Sensing WiFi Packets in the Air: Practicality and Implications in Urban Mobility Monitoring

Yohan Chon, Suyeon Kim, Seungwoo Lee, Dongwon Kim, Yungeun Kim, Hojung Cha
Department of Computer Science, Yonsei University, Seoul, Korea
{yohan,sykim,swlee,dwkim,ygkim,hjcha}@cs.yonsei.ac.kr

ABSTRACT

Mobile sensing systems employ various sensors in smartphones to extract human-related information. As the demand for sensing systems increases, a more effective mechanism is required to sense information about human life. In this paper, we present a systematic study on the feasibility and gaining properties of a crowdsensing system that primarily concerns sensing WiFi packets in the air. We propose that this method is effective for estimating urban mobility by using only a small number of participants. During a seven-week deployment, we collected smartphone sensor data, including approximately four million WiFi packets from more than 130,000 unique devices in a city. Our analysis of this dataset examines core issues in urban mobility monitoring, including feasibility, spatio-temporal coverage, scalability, and threats to privacy. Collectively, our findings provide valuable insights to guide the development of new mobile sensing systems for urban life monitoring.

Author Keywords

Crowdsensing; Smartphone Sensing; WiFi Monitor Mode

ACM Classification Keywords

H.4.m. Information Systems Application: Miscellaneous

General Terms

Design; Experimentation; Measurement; Performance

INTRODUCTION

The proliferation of smartphones keeps people constantly tethered to their phones even while walking, driving, eating, and having a conversation. This activity constantly generates network packets from smartphones, since 82% of phone usage is related to Internet use such as communications, web browsing, watching media content, and online gaming [10]. People typically use Internet applications through WiFi or cellular interface. If mobile users are within WiFi coverage, they prefer redirecting data traffic through WiFi networks in order to reduce data charges. Users often seek public WiFi access points (APs), and smartphones are typically configured to scan WiFi APs automatically. Consequently, smartphones continuously send WiFi packets to surrounding environments, whatever the user may do with them.

A WiFi packet contains a media access control (MAC) address assigned to network interfaces that uses the IEEE 802 network protocols. The MAC address can be used as an identity of the smartphone user, since the address is a unique identifier of each mobile device and people always carry their smartphones. Considering periodic packet transmission from smartphones, we can say that mobile users continuously notify the surrounding environments of their existence (i.e., MAC address). Conceptually, the WiFi packet is expressed as an "*I'm here!*" message from mobile users. This is a basic mechanism in wireless communication that is commonly used today. However, we argue that WiFi packets in the air may not only derive precious context information about urban mobility, but also threaten the privacy of mobile users.

Our study originated from one interesting idea: WiFi packets are everywhere and can be used to identify mobile users, then *what happens if mobile sensing systems sense ambient WiFi packets in the air?* Mobile sensing research communities employ various sensors in smartphones to extract human-related information such as location, activity, or environment. What would be an additional gain if we monitor surrounding WiFi packets by crowdsensing? Would we extract richer contexts about mobile users? Would it be useful for service providers or individual users? In a way, we can infer the identity of mobile users from the MAC address in WiFi packets. Then, would it invade an individual's privacy? How harmful would this be to mobile users? Should we design a security scheme to prevent such threats in using WiFi?

In this paper, we report on the systematic study of sensing WiFi packets by smartphones. A key aim of our study is to demystify the potential and the threat that exist in sensing ambient packets by smartphone-based crowdsensing. We enable the *WiFi monitor mode* in commodity smartphones to sense WiFi packets, although the packets are not associated with a user's phone. We validate that this method enables the estimation of urban mobility (i.e., large-scale user movements in a city) by using only a small number of participants. In previous work on mobility monitoring [13, 27, 22, 5], the collected data from 100 users estimates the mobility of exactly 100 users, since the system focuses on the mobility of the phone owner. With the use of the WiFi monitor mode, we found that the coverage of mobility monitoring is greatly increased, since the method collects the mobility of the phone

Ubicomp '14, September 13-17 2014, Seattle, WA, USA

http://dx.doi.org/10.1145/2632048.2636066

owner as well as the surrounding users. For example, the mobile sensing from 100 participants may enable the mobility monitoring of a few thousands users in a city. This advantage makes urban mobility monitoring feasible, since the scheme greatly reduces the burden to motivate user participation.

To validate our idea, we deployed a mobility monitoring system into 25 university students and enabled the WiFi monitor mode in smartphones to sense WiFi packets. The phones periodically sensed WiFi packets in the air as well as their locations. During seven weeks of deployment, we collected a dataset including approximately four million WiFi packets from more than 130,000 unique devices, making 8,434 place visits in Seoul, Korea. We use this dataset to validate the major advantage of sensing ambient packets: the system can monitor the mobility of a large-scale number of people from the data collected by a small number of users. We additionally investigated fundamental issues about sensing the WiFi packets, including feasibility–examining the cost and potential of packet sensing; coverage and scalability–estimating coverage of mobility monitoring in a city; and privacy concerns–considering privacy invasion by packet sensing.

This paper makes the following contributions:

- We conducted a large-scale study of sensing WiFi packets in the air by smartphone-based crowdsensing.
- We validated that the packet sensing is effective to estimate urban mobility via only a small number of participants.
- We analyzed diverse issues related to sensing WiFi packets such as feasibility, coverage, scalability, and threats to privacy in spreading WiFi packets in use today.

PRELIMINARIES AND MOTIVATION

The nomadic characterization of urban mobility is a fundamental resource for various applications such as traffic estimation [20], urban planning [26], and context-aware local search [25]. The urban mobility could be estimated by cameras on roadsides or by the traffic on cell towers, but the method requires costly infrastructures or provides coarse-grained mobility (i.e., several kilometers). Smartphone-based WiFi packet sensing can enlarge the coverage of mobility estimation, since mobile users could provide information by device-side crowdsensing manners. In addition, the packet sensing provides the fine granularity (i.e., a few meters) by using the WiFi fingerprint. These advantages expand the domain of applications from the coarse-grained outdoor regions to fine-grained indoor places. In what follows, we elaborate on motivations by describing the cost and the potential of monitoring WiFi packets in the air.

WiFi Packets from Smartphones

Our main concept is that the users' smartphones monitor the WiFi packets that are transmitted from the surrounding smartphones in daily life. Thus, the system can collect a large amount of data if (1) the user usually sets WiFi on (i.e., WiFi usage time), and (2) the smartphone frequently transmits WiFi packets (i.e., frequency of packet transmission).

Sensing WiFi packets is effective if mobile users turn on WiFi in casual daily life. To explore the WiFi usage time, we deployed a simple application to 25 smartphone users for one week in order to record the WiFi usage times and the timestamps of WiFi scanning in daily life. On average, the participants turned on WiFi for 12.4±5.9 hours and connected to APs for 8.5±6.1 hours in a day. The reason of this high usage time is to seek WiFi APs in order to redirect data traffic through a WiFi network. We then investigated how frequently smartphones transmit WiFi packets. The results from 11 smartphone models show that 82% of smartphones scanned the surrounding WiFi APs every 60 seconds at the most (130 seconds for 90% of cases). These results indicate that our method could capture the surrounding users if two smartphones are located within WiFi communication range for at least 60 seconds.

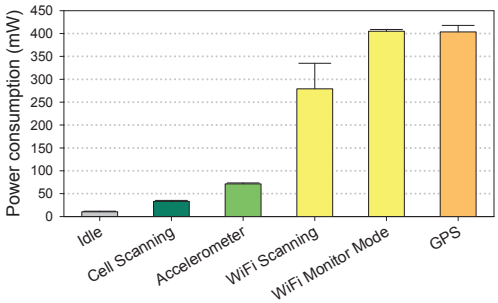

Figure 1. Energy profile of several sensors in Nexus One. All measurements include energy consumption of CPU for data reading and storing. The WiFi monitor mode consumes similar energy with GPS.

Energy Consumption of WiFi Monitor Mode

In what follows, we explore the overhead of WiFi monitor mode. We use the WiFi monitor mode to sense WiFi packets, although the packets are not associated with a user's phone. If the WiFi monitor mode significantly drains the battery lifetime, its usage is not feasible in practice. We measured the average power consumption of WiFi monitor modes and other sensor usage such as accelerometer, cell, WiFi scanning, and GPS. To maintain the measurement consistency, we used the Monsoon power monitor to deliver a constant voltage of 4.0V and halted all applications except for a system service. Figure 1 shows the energy consumption of several sensors compared with the activation of the WiFi monitor modes in the Nexus One. The scanning interval of WiFi sensing is 10 seconds, and we used the lowest sampling rate with 10% duty cycle for the accelerometer. The WiFi monitor mode (404.9mW) consumes more energy than WiFi scanning (279.2mW) and it is similar to the consumption of GPS sensing (403.5mW). Similar to GPS, the WiFi monitor mode drains a fully charged 1500mAh battery within 14.8 hours. Considering that GPS is commonly used in commercial smartphones, the WiFi monitor mode is a viable scheme, but it requires an adaptive sensing scheduling to preserve battery lifetime.

Potential of Sensing WiFi Packets

We explored the number of unique devices observed by a user in a day. The high number of observed devices indicates that the system has a great potential to estimate large-scale mobility despite a small number of data collectors. Figure 2(a) shows the trace of the devices observed in a day. Intuitively, high peaks appeared in moments when the user is moving

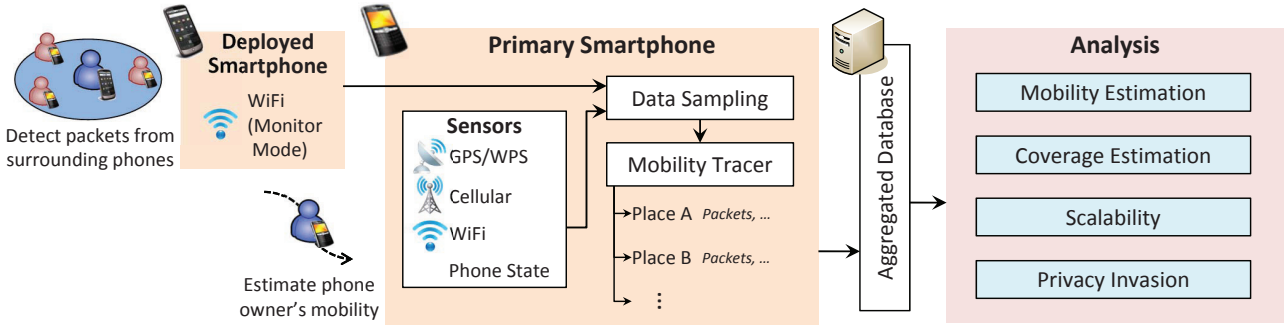

Figure 3. Overview of our study. The system utilized a representative sensing scheme to trace a user's mobility and enabled WiFi monitor mode to capture the ambient WiFi packets from surrounding phones.

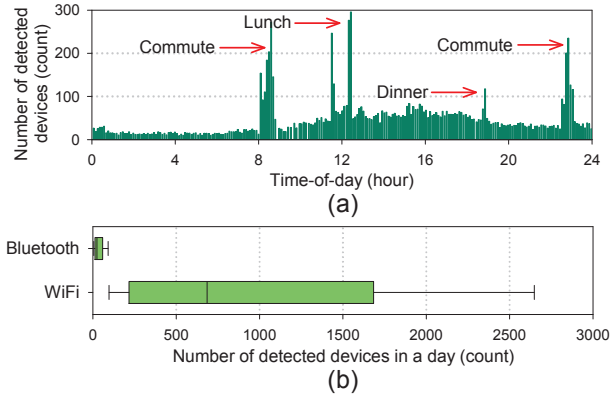

Figure 2. (a) Diurnal distribution of sensed packets by WiFi monitor mode. High peaks appeared at commute, lunch, and dinner time; (b) Box plot of the number of observed unique devices in a day by Bluetooth and WiFi. Box indicates lower quartile, median, upper quartile and whisker indicates 10% and 90% of observation. WiFi captured significantly higher number of devices than Bluetooth.

(e.g., commute route) or in crowded places (e.g., restaurants for lunch). The user detected more devices during working hours than at night. Figure 2(b) shows the number of devices observed for three business days by three university students and by three researchers. The number of observed devices varies from 30 to 3,100 (the average is 1,019 devices and the median is 685) per day due to the varied life patterns of the participants. The results show that one user encountered approximately 685 unique devices in a day within WiFi communication range. Previous work [9, 19] commonly used Bluetooth to detect surrounding smartphones, but we found that the participants detected an average of 28 unique devices in a day by Bluetooth. The reason is that Bluetooth has a shorter communication range than WiFi, and the device is discoverable only at the moment a user enabled the Bluetooth scanning. The alternative method is the use of cellular information such as data packets or call detail records. Sensing the cellular packet may detect significantly higher number of devices than the WiFi packet. However, the method requires specialized devices and the range is too wide (a few kilometers) to estimate the place-level granularity.

Consequently, the preliminaries reveal that the WiFi environment has great potential for monitoring urban mobility without modifying the smartphones of surrounding users. The WiFi-enabled smartphone periodically transmits the WiFi packets with a MAC address, and people living in urban areas contact numerous smartphones within a radio communication range during their daily lives.

STUDY DESIGN

We now describe the design of our study and the system we deployed for estimating urban mobility.

Overview. Figure 3 illustrates the overview of our study. We utilized a sensing scheme in the literature [5, 13] to estimate a user's mobility and added a feature to sense WiFi packets from surrounding smartphones. Participants installed a smartphone application that traces his/her mobility using a cellular module, WiFi, WPS, and GPS. They were additionally given the Nexus One smartphone or the WiFi module (i.e., Netgear WG111v2) to sense the ambient WiFi packets. The server infrastructure collects, stores, and analyzes data received from participants. All uploaded data, along with WiFi packets, were segmented into places by using WiFi fingerprints and then stored for analysis. We analyzed the core issues in mobility monitoring by sensing WiFi packets, such as feasibility, coverage, scalability, and threats to privacy.

Participants were told that the system's aim was to gather the mobility information of a participant as well as the surrounding users. We explained the mechanism of communicating WiFi packets in commercial smartphones. We requested that participants carry the WiFi monitor mode-enabled phone or the WiFi module with them at all times, keeping the device charged and powered on for most of the day.

Recruitment. To recruit the participants, we advertised via the university website and placed posters around campus. The interested students visited the lab and received a comprehensive description of the experiment. The participants received a payment of 100 USD for seven weeks of data collection, and they were allowed to quit the experiment at any time.

Protection of Human Subjects. We followed the policy of the National Research Foundation, which carefully examined this study's design. Our recruitment process and the policy regarding the experiment were approved by the Institutional Review Board (No.1040917-201312-SB-130-02). We did not link collected data with any information that relies on the users identity, and the MAC addresses of mobile devices were hashed for anonymization. The WiFi monitor mode is implemented differently from the sniffing mode: the content

of the WiFi packets was filtered out at the firmware, and only the MAC address of source and destination, signal strength, channel, and packet types were recorded. Participants were free to remove any collected data, and the data were accessible only to researchers who were part of the project.

Deployment Location. We conducted our study at Yonsei University in Seoul, Korea. The University has 18,000 undergraduate students, 11,000 graduate students, and 5,000 faculty members on campus. The University is located in *West-Gate* district, where the settled population is 387,000. Seoul is the largest metropolis in Korea, with a population of 10,413,000 [24]. The WiFi network is available almost everywhere in Seoul, even on the subway. In Seoul, 67% of people use smartphones, which is one of the highest adaption rates in the world [21].

Mobility Estimation

We estimate a user's *place-centric mobility*, defined as a location is recognized with room-level accuracy in indoor as well as outdoor environments; the movement patterns are traced for an entire day. This *place-centric* data is widely considered to provide place-related information where users spent the majority of their time [14, 7]. Our study focuses on the mobility of phone owners as well as surrounding users. This issue was not considered in previous work about mobility monitoring [13, 5, 27, 8].

Sensing Scheduling. We adopt the sensor sampling policy proposed in [5]; this cell-based scheduling is based on a user's everyday routine mobility patterns. The basic idea is that a high-power sensor (i.e., WiFi, GPS) is activated only if the data from a low-power sensor (i.e., cell-tower connections) show strong signs about location change. Thus, WiFi scanning occurs only when the user is likely to be in motion or visiting new locations. When a user follows previously observed mobility patterns, the scheme re-uses the stored location information without activating high-power sensors. To sense the WiFi packets, we periodically activate the WiFi monitor mode with a fixed time interval (i.e., two minutes) for collecting full coverage of data. The adaptive sensing scheduling of WiFi monitor mode is beyond the scope of our paper.

Place Detection. We used WiFi fingerprint matching to recognize logical places (such as a store or a user's home) and revisited places. The basic operation is that if a user is stationary, the signal fingerprints of surrounding APs are relatively similar to each other. Whenever a WiFi fingerprint is encountered that is previously unseen, a new place is assumed to have been discovered. Similarly, previously visited places are recognized based on the similarity of the WiFi fingerprints. This approach is widely used in the literature. We adopted the Tanimoto Coefficient [12] as the WiFi similarity function in this process, defined as:

$$\mathcal{S} = \begin{cases} \text{different (move)} & \text{, if } \frac{\vec{f_i}\cdot\vec{f_j}}{\|\vec{f_i}\|^2+\|\vec{f_j}\|^2-\vec{f_i}\cdot\vec{f_j}} \leq \varphi \\ \text{same (stationary)} & \text{, else} \end{cases}$$

where $\vec{f}$ is the WiFi vector containing the MAC address and signal strength, φ is the similarity threshold, and the output is a similarity estimated between 0.0 to 1.0. We set the φ to 0.7, as suggested by [13, 7].

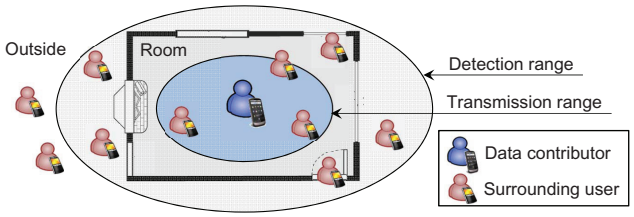

Figure 4. Conceptual view of co-location detection of surrounding users. The method uses the signal strength of packets within transmission range to choose mobile devices in the room.

Co-location Detection. The mobility of surrounding users comprises the place and the transition between visited places. We should infer the mobility information from passively received WiFi packet data, since the surrounding users do not actively sense their own locations.

We denote the phone owner, who senses WiFi packets by WiFi monitor mode, as *data contributor*. The data contributor senses a set of WiFi packets $P_t = \{p_1, p_2, \cdots, p_n\}$ at certain time t, where the packet p_i contains the MAC address of source and destination, channel (from 1 to 13), packet type (e.g., beacon request, beacon response, etc.), and signal strength. Among the scanned MAC addresses, we filter out one of the static WiFi APs by using a stored WiFi fingerprints. Then, we consider that each MAC address indicates the mobile device of surrounding users. When the data contributor is staying at a certain place, the surrounding user with the scanned MAC address is likely to be staying at the same place. However, the simple detection derives the overestimation, since the WiFi monitor mode can capture the packets from smartphones located in certain places as well as outside of the room, as shown in Figure 4. Thus, the simple use of detected packets may overestimate the number of visitors in places. To choose the devices in places, we utilized the consecutive detection of packets and the signal strength, expressed as:

$$\mathcal{U} = \begin{cases} \text{stationary} & \text{, if } p_e^t - p_s^t > \alpha \text{ and } p^s > \theta \\ \text{move} & \text{, else} \end{cases}$$

where p^s is the signal strength of the packet, p_s^t and p_e^t are the timestamps of the first packet and the last packet respectively in consecutive packets from certain devices. In other words, we consider that the surrounding user is staying in the places only if (1) the multiple packets are observed for a certain amount of time (i.e., five minutes) and (2) the signal strength of the packet is stronger than a certain threshold θ. Here, the use of constant value θ may cause a bias of estimation due to the different size of the place. We employed the signal strength of packets from the devices within the transmission range to adaptively choose the threshold. The idea is that the devices within the transmission range would be staying in the same space with the data contributor, as shown in Figure 4. To detect devices within the transmission range, we temporarily activate the soft AP mode with the service set identifier of the popular network operator. The surrounding phones then try to connect into the contributor's phone,

Figure 5. User interface of smartphone application. The application displays a user's mobility on (a) the Google map and (b) the list. (c) The user annotates the place name along with the list of places in Facebook.

since the commercial smartphones automatically connect to the previously connected APs [20]. The scheme then uses the distribution of the signal strength among the devices successfully connected to the contributor's smartphone in order to decide θ.

IMPLEMENTATION

Our system consists of a smartphone application and server infrastructure. For primary smartphones, we implemented the application on the Android SDK 4.0 running on commercial smartphones equipped with GSM/CDMA, WiFi, and GPS. The application indicates a user's visited places on the map, and the list interfaces to confirm/modify the mobility information, as shown in Figure 5. It also allows users to annotate a logical place as a particular place name in social network. For deployed smartphones, we modified the WiFi firmware and driver to enable the WiFi monitor mode in Nexus One. We implemented drivers as kernel module and integrated it into the Linux kernel 2.6.35. The application runs on the Android SDK 2.3 to periodically activate the WiFi monitor mode. For server infrastructure, we used the Windows Azure Cloud Service with 8 cores and 10TB of storage. The collected data were uploaded to the server with the access key.

STUDY FINDINGS AND IMPLICATIONS

We present key results and implications from our study that touch on the three major themes: (1) richer contexts in urban mobility estimation, (2) coverage and scalability, and (3) threats to privacy. We omitted the analysis of WiFi fingerprint-based place learning, since we adopted the proposed method in the literature [13, 7]. In brief, the adopted fingerprint technique correctly recognized 91% of the place visits.

Data Collection

We collected data traces from 25 university students over a seven-week period in Seoul, Korea. Among the 50 applicants, we gave priority to participants who took a class in the same building. Thus, they encountered within WiFi range at least once per week. Participants installed our smartphone application on their primary phones. The participants labeled the names of visited places to determine the ground truth regarding mobility. The data traces include 1,122 places with 8,434 stays and 131,351 unique devices with approximately

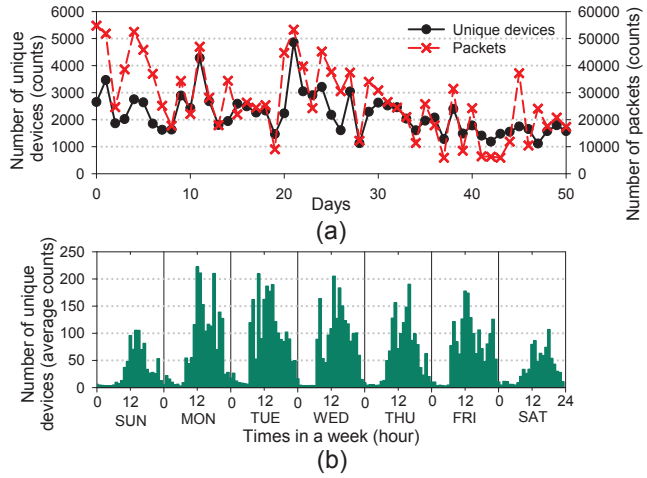

Figure 6. Diurnal patterns of detected devices according to (a) days and (b) hour-of-week.

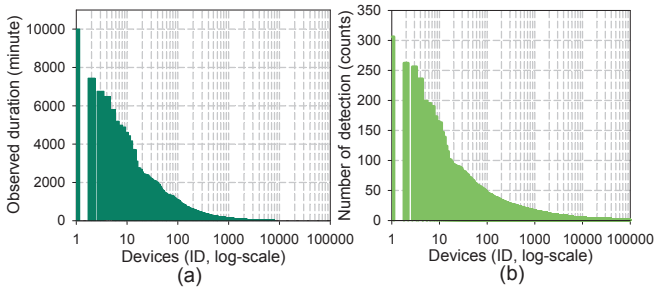

Figure 7. The distribution of detected devices (a) with observed duration and (b) with number of detections. The pattern follows a long-tail distribution.

four million WiFi packets. On average, a participant collected traces for 53 days and detected 5,254 unique devices.

Contexts in the Ambient WiFi Packets

We analyze the detection patterns of WiFi packets and investigated the diverse contexts in sensed ambient packets. Figure 6(a) shows the daily detection patterns of devices and packets. The participants detected, on average, about 2,200 devices with 27,000 packets per day. Figure 6(b) illustrates the diurnal patterns according to hour-of-week. The peak values appeared at lunchtime on weekdays and decreased slightly during the weekend. The users observed more devices on Friday nights than other days. This tendency directly reflected the human life pattern: people tend to encounter many people during workdays and meal times.

We investigate how much time each device was observed by our participants. The detected devices with observed duration follow the long-tail distribution, as shown in Figure 7(a). The top 20% of devices show 87% of observed durations, and each device in long-tail (i.e., 80% of devices) is detected for less than five minutes. Actually, among 131,351 devices in collected data, 48% of devices (i.e., 62,448) were detected by only one packet. Figure 7(b) shows how many devices were repeatedly discovered. We consider the consecutive packets (e.g., 20 packets for five minutes) as one-time detection. Approximately 36,000 devices were observed more than one time, and the top 20% of devices contained 53% of detections. The result indicates that the participants frequently observed

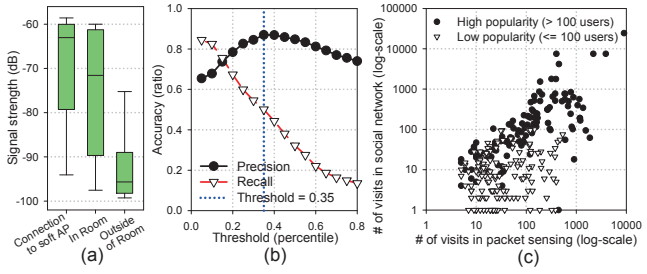

Figure 8. (a) Distribution of signal strength in the packets; (b) Parameter optimization for co-location detection. The use of signal strength is feasible for co-location detection; (c) Correlation between the number of visits counted by packet sensing and social network. The packet sensing generates similar results with social network in popular places and higher number of visits in unpopular places.

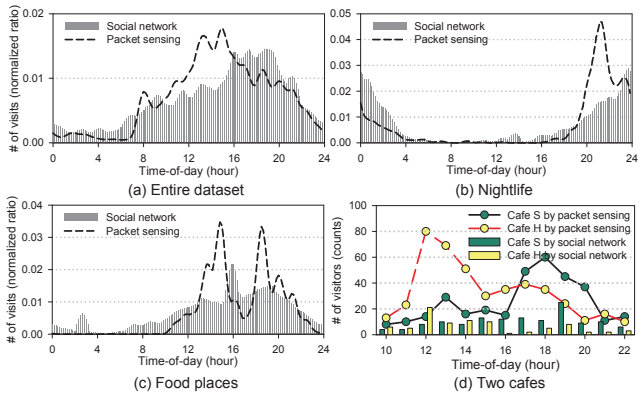

Figure 9. Number of visits counted by packet sensing and social network in (a) entire dataset, (b) nightlife places, (c) food places, and (d) unpopular two cafes. Our methods follows the trend of social network, yet provides detailed reports in low popularity places.

a certain number of people during their daily routines and encountered a large number of people within a short time.

We examine the performance of co-location detection and then validate the results by comparing with social network check-in data. To determine the threshold for co-location detection, we measured the signal strength from 10 smartphones located in the inside/outside of three places (i.e., office, classroom, and a reading room). Figure 8(a) shows the distribution of the signal strength received at the center of each place. Intuitively, the packets from the devices inside the room had stronger signals than the packets from the devices outside the room. We observed the strongest signals from the devices connected to the soft AP mode, since the connection occurrs with the device closely located to the phone with activating soft AP mode. The results indicate that the use of the signal strength is feasible to adaptively choose the smartphone located inside the places.

Figure 8(b) shows the performance of co-location detection according to the threshold. We found that the signal of specific percentile in the distribution is reasonably effective rather than the constant value since the signal strength would be different according to the density of people or the size of room. We can derive the most accurate precision (0.87) if we set a threshold of 0.35 percentile in the distribution. Although this threshold is effective in lowering false positives, it also causes many true positives to be ignored (i.e., low recall). However, we set a threshold as 0.35 percentile since the false-positive is a more critical error than false-negative in our scheme. The scheme generated a reasonably accurate detection (the precision is 0.76 and the recall is 0.39) from the 8,434 stays in real traces of 25 participants. The low accuracy is generated due to short-time visits or large open spaces, including cafeterias, lobbies, and theaters. The result indicates that our scheme is feasible to estimate co-locations of surrounding users. The problem of low recall could be solved by using the large volume of data collected by crowdsensing.

We validate the estimation results in comparison with social network check-ins data. Figure 8(c) shows the number of visits in places, counted by packet sensing and social network over four months. The proposed system derives similar results with social network in high popularity places (i.e., the number of visitors in social network is more than 100), but our scheme overestimates the number of visits in low popularity places. The Pearson correlation coefficient between the results from packet sensing and social network is 0.78 in popular places, but it is only 0.29 in unpopular places. The reason of low correlation in unpopular places is that social network users rarely performed check-ins at these places due to low social motivations (e.g., sharing experiences, receiving rewards), but our method passively estimated the surrounding users without user intervention. The result indicates that the proposed system enhances the contexts of places since the method provides unbiased reports about the number of visitors in places while the social network may miss such information in unpopular places.

The correlation analysis with temporal domain also indicates the advantage of WiFi packet sensing. Figure 9(a-c) shows the number of visits estimated by packet sensing and social network according to time-of-day. The Pearson correlation coefficient between two dataset is 0.75, and the correlation decreased to 0.43 ± 0.22 when the value is calculated according to place categories. The low value is caused by high peaks at certain time as shown in Figure 9(b, c), which were detected when our participants visited exceptionally crowded places. This tendency could be used to detect abnormality of mobility. In Figure 9(d), we present the number of visitors in two cafes which are not popular in social network (i.e., the number of visitors is less than 100). Our method provides detailed reports about the number of visitors: one cafe is crowded at lunchtime while the other is crowded at dinnertime. However, the social network provides relatively constant check-ins because of low popularity. The results indicate that our method follows the trend of social network data, but the scheme could further estimate the rich contexts such as the abnormality of mobility or the unbiased number of visits in low popularity places.

We now visualize the additional gain of sensing ambient packets for estimating the dynamics of urban mobility. Figure 10 shows the *food* places and the *nightlife* places visited by our participants in one week. The red circle is a place and the size of the circle shows the number of visitors in each place. The difference between Figure 10(a) and Figure 10(b,

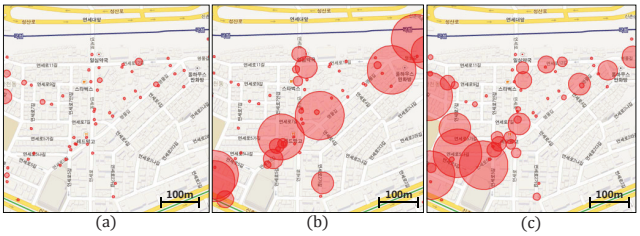

Figure 10. Distribution of visited food places and nightlife places in 600 m x 550 m region. The circles are places and the size of the circle indicates the number of visitors; (a) without packet sensing during the entire day, (b) 11 am to 2 pm with packet sensing, and (c) 5 pm to 11 pm with packet sensing. Our methods estimated rich dynamics of urban mobility according to time-of-day.

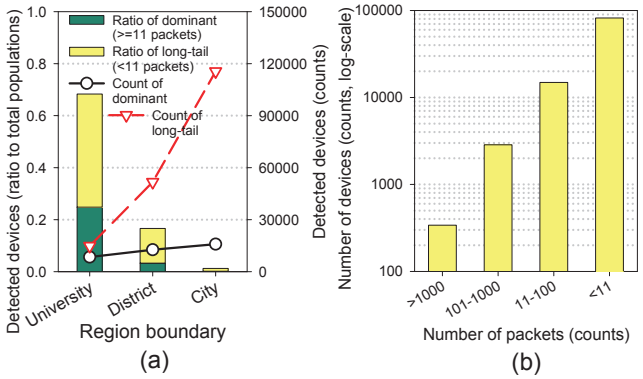

Figure 11. (a) Coverage of population and (b) devices according to number of detected packets. The system estimated incredibly high population coverage despite of a small number of users.

c) clearly present the additional gain of packet sensing. Without the packet sensing (in Figure 10(a)), the size of the circles is almost the same (i.e., a few visitors). However, the packet sensing clearly presents the diverse populations according to lunch and nighttime, as shown in Figure 10(b, c). The pubs are mostly located on the left side of the regions, and the results reflect this characteristic. In other words, the packet sensing enables the understanding of city dynamics by using a small number of users. In the active tracking of mobility (i.e., phone owners track only their own mobility) [13, 5, 27, 8], this characteristic is hard to estimate and would be biased depending on the number of participants. However, our method derives an impartial report about visited places, since the surrounding users are passively estimated.

Consequently, the sensing of ambient WiFi packets is feasible for estimating surrounding users and co-locations. The common use of WiFi enables our method to collect richer contexts and unbiased reports about urban mobility by using only a small number of participants.

Coverage and Scalability

In this section, we investigated the following coverage-related issues: (1) How many users are covered by 25 participants? (2) How many participants will be needed to scale up the spatial coverage? (3) How extensively are devices/places covered in the temporal aspect?

To estimate the population coverage, we categorized the locations into University, district, and city. The participants

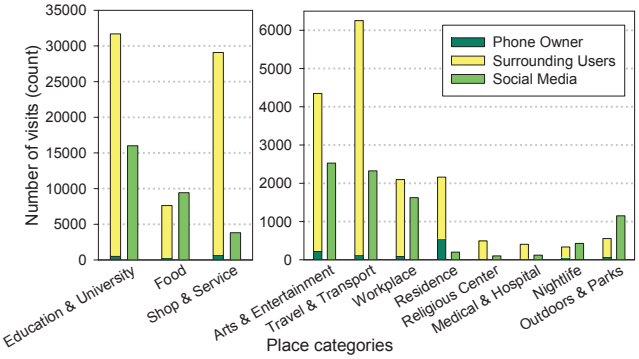

Figure 12. Visit counts of places in crowdsensing and social media according to categories. Our method estimated 2.1 times more visits than that of social network.

detected a total of 131,351 devices during the study, and approximately 18,000 devices (i.e., 18%) were detected by more than 10 packets, as shown in Figure 11(b). We defined these 18% of devices as the dominant devices and the rest as long-tail devices. The number of all detected devices is 1.8% of the entire population in the city; it is 0.2% if we consider only the dominant devices, as shown in Figure 11(a). However, a higher level of coverage is found when the boundaries of the regions are taken into consideration. The system covers 24.8% of the university's population and 3.3% of the district's population when considering the dominant devices. The result indicates that the packet sensing can derive incredibly high population coverage in terms of a small number of users.

We investigate how this coverage is distributed in the places people visit. Figure 12 compares the place visits made in our dataset to social network check-ins made over four months. Obviously, the phone owner made far fewer place visits than SNS users due to a much lower number of users. However, with the number of surrounding users, our participants (25 users) estimated an appreciably higher number of visits (i.e., 2.1 times more visits) than that of social network (31,000 users). Specifically, they collected 2.9 times more place visits than social network in several place categories where students are likely to visit (i.e., education, shopping, entertainment, transport). The reason is that check-ins in social network are manually performed, but our scheme automatically senses the existence of surrounding users. This shows that sensing WiFi packets can effectively collect populations in certain places.

We then investigate scaling up the spatial coverage of the sensing WiFi packets. Figure 13 shows the estimation model of the number of users and population coverage using the regression method. The generated model follows the power distribution with a 0.98 R^2 value and a 0.1% standard error. The model reveals that 200 participants may cover 70% and 13% of the population at the University and the district, respectively. Considering that 200 users are only 0.5% of the entire population of the University, this coverage is incredibly high. To predict the coverage in the entire city, we applied the estimation model to the 25 districts that make up Seoul. In this scenario, the model predicts that 5,200 users would be necessary to cover 13% of the population of the city. This is only 0.05% of the entire population. One key limitation of this

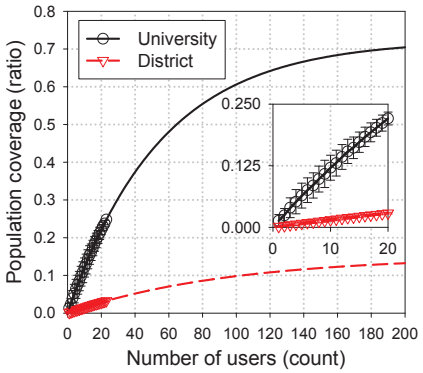

Figure 13. Estimation model about the number of users with the population coverage. The model reveals that 200 participants may cover 70% and 13% of the population at the University and at the district, respectively.

result is that we only used 25 participants; thus, an experiment with a large number of users is required to practically examine the model.

We examine the temporal coverage of packet sensing in terms of the inter-contact time and the contact duration. We defined the inter-contact time as the time elapsed between two detections of a given device; the contact duration is measured by the duration of the continuous detections of given devices. Figure 14(a) exhibits the empirical distribution of the inter-contact time according to the number of detected packets. For example, a 10-hour inter-contact time indicates that our participants observed a single person's smartphone 10 hours later from the previous observation. On average, the inter-contact time of dominant devices (devices with more than 10 packets) is approximately 82 hours, which is more than three days. This inter-contact time greatly decreases when we consider the most dominant devices (devices with more than 1,000 packets): the participants detected these 340 devices with less than 15.8 hours interval in 80% of cases. We expect these devices are carried by the acquaintances (e.g., family, friends, or co-workers) of our participants. In addition, the inter-contact time is approximately divided into two regimes around 24 hours. This 24-hour inter-contact time indicates that a user repeatedly observed the devices with a longer-than-a-day interval, which is understandable given the natural pattern of human life. In terms of contact duration, 90% of the detections lasted less than 30 minutes, and the duration of the most dominant devices slightly increased (i.e., 48 minutes). This low temporal coverage is understandable, given the small number of participants, but highlights the difficulty in providing continuous mobility information about surrounding users.

Finally, we explore temporal place coverage. Here, we considered only 260 popular places where more than 100 SNS users performed check-ins. We defined the temporal coverage as the coverage from 9 a.m. to 11 p.m., since temporal coverage at nighttime is meaningless in public places. We discretized the time into hourly intervals, and one slot was covered if at least one participant visited the place. For example, the 10 a.m. time slot at place A is covered if the user visited place A between 10 a.m. and 11 a.m. Thus, 100% temporal

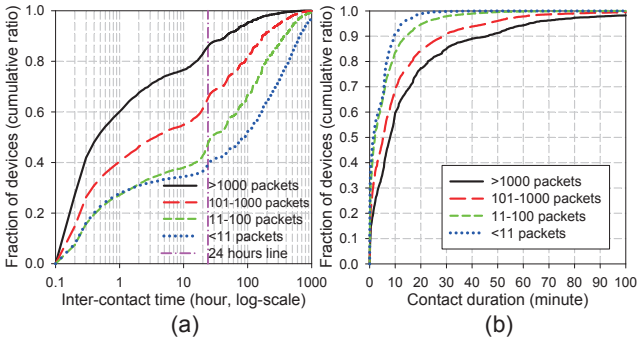

Figure 14. (a) Inter-contact time and (b) contact duration according to the number of detected packets.

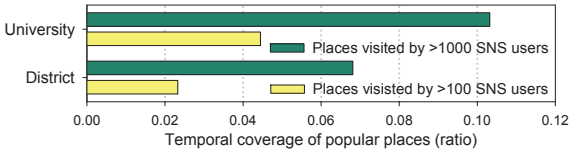

Figure 15. Temporal coverage of popular places located in the University and the district. The popular places are temporally covered only 6.3% on average.

coverage means that the participants visited all popular places every hour. Figure 15 shows the temporal coverage obtained during the study. The popular places are temporally covered 6.3% on average. The value increased to 10.8%, although we considered only 17 places within the university that more than 1,000 SNS users visited. The result reveals that obtaining high temporal place coverage is more challenging than the spatial coverage.

Individual Privacy Threats

The MAC address of mobile devices can be used as the identity of smartphone users. In the following set of analyses, we investigated the possibility of privacy threats derived from the sensing ambient packets by crowdsensing.

In this study, our participants observed approximately 130,000 unique devices in the city. Here, the identification of these 130,000 users is not feasible. For example, although we know the specific MAC address (e.g., 00:11:22:33:44:55), it is almost impossible to know the owner of this MAC address. However, we can easily know the MAC address of our acquaintance by using WiFi monitor mode (e.g., we can find the MAC address of a friend's smartphone). This is simply solved by physically encountering the specific users within the WiFi communication ranges for a few hours. The one-time detection is sufficient, since the MAC address is a mobile device's fingerprint (i.e., it is not changed until the user changes his/her phone).

We investigate how much time the surrounding users were observed by our participants. The long duration indicates that we can easily trace a certain user without their intention. Figure 16(a) shows the distribution of the longest observed duration of devices over the span of a day by each participant, ordered by duration time. For example, a two-hour long duration means that a participant observed the specific device for two hours at the most in a day. On average, the participant observed the certain surrounding user for three hours in a day,

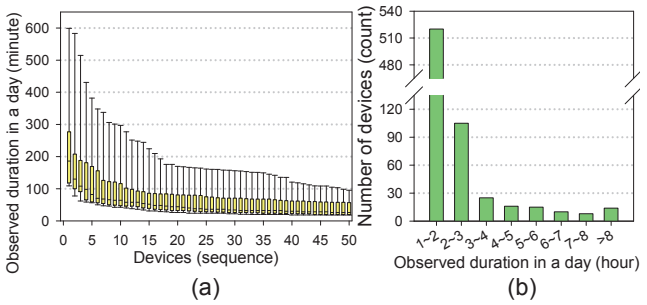

Figure 16. The observation time of specific users in a day collected by (a) one user and (b) all participants. Each participant detected more than 20 devices for longer than one hour in a day.

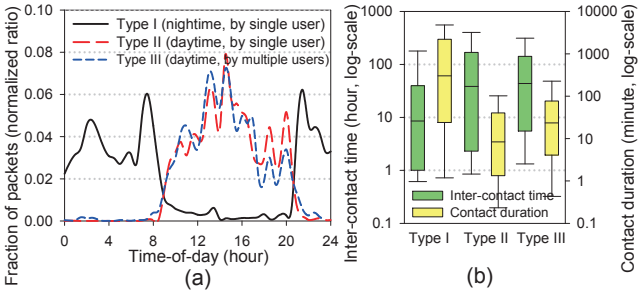

Figure 17. (a) The detection time, (b) inter-contact time, and contact duration of detected devices according to three types. On average, one user detected 2.4 type I devices, 3.9 type II devices, and 35.7 type III devices.

and at most for 10 hours. Each participant detected more than 20 devices for longer than one hour in a day, and 25% of participants observed more than 10 devices for longer than two hours. We further investigated the duration of observation in crowdsensing. Among the approximately 130,000 unique devices, 713 devices were observed for more than one hour in a day, as shown in Figure 16(b). Fourteen devices were observed for more than 8 hours in a day (i.e., almost the entire working day). The results indicate that *the crowd* can trace a certain user's everyday locations by opportunistically sensing WiFi packets in the air.

Our final concern was to investigate what type of devices were repeatedly detected. We found that the devices with more than 100 packets can be categorized into three types according to the detection time and the number of observers, as shown in Figure 17(a). The devices in type I are mostly observed at nighttime by a single participant; type II and III devices are mostly observed at day time, but type II devices are detected by single user while devices in type III are detected by multiple users. The inter-contact time of type I devices is significantly less than that of type II and III devices, while participants encountered the type I devices longer than type II and III, as shown in Figure 17(b). On average, one user detected 2.4 type I devices, 3.9 type II devices, and 35.7 type III devices. The device of type I was discovered more often (average 17.4 days) than type II (6.0 days) and III devices (9.9 days). We expect that the device owners of type I are family or friends who share residence with the participants, and people in types II and III are friends, co-workers, or familiar strangers of the participants. The reason for the high number of type III devices is that our participants routinely visited the campus where they encountered many students. The result indicates that the packet sensing from one user may expose the privacy of about of 40 his/her acquaintances. This tendency has pros and cons: the participation of one user greatly increases the population coverage, but he/she may provide the mobility information of acquaintances without their permission. The in-depth analysis on this issue requires large-scale experiments with being able to identify the ground truth of human relationships among the device owners.

DISCUSSION

In what follows, we describe the limitations of the present work along with future research directions. In this study, we covered a wide range of critical issues in sensing WiFi packets in the air. We believe that many of our findings highlight new directions of study and provide a starting point for designing subsequent studies.

Diverse Populations. Despite the scale of our study, the participant population is not particularly diverse. Our data was collected in a single city in Korea using university students. Moreover, smartphone users in South Korea are the most active phone users in the world [21]. Thus, the numbers in our results (e.g., WiFi usage time, coverage, and scalability) may be overestimated, since students are very likely to have more regular patterns than other people and to be co-located with other active participants in the study. As a result, some of our findings still remain to be verified by performing similar experiments elsewhere with more diverse populations.

Scheduling of WiFi Monitor Mode. We demonstrated that the energy consumption of the WiFi monitor mode is similar to GPS sensing. Similar to the diverse scheduling policies of GPS [13, 5, 8, 22], the adaptive sensing scheduling of the WiFi monitor mode is necessary. Considering that the WiFi monitor mode disables WiFi communication for Internet use, the policy is not trivial since it should consider not only the owner's smartphone usage, but also the surrounding environments to maximize the number of detected devices within minimum energy consumption. We expect that our work may induce extensive work in scheduling the WiFi monitor mode.

Estimation of Transition between Places. In this study, we validated that sensing ambient packets by smartphones is effective for collecting mobility information in places. We do not touch the estimation of transition between places, which requires the continuous observation of specific users. This is more challenging than the estimation of place visits because a specific user would not regularly come into the WiFi communication range of data contributors. Actually, among the 130,000 observed devices in our dataset, approximately 62,000 devices were detected only once during the entire period of study. This sparsity clearly shows the challenge of transition estimation. Transition information, along with place information, is important for fully understanding urban mobility. Thus, the intelligent model should be designed to interpolate the transition information from the sparse observation of packet sensing.

Privacy Concern. This study is directly related to the controversial question of privacy concerns in mobile sensing: If

smartphones sense the data about surrounding environments, including other people, then *who has the right of this data, the smartphone owner or the person being monitored?* We found that the continuous detection of a specific MAC address can be used for monitoring a user's location without his or her agreement. Similar to anonymizing an image or video data (i.e., blurring the face), anonymization of a MAC address is necessary in practice. Previous work [23] showed that the simple hashing of a MAC address is very easy to de-anonymize. Thus, a strict privacy-preserving scheme for the MAC addresses should be designed. For example, smartphones could periodically change the virtual MAC address or stop the WiFi probes if they knew they were being tracked. The sniffing mode is illegal in many European countries and possibly in the United States (depending on who is doing the sniffing). However, the contents in the packet are not necessary for sensing urban mobility; we only used the MAC address and the signal strength of the packets. In addition, a MAC address can be transformed into the form of another identifier to prevent the link between the MAC address and the user's identity. This implies that we may enable packet sensing in practice by using the *semi-WiFi monitor mode* (i.e., filtering the packet contents and recording only headers) and the strict anonymization of a MAC address.

RELATED WORK

In this section, we describe prior work while highlighting the novel contributions of our work.

Mobile Crowdsensing. Crowdsensing is an active area of interest and has been applied to a variety of different applications such as image searching [28] and ambience fingerprinting [2]. Recently, a number of projects have examined large-scale crowdsensing [6, 4]. These works [2, 6, 4] utilize large-scale smartphone-collected sensor data from users, including accelerometer, GPS, WiFi, camera, and microphone. Unlike prior research, our study performs a different type of crowdsensing and makes completely different contributions. The WiFi sensor is commonly used for detecting surrounding WiFi APs, but we utilize it to sense ambient WiFi packets, although the packets are not associated with a user's phone. We investigated the cost, privacy threats, and additional gain of packet-sensing for urban mobility monitoring.

Mobility Monitoring. We claimed that WiFi packet-sensing is effective for urban mobility monitoring in a passive manner. Passive mobility monitoring employs a specialized infrastructure to trace users in the area of interest. Previous work deployed special-purpose sensors, such as magnetic loops [17] or cameras [1], into streets or vehicles. These methods require costly infrastructures and specifically focus on traffic measurements. Several studies used configured WiFi APs [20] and/or Bluetooth devices [3, 15] to estimate the vehicle trajectory or populations in certain areas. In contrast to prior work, we used off-the-shelf smartphones without configured infrastructures. The coverage area of our system is therefore wider because our participants are mobile users, not static infrastructures. Several works [11, 27] used the call detail records from a cellular network to estimate urban mobility. These systems took an aggregated approach on the operator side with coarse-granularity (i.e., the range of cell towers), whereas we aimed to estimate a finer granularity (i.e., indoor places) by device-side crowdsensing. To the best of our knowledge, this WiFi packet-sensing by smartphone-based crowdsensing is the first attempt in the literature.

Proximity Sensing. Various techniques have emerged to measure the proximity of mobile users based on available technologies in smartphones (i.e., GPS, cell, WiFi, and Bluetooth). Previous studies [16, 18] have used WiFi fingerprints or GSM readings to explore the proximity of mobile users. Eagle et al. [9] employed Bluetooth to infer the social interactions among users, and Tang et al. [25] used WiFi in laptops to detect surrounding computers in order to infer space availability in coffee shops. Recently, by using Bluetooth and WiFi fingerprints, Liu and Striegel [19] explored the proximity between 200 students for two years. Compared to these works, we used the ambient WiFi packets that were opportunistically transmitted from unmodified smartphones. Our method does not require either pairing or installation of additional applications to surrounding smartphones. Thus, our method is practical and the coverage is much higher than that of previous work. We expect that the WiFi monitor mode could be widely used for proximity sensing.

CONCLUSION

In this paper, we presented a detailed study about smartphone-based crowdsensing that senses WiFi packets in the air. We used the WiFi monitor mode that can be implemented into commercial smartphones in use today. We collected 1,113 user-days data from 25 study subjects who detected a total of approximately four million WiFi packets in Seoul, Korea. Through our analysis from the experimental study, we validated that the use of sensing WiFi packets enables the collection of richer contexts about urban mobility using only a small number of participants (25 users). The collected dataset captures a surprisingly high number of mobile users in urban areas (i.e., about 130,000 users) and richer information about the place visits than that of social network (i.e., 2.1 times larger visit counts). This advantage may lead the new era of mobile sensing, since it effectively collects impartial data about mobility and significantly increases the coverage of mobility monitoring in a city. However, the low temporal coverage highlights the difficulty in providing continuous mobility information about the surrounding users. In addition, the packet sensing can invade users' privacy. The dataset traces the location of 713 devices for more than one hour in a day, and 14 devices were observed for more than 8 hours in a day. The participants repeatedly observed about 40 devices of their acquaintances. In other words, the locations of surrounding users would be exposed without their permission. This issue should be resolved in practice by anonymizing MAC addresses and using semi-WiFi monitor mode. We believe our analysis and findings will provide valuable insights not only for builders of crowdsensing applications, but for other closely-related sensing systems that rely on a close engagement with urban life. We plan to integrate the static APs at a few popular places or public buses into the system in order to increase the spatio-temporal coverage of urban mobility.

ACKNOWLEDGMENTS

This work was supported by the Microsoft Research and the National Research Foundation of Korea grant funded by the Korean government, Ministry of Education, Science and Technology (No.2011-0006464, No.2013-027363).

REFERENCES

1. Anagnostopoulos, C.-N., Anagnostopoulos, I., Psoroulas, I., Loumos, V., and Kayafas, E. License plate recognition from still images and video sequences: A survey. *Intelligent Transportation Systems, IEEE Transactions on 9*, 3 (2008), 377–391.
2. Azizyan, M., Constandache, I., and Roy Choudhury, R. Surroundsense: mobile phone localization via ambience fingerprinting. In *Proceedings of 15th Annual International Conference Mobile Computing and Networking*, MobiCom'09, ACM (2009), 261–272.
3. BLIP System. BLIP system. **http://bliptrack.com**.
4. Chon, Y., et al. Understanding the coverage and scalability of place-centric crowdsensing. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '13, ACM (2013), 3–12.
5. Chon, Y., Kim, Y., Shin, H., and Cha, H. Adaptive duty cycling for place-centric mobility monitoring using zero-cost information in smartphone. *Mobile Computing, IEEE Transactions on* (2013).
6. Chon, Y., Lane, N. D., Li, F., Cha, H., and Zhao, F. Automatically characterizing places with opportunistic crowdsensing using smartphones. In *Proceedings of the 14th International Conference on Ubiquitous Computing*, Ubicomp'12, ACM (2012), 206–212.
7. Chon, Y., Talipov, E., Shin, H., and Cha, H. Mobility prediction-based smartphone energy optimization for everyday location monitoring. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, SenSys'11, ACM (2011), 82–95.
8. Constandache, I., et al. Enloc: Energy-efficient localization for mobile phones. In *INFOCOM'09*, IEEE (2009), 2716–2720.
9. Eagle, N., Pentland, A. S., and Lazer, D. Inferring friendship network structure by using mobile phone data. *Proceedings of the National Academy of Sciences 106*, 36 (2009), 15274–15278.
10. Falaki, H., et al. Diversity in smartphone usage. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*, MobiSys'10, ACM (2010), 179–194.
11. Isaacman, S., et al. Human mobility modeling at metropolitan scales. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*, MobiSys '12, ACM (2012), 239–252.
12. Jaccard, P. The distribution of the flora in the alpine zone. *New Phytologist 11*, 2 (1912), 37–50.
13. Kim, D. H., et al. Sensloc: sensing everyday places and paths using less energy. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, ACM (2010), 43–56.
14. Klepeis, N., et al. The National Human Activity Pattern Survey (NHAPS): a resource for assessing exposure to environmental pollutants. *J. Expo. Anal. Environ. Epidemiol. 11*, 3 (May-Jun 2001), 231–252.
15. Kostakos, V., O'Neill, E., Penn, A., Roussos, G., and Papadongonas, D. Brief encounters: Sensing, modeling and visualizing urban mobility and copresence networks. *ACM Trans. Comput.-Hum. Interact. 17*, 1 (Apr. 2010), 2:1–2:38.
16. Krumm, J., and Hinckley, K. The nearme wireless proximity server. In *UbiComp 2004: Ubiquitous Computing*, N. Davies, E. Mynatt, and I. Siio, Eds., vol. 3205 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2004, 283–300.
17. Kwong, K., Kavaler, R., Rajagopal, R., and Varaiya, P. Arterial travel time estimation based on vehicle re-identification using wireless magnetic sensors. *Transportation Research Part C: Emerging Technologies 17*, 6 (2009), 586 – 606.
18. Li, K. A., et al. Peopletones: A system for the detection and notification of buddy proximity on mobile phones. In *Proceedings of the 6th International Conference on Mobile Systems, Applications, and Services*, MobiSys '08, ACM (2008), 160–173.
19. Liu, S., and Striegel, A. D. Exploring the potential in practice for opportunistic networks amongst smart mobile devices. In *Proceedings of the 19th Annual International Conference on Mobile Computing and Networking*, MobiCom '13, ACM (2013), 315–326.
20. Musa, A. B. M., and Eriksson, J. Tracking unmodified smartphones using wi-fi monitors. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, SenSys '12, ACM (2012), 281–294.
21. Nielsen Research. The mobile consumer report 2013. *Mobile World Congress* (2013).
22. Paek, J., Kim, J., and Govindan, R. Energy-efficient rate-adaptive gps-based positioning for smartphones. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*, MobiSys '10, ACM (2010), 299–314.
23. Pang, J., Greenstein, B., Gummadi, R., Seshan, S., and Wetherall, D. 802.11 user fingerprinting. In *Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking*, MobiCom '07, ACM (2007), 99–110.
24. Seoul Metropolitan Government. Statistics of seoul. **http://stat.seoul.go.kr**.

25. Tang, K. P., Keyani, P., Fogarty, J., and Hong, J. I. Putting people in their place: An anonymous and privacy-sensitive approach to collecting sensed data in location-based applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, ACM (2006), 93–102.

26. Wiesenthal, T., Leduc, G., Cazzola, P., Schade, W., and Köhler, J. Mapping innovation in the european transport sector. *An assessment of R&D efforts and priorities, institutional capacities, drivers and barriers to innovation. JRC Scientific and Technical Report.* (2011).

27. Xu, Q., Gerber, A., Mao, Z. M., and Pang, J. Acculoc: practical localization of performance measurements in 3g networks. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*, MobiSys '11, ACM (2011), 183–196.

28. Yan, T., Kumar, V., and Ganesan, D. Crowdsearch: Exploiting crowds for accurate real-time image search on mobile phones. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services*, MobiSys '10, ACM (2010), 77–90.