile phone may only receive signals from a small number of satellites. Some received GPS signals might be from the reflecting walls leading to the multipath problem. In such scenarios, the localization accuracy degrades dramatically. In the indoor environment, there is normally no line-of-sight path to satellites. As a result, the mobile phone takes minutes to scan the satellites without finding any strong signals from satellites, and the localization error can be up to 400m or may not even get a location fix. In addition to the inaccuracy, it usually causes high responsive latency and extra power consumption. As GPS performance differs significantly in the indoor and outdoor environments, mobile phones greatly benefit from a priori knowledge of the ambient environment types.

4.3.2. GPS Availability and Localization Accuracy. We evaluate the localization accuracy and energy consumption of a mobile phone GPS module along with a walking path in our experiment. Figure 14(a) plots the experiment path on our campus. We mark the route segments from A to S. The total length of the walking path is approximately 1,600m, with 620m outdoor, 380m semioutdoor, and 600m indoor segments, respectively. We query the GPS for location information when we travel along the circular path with different mobile phone models under different weather conditions during the 1-month experiments.

Figure 14(b) plots the number of visible satellites as well as the signal to noise ratio (SNR) of the GPS signals in the indoor, semioutdoor, and outdoor environments, respectively. In the left-hand portion of Figure 14(b), we find that in the indoor environment, fewer than two GPS satellites are visible, although the mobile phones can sometimes capture slightly more GPS signals near windows. In the outdoor environment, mobile phones normally receive signals from more than six GPS satellites, even on cloudy and rainy days. The number of observed satellites varies in between in the semioutdoor environment (e.g., corridors and paths in the shadow of buildings). The right-hand portion of Figure 14(b) plots the SNR of the received GPS signals. The SNR value is a normalized value from the Android API indicating the SNR of the received satellite signal. An SNR greater than 20 is usually high enough for the mobile phone to get a location fix, and typically, the greater, the better. As well in the right-hand portion of Figure14(b), we observe that in the indoor environment, the SNR of GPS signals varies from 0 to 10. In the outdoor environment, the SNR becomes much higher, varying from 25 to 42 due to the clear line-of-sight paths between the phones and GPS satellites. In the semioutdoor environment, although we may sometimes observe GPS signals from more than four satellites, typically the SNR of GPS signals is not high enough to ensure accurate localization.

Figure 14(c) plots the summarized GPS localization error against the number of visible satellites. We find that the GPS modules can obtain more accurate localization results with more visible satellites. According to the experiment results, with fewer than 4 visible satellites, GPS service is generally unavailable. The GPS module is able to work with more than 4 visible satellites. However, even with 4 satellite signals, the localization accuracy varies dramatically in our experiment. With more than 6 visible satellites, the localization error is around 20m. We also observe that more visible

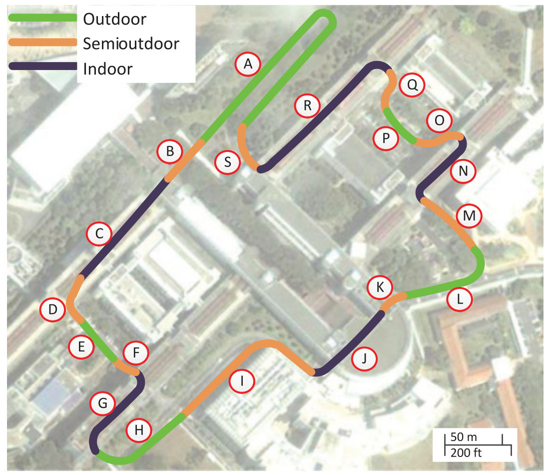

(a) GPS experiment trace in the university campus. We divide the trace into several segments based on the ground truth.

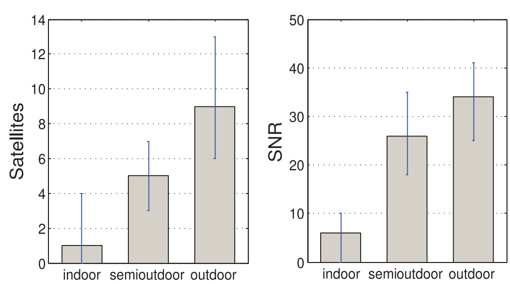

(b) Number of visible satellites and SNR of GPS signals in different environments. We see the number is related to the ambient environment type.

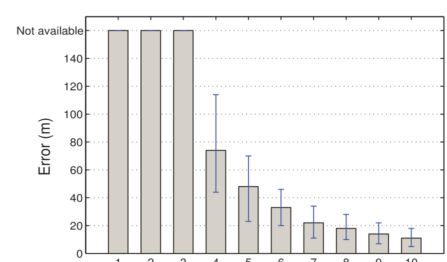

(c) GPS accuracy according to the number of available satellites More satellites usually result in more accurate localization.

Fig. 14. Indoor/outdoor-dependent GPS performance.

satellites (e.g., >9) yield less marginal improvements in the localization accuracy. With 10 GPS satellites, the localization error can be within 10m.

In summary, the experiment results demonstrate that GPS availability and localization accuracy are highly correlated to the environment types. Yet solely reading such availability from the GPS module itself can be up to minutes and consume much extra energy in scanning the satellites.

4.3.3. IODetector-Augmented GPS: IO-GPS. We can simply leverage IODetector to infer the GPS with accurate indoor/outdoor awareness. In our IOdetector-augmented GPS (IO-GPS) scheme, mobile applications invoke IODetector for the indoor/outdoor detection before switching on the GPS module. If the mobile phone is outdoors, the applications can confidently call the GPS for an outdoor localization; if it is indoors, the applications may postpone the GPS localization and resort to a variety of alternative indoor localization techniques [Yang et al. 2012]. In this experiment, we track the localization accuracy and energy consumption of the traditional GPS and the IO-GPS scheme.

IO-GPS localization accuracy. We follow the path in Figure 14(a) at a walking speed and collect the GPS localization data. We repeat the experiment 10 times and report

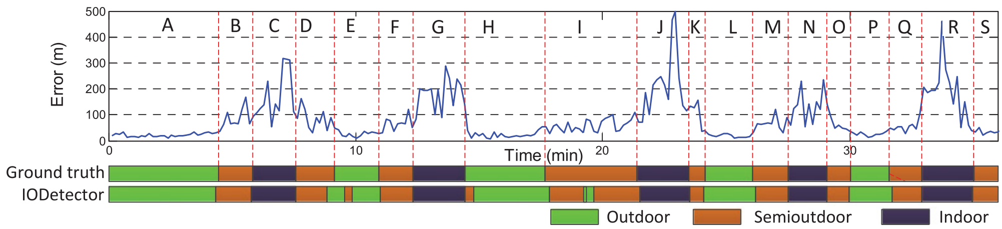

Fig. 15. GPS localization accuracy of an example instance along the walking path. The user uniformly walks along the path, and the environment detection results are stored on the user's mobile phone. The ground truth data is collected manually, and the GPS localization error is calculated offline.

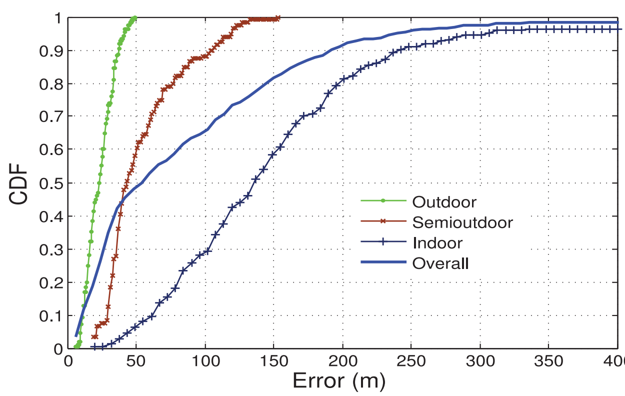

Fig. 16. CDF of localization error.

the average results. We use the stateful IODetector to estimate the environment type. Figure 15 presents one example instance. In this figure, we observe that the GPS localization error varies across different path segments. We see apparent variation on GPS localization error due to the indoor/outdoor environment transition. For example, when we move from segment G to H (indoors → outdoors), we see a big dive of the localization error; when we move from P to Q and then to R (outdoors → semioutdoors → indoors), we observe a two-stage jump of the localization error. Consistent with the preceding measurement, in the indoor environment the GPS localization error is much larger than that in the semioutdoor or outdoor environment. The path segment J has a particularly high error because the segment is underground and the GPS component detects almost no satellite signals.

Figure 16 summarizes the localization error from the 10 experiments, and we take a fine look at the localization error in outdoor, semioutdoor, and indoor areas, respectively. The median localization error in the outdoor environment is around 24m with the maximum error within 50m in our experiments. In the semioutdoor areas, the median error is around 44m, whereas the 90th percentile can be up to 100m. In the indoor environment, the median localization error is around 140m with the 90th percentile of 235m. The overall localization error presents the performance of the traditional GPS. We find that the median localization error is around 55m with a long tail up to 400m. In our experiment setting, users walk around campus with comparable route segments inside buildings and in outdoor environments. Yet the research on human activity patterns show that people spend around 89% of the time in the indoor environment [Klepeis et al. 2001].

Table VI. Battery Consumption Comparison (in Hours)

Environment	Samsung i9100	HTC Desire	HTC Sensation
Indoor GPS	9.2	6.6	8.7
Semioutdoor GPS	9.8	7.2	9.7
Outdoor GPS	10.1	7.3	9.8
Overall GPS	9.8	7.2	8.9
IO-GPS	15.3	12.7	14.5

Without discriminating the indoor/outdoor environment, blindly using the traditional GPS scheme would perform similarly to that in indoor cases most of the time. Augmented by IODetector, the IO-GPS performance would be closer to that in the outdoor/semioutdoor environment. In the bottom portion of Figure 15, we compare the indoor/outdoor detection results with the ground truth along the experiment path. We see that IODetector provides promising detection accuracy. In particular, IODetector successfully detects the indoor cases from the semioutdoor and outdoor cases. For outdoor/semioutdoor detection at some places, IODetector cannot provide the most accurate result. We revisit places such as the path segments D and E, where IODetector misclassifies the semioutdoor and outdoor environments. We find that D and E are located at a corner passing by a two-storey building. It is even difficult to manually label such places as ground truth, yet we believe that the misclassification results of IODetector in such corner cases would introduce little influence to the GPS localization service.

Energy consumption. We measure the power consumption when we run the GPS module during the experiment. We measure the battery life with the screen set to the minimum brightness. Table VI summarizes the battery life of three different mobile phone models in different environments. We also present the battery duration for running the stateful IODetector for the indoor/outdoor detection. In Table VI, the first three rows show the energy consumption of mobile phones when the GPS is turned on for indoor, semioutdoor, and outdoor environments. We also measure the energy consumption of “Overall GPS” and IO-GPS. In IO-GPS, the GPS module is turned off when the user is in indoor environments. In the “Overall GPS” case, the GPS module is always on regardless of the ambient environment type. We find that the GPS drains the battery rapidly in all environments. The energy consumption of the GPS is especially high in the indoor environment, where the GPS module continuously scans the satellite signals and rapidly depletes the battery energy. IO-GPS can significantly save the battery power that was wasted before because the mobile phone cannot sense and receive enough GPS satellite signals to get a location fix. With the awareness of the indoor/outdoor environment, IO-GPS avoids unnecessarily switching on the GPS module and saves the energy consumption in the indoor environment.

5. RELATED WORK

Although there have not yet been generic approaches proposed for explicit indoor/outdoor detection, there exists a wide body of related works that implicitly deal with such a problem.

Environment detection. GPS lock status can be used to indirectly infer the ambient environment [Ravindranath et al. 2011], but it usually incurs substantial energy cost and high latency. ABL [Lane et al. 2007] proposes the approach that allows mobile sensors to localize themselves by exploiting their ambient physical environment signals. FLIGHT [Li et al. 2012] explores the fact that the light intensity changes with a stable period in the indoor environment and uses the feature to perform clock calibration. TempIO [Krumm and Hariharan 2004] classifies the ambient environment by comparing the environment temperature with the current outdoor temperature through the network query. Yet temperature sensors are not widely available on current mobile phones. Along with many other sensing recourses, the temperature sensor, if available on mobile phones, can be used to complement our work. TagSense [Qin et al. 2011] classifies the ambient environments to automatically annotate images during the picture click. Some works in image processing and pattern recognition [Payne and Singh 2005; Szummer and Picard 1998] also study the problem of classifying images according to ambient environments. Those works can provide partial indication of indoor/outdoor environment. As taking photos normally incurs substantial human effort and energy cost, we can hardly rely on such classification approaches to build a generic and automatic indoor/outdoor detection service.

Localization and tracking. Many works study GPS/GSM/WiFi localization schemes. StarTrack [Ananthanarayanan et al. 2009] provides a comprehensive set of APIs for the development of mobile localization and tracking applications. Zhou et al. [2012a] use cell tower sequences to track the buses and make bus arrival time predictions for the waiting passengers. LANDMARC [Ni et al. 2004] proposes a location sensing prototype system that uses RFID technology for locating objects inside buildings. EnTracked [Kjærgaard et al. 2009] focuses on outdoor pedestrian tracking using a lightweight accelerometer to trigger GPS to reduce power consumption. Jurdak et al. [2010] complement GPS duty cycling with short-range radio contacts to balance positioning accuracy and energy consumption. VTrack [Thiagarajan et al. 2009] studies reducing energy consumption by using inaccurate WiFi positioning schemes to measure road traffic conditions. Chung et al. [2011] present an accurate positioning system based on the magnetic signatures in the indoor environment. Liu et al. [2012] design a cloud-offloaded GPS solution that allows a sensing device to merely collect little GPS signal for postprocessing with cloud computation. Such a method may lower the energy consumption of GPS sensing and make GPS-based localization more energy efficient. The proposed approach suffers substantial detection delay that may not be feasible for real-time applications. Walkie-Markie [Shen et al. 2013] is an indoor pathway mapping system that is based entirely on crowdsourcing and leverages user trajectories obtained from ordinary pedestrians and their mobile phones.

The preceding approaches primarily focus on obtaining accurate physical locations and track the targeted objects. They can potentially benefit from the indoor/outdoor awareness of IODetector—that is, adaptively switching on/off the GPS modules in localization.

Context awareness and activity recognition. A number of works have studied the use of sensors to recognize user activities and detect ambient context. Yan et al. [2012] design and build FALCON to remedy slow app launch using contexts to predict the next app to launch. CenceMe [Miluzzo et al. 2008] exploits sensors on mobile phones to automatically infer people's ambient context and then allows users to share that through social networks. Mercury [Lorincz et al. 2009] monitors patients using wearable sensors in indoor medical environments. EEMSS [Wang et al. 2009] presents an energy-efficient sensor management framework that uses a minimum number of sensors on mobile devices to monitor user status. Jigsaw [Lu et al. 2010] supports continuous sensing applications on mobile phones to infer human activities and ambient context. PBN [Keally et al. 2011] proposes a user activity detection system using sensors on both mobile phones and on-body wireless sensors. Such works either implicitly assume the activity context or passively infer the ambient context. Unlike those works, our work proactively detects the indoor/outdoor environment using various lightweight sensors (e.g., light sensor, cellular signal, and magnetism sensor) without any remote supports.

Christodoulidis et al. [2012] describe a software video processing and analysis system to assist the near real-time detection of human activity. SoundSense [Lu et al. 2009] classifies general sound types (e.g., music, voice) to achieve context recognition. SensLoc [Kim et al. 2010] collects WiFi beacons to extract useful patterns to infer contextual information. Kobe [Chu et al. 2011] aids the mobile classifier development by automatically extracting high-level semantics from raw sensory data while balancing energy, latency, and accuracy. Our work primarily differs from them in that IODetector instantly detects the primitive ambient context without any labor-intensive site survey and user feedback. Those works may benefit from IODetector by taking the indoor/outdoor information as a primary filter for context recognition.

6. CONCLUSIONS

We present the design and implementation of an indoor/outdoor environment detection system that efficiently takes input from a variety of lightweight sensors to derive the indoor/outdoor information. By intelligently aggregating the subdetectors, IODetector achieves prompt and accurate detection results in various times and environments. We comprehensively test IODetector through a prototype implementation and evaluate the system based on different Android mobile phone models. We particularly conduct a case study where we make use of IODetector results to infer the GPS availability and accuracy under various indoor/outdoor environments.

REFERENCES

Ganesh Ananthanarayanan, Maya Haridasan, Iqbal Mohomed, Doug Terry, and Chandramohan A. Thekkath. 2009. StarTrack: A framework for enabling track-based applications. In *Proceedings of ACM MobiSys*. 207–220.

Martin Azizyan, Ionut Constandache, and Romit Roy Choudhury. 2009. SurroundSense: Mobile phone localization via ambience fingerprinting. In *Proceedings of ACM MobiCom*. 261–272.

Paramvir Bahl and Venkata N. Padmanabhan. 2000. RADAR: An in-building RF-based user location and tracking system. In *Proceedings of IEEE INFOCOM*. 775–784.

Stefano Burigat and Luca Chittaro. 2011. Pedestrian navigation with degraded GPS signal: Investigating the effects of visualizing position uncertainty. In *Proceedings of ACM MobileHCI*. 221–230.

Argiris Christodoulidis, Konstantinos K. Delibasis, and Ilias Maglogiannis. 2012. Near real-time human silhouette and movement detection in indoor environments using fixed cameras. In *Proceedings of ACM PETRA*. Article No. 1 DOI:http://dx.doi.org/10.1145/2413097.2413099

David Chu, Nicholas D. Lane, Ted Tsung-Te Lai, Cong Pang, Xiangying Meng, Qing Guo, Fan Li, and Feng Zhao. 2011. Balancing energy, latency and accuracy for mobile sensor data classification. In *Proceedings of ACM SenSys*. 54–67.

Jaewoo Chung, Matt Donahoe, Chris Schmandt, Ig-Jae Kim, Pedram Razavai, and Micaela Wiseman. 2011. Indoor location sensing using geo-magnetism. In *Proceedings of ACM MobiSys*. 141–154.

Sibren Isaacman, Richard Becker, Ramón Cáceres, Margaret Martonosi, James Rowland, Alexander Varshavsky, and Walter Willinger. 2012. Human mobility modeling at metropolitan scales. In *Proceedings of ACM MobiSys*. 239–252.

Markus Junker, Rainer Hoch, and Andreas Dengel. 1999. On the evaluation of document analysis components by recall, precision, and accuracy. In *Proceedings of ICDAR*. 713–716.

Raja Jurdak, Peter Corke, Dhinesh Dharman, and Guillaume Salagnac. 2010. Adaptive GPS duty cycling and radio ranging for energy-efficient localization. In *Proceedings of ACM SenSys*. 57–70.

Matthew Keally, Gang Zhou, Guoliang Xing, Jianxin Wu, and Andrew Pyles. 2011. PBN: Towards practical activity recognition using smartphone-based body sensor networks. In *Proceedings of ACM SenSys*. 246–259.

Donnie H. Kim, Younghun Kim, Deborah Estrin, and Mani B. Srivastava. 2010. SensLoc: Sensing everyday places and paths using less energy. In *Proceedings of ACM SenSys*. 43–56.

Mikkel Baun Kjærgaard, Jakob Langdal, Torben Godsk, and Thomas Toftkjær. 2009. EnTracked: Energy-efficient robust position tracking for mobile devices. In *Proceedings of ACM MobiSys*. 221–234.

Neil E. Klepeis, William C. Nelson, Wayne R. Ott, John P. Robinson, Andy M. Tsang, Paul Switzer, Joseph V. Behar, Stephen C. Hern, and William H. Engelmann. 2001. The national human activity pattern survey

(NHAPS): A resource for assessing exposure to environmental pollutants. *Journal of Exposure Analysis and Environmental Epidemiology* 11, 3, 231–252.

John Krumm and Ramaswamy Hariharan. 2004. TempIO: Inside/outside classification with temperature. In *Proceedings of the 2nd Workshop on Man-Machine Symbiotic Systems*.

Nicholas D. Lane, Hong Lu, and Andrew T. Campbell. 2007. Ambient beacon localization: Using sensed characteristics of the physical world to localize mobile sensors. In *Proceedings of the 4th Workshop on Embedded Networked Sensors*. 38–42.

Nicholas D. Lane, Emiliano Miluzzo, Hong Lu, Daniel Peebles, Tanzeem Choudhury, and Andrew T. Campbell. 2010. A survey of mobile phone sensing. *IEEE Communications Magazine* 48, 9, 140–150.

Zhenjiang Li, Wenwei Chen, Cheng Li, Mo Li, Xiang-Yang Li, and Yunhao Liu. 2012. FLIGHT: Clock calibration using fluorescent lighting. In *Proceedings of ACM MobiCom*. 329–340.

Robert LiKamWa, Bodhi Priyantha, Matthai Philipose, Lin Zhong, and Paramvir Bahl. 2013. Energy characterization and optimization of image sensing toward continuous mobile vision. In *Proceedings of ACM MobiSys*. 69–82.

Jie Liu, Bodhi Priyantha, Ted Hart, Heitor S. Ramos, Antonio A. F. Loureiro, and Qiang Wang. 2012. Energy efficient GPS sensing with cloud offloading. In *Proceedings of ACM SenSys*. 85–98.

Siyuan Liu, Yunhuai Liu, Lionel M. Ni, Jianping Fan, and Minglu Li. 2010. Towards mobility-based clustering. In *Proceedings of KDD'10*. 919–928.

Siyuan Liu, Yisong Yue, and Ramayya Krishnan. 2013. Adaptive collective routing using Gaussian process dynamic congestion models. In *Proceedings of KDD'13*. 704–712.

Konrad Lorincz, Bor-rong Chen, Geoffrey Werner Challen, Atanu Roy Chowdhury, Shyamal Patel, Paolo Bonato, and Matt Welsh. 2009. Mercury: A wearable sensor network platform for high-fidelity motion analysis. In *Proceedings of ACM SenSys*. 183–196.

Hong Lu, Wei Pan, Nicholas D. Lane, Tanzeem Choudhury, and Andrew T. Campbell. 2009. SoundSense: Scalable sound sensing for people-centric applications on mobile phones. In *Proceedings of ACM MobiSys*. 165–178.

Hong Lu, Jun Yang, Zhigang Liu, Nicholas D. Lane, Tanzeem Choudhury, and Andrew T. Campbell. 2010. The Jigsaw continuous sensing engine for mobile phone applications. In *Proceedings of ACM SenSys*. 71–84.

Emiliano Miluzzo, Nicholas D. Lane, Kristóf Fodor, Ronald Peterson, Hong Lu, Mirco Musolesi, Shane B. Eisenman, Xiao Zheng, and Andrew T. Campbell. 2008. Sensing meets mobile social networks: The design, implementation and evaluation of the CenceMe application. In *Proceedings of ACM SenSys*. 337–350.

Lionel M. Ni, Yunhao Liu, Yiu Cho Lau, and Abhishek P. Patil. 2004. LANDMARC: Indoor location sensing using active RFID. *ACM Wireless Networks* 10, 6, 701–710.

Andrew Payne and Sameer Singh. 2005. Indoor vs. outdoor scene classification in digital photographs. *Pattern Recognition* 38, 10, 1533–1545.

Chuan Qin, Xuan Bao, Romit Roy Choudhury, and Srihari Nelakuditi. 2011. TagSense: A smartphone-based approach to automatic image tagging. In *Proceedings of ACM MobiSys*. 1–14.

Lenin Ravindranath, Calvin Newport, Hari Balakrishnan, and Samuel Madden. 2011. Improving wireless network performance using sensor hints. In *Proceedings of USENIX NSDI*. 21.

Anthony Rowe, Vikram Gupta, and Ragunathan (Raj) Rajkumar. 2009. Low-power clock synchronization using electromagnetic energy radiating from AC power lines. In *Proceedings of ACM SenSys*. 211–224.

Guobin Shen, Zhuo Chen, Peichao Zhang, Thomas Moscibroda, and Yongguang Zhang. 2013. Walkie-Markie: Indoor pathway mapping made easy. In *Proceedings of USENIX NSDI*. 85–98.

SingTel. 2013. Home Page. Retrieved August 17, 2014, from http://info.singtel.com/.

Adam Smith, Hari Balakrishnan, Michel Goraczko, and Nissanka Priyantha. 2004. Tracking moving devices with the cricket location system. In *Proceedings of ACM MobiSys*. 190–202.

Martin Szummer and Rosalind W. Picard. 1998. Indoor-outdoor image classification. In *Proceedings of the IEEE Workshop on Content-Based Access of Image and Video Databases*. 42–51.

China Telecom. 2012. Home Page. Retrieved August 17, 2014, from http://en.chinatelecom.com.cn.

Arvind Thiagarajan, James Biagioni, Tomas Gerlich, and Jakob Eriksson. 2010. Cooperative transit tracking using smart-phones. In *Proceedings of ACM SenSys*. 85–98.

Arvind Thiagarajan, Lenin Ravindranath, Katrina LaCurts, Samuel Madden, Hari Balakrishnan, Sivan Toledo, and Jakob Eriksson. 2009. VTrack: Accurate, energy-aware road traffic delay estimation using mobile phones. In *Proceedings of ACM SenSys*. 85–98.

Nishith D. Tripathi, Jeffrey H. Reed, and Hugh F. VanLandingham. 1998. Handoff in cellular systems. *IEEE Personal Communications* 5, 6, 26–37.

Yi Wang, Jialiu Lin, Murali Annavaram, Quinn A. Jacobson, Jason Hong, Bhaskar Krishnamachari, and Norman Sadeh. 2009. A framework of energy efficient mobile sensing for automatic user state recognition. In *Proceedings of ACM MobiSys*. 179–192.

Wikipedia. 2013a. Earth's Magnetic Field. Retrieved August 17, 2014, from http://en.wikipedia.org/wiki/Earth_magnetic_field.

Wikipedia. 2013b. Lux. Retrieved August 17, 2014, from http://en.wikipedia.org/wiki/Lux.

Tingxin Yan, David Chu, Deepak Ganesan, Aman Kansal, and Jie Liu. 2012. Fast app launching for mobile devices using predictive user context. In *Proceedings of ACM MobiSys*. 113–126.

Zheng Yang, Chenshu Wu, and Yunhao Liu. 2012. Locating in fingerprint space: Wireless indoor localization with little human intervention. In *Proceedings of ACM MobiCom*. 269–280.

Xiaohong Zhang, Xingxing Li, and Fei Guo. 2011a. Satellite clock estimation at 1 Hz for realtime kinematic PPP applications. *GPS Solutions* 15, 4, 315–324.

Zengbin Zhang, Xia Zhou, Weile Zhang, Yuanyang Zhang, Gang Wang, Ben Y. Zhao, and Haitao Zheng. 2011b. I am the antenna: Accurate outdoor AP location using smartphones. In *Proceedings of ACM MobiCom*. 109–120.

Pengfei Zhou, Yuanqing Zheng, and Mo Li. 2012a. How long to wait?: Predicting bus arrival time with mobile phone based participatory sensing. In *Proceedings of ACM MobiSys*. 379–392.

Pengfei Zhou, Yuanqing Zheng, Zhenjiang Li, Mo Li, and Guobin Shen. 2012b. IODetector: A generic service for indoor outdoor detection. In *Proceedings of ACM SenSys*. 113–126.

Received February 2013; revised May 2014; accepted June 2014