WiFi Networks in Metropolises: From Access Point and User Perspectives

Lei Zhang, Liting Zhao, Zhi Wang, and Jiangchuan Liu

The authors have closely collaborated with a leading network service provider to collect massive information about wireless APs and their users in four metropolises. In this article, based on the large-scale dataset, they attempt to answer the critical questions on how the APs are deployed and how they are utilized in urban areas.

Abstract

Although WiFi was initially designed for providing local access to the Internet, today's expansive and explosive deployment of WiFi networks has enabled nearly ubiquitous access for mobile users in many urban areas. However, these WiFi networks have been woefully undermeasured and underinstrumented during the wild expansion. We have closely collaborated with a leading network service provider to collect massive information about wireless APs and their users in four metropolises. In this article, based on the large-scale dataset, we attempt to answer the critical questions on how the APs are deployed and how they are utilized in urban areas. At the macro level, we depict the coverage of WiFi networks and the usage patterns, and by carefully classifying the APs, we unveil rich geographical features of today's urban WiFi networks. At the micro level, we identify the implicit social relationships among WiFi users, and uncover the underlying social communities that have great potential for network optimization.

Introduction

Modern cities are ever increasingly expanding in terms of geographical coverage, residential population, and social and economic functionalities. With such speed of urbanization, today's metropolises are facing significant challenges in informational, intelligent, and integrated management. To build appreciable and measurable *smart cities* with seamless interconnection and interoperability, ubiquitous Internet access has become one of the underlying fundamentals. While *wireless local area network* (WLAN) technology was initially designed to provide local access for a limited number of users, the expansive and explosive deployment of 802.11 WiFi networks and the universal availability of WiFi interfaces in state-of-the-art mobile terminals has led to great Internet coverage nowadays, especially in urban areas. Besides business-oriented restaurants and hotels, such public service sectors as universities and community centers have deployed WiFi networks on tremendous scales as well. Many Internet service providers, together with regional governments, have also initiated plans toward city-wide WiFi coverage. Urban users reportedly have registered WiFi accesses 70 percent of the time [1], and are often exposed to multiple *access points* (APs) in one location.

Given that nearly 80 percent of mobile data usage is more nomadic than highly mobile [2], the massive urban WiFi networks have become major carriers, offloading much of the traffic from the conventional wired networks and the more expensive cellular mobile networks [3]. Considerable research efforts have been made on leveraging WiFi networks to balance traffic load, improve energy efficiency, and assist content delivery for mobile users. However, during this era of wild expansion, it remains unclear how the WiFi APs are deployed in modern cities, not to mention how they are connected and utilized by the massive number of users, and their dynamics over time. Pioneering studies on the deployed WiFi networks and the corresponding user access patterns have quite limited network scale [4], geographical area (e.g., campus only) [5], target application [6], and hardware platform [7]. Even though the state-of-the-art WiFi measurements [7] attempt to provide relatively comprehensive understanding from different layers and aspects, the social implications of WiFi users are seldom addressed.

In this article, based on a large-scale dataset (about *8 million WiFi APs and 27 million connection records from 6.4 million active users per day*) collected during a one-month period (March 12 to April 12, 2015) in four metropolises, we attempt to answer the questions of how WiFi networks are deployed and how they are utilized in urban areas. At the macro level, we depict the coverage of WiFi networks and the usage patterns, and by carefully classifying the wireless APs, we unveil rich geographical features of today's urban WiFi networks. At the micro level, we identify the implicit social relationships among WiFi users, revealing the underlying social communities that have great potential for network optimization. In particular, we find that:

- Today's WiFi networks are densely deployed in urban areas, which are driven by and naturally reflect the intensity of human social and economic activities.
- Business and public WiFi networks' deployments are more concentrated, whereas residential WiFi networks are more evenly distributed in urban areas.
- A majority of WiFi users have regular access patterns during workdays, and most user-AP connections have durations that are either ultra-short (less than 5 minutes) or quite long (longer than 1 hour).

Digital Object Identifier:
10.1109/MCOM.2017.1600262

Lei Zhang and Jiangchuan Liu are with Simon Fraser University; Liting Zhao and Zhi Wang are with Tsinghua University.

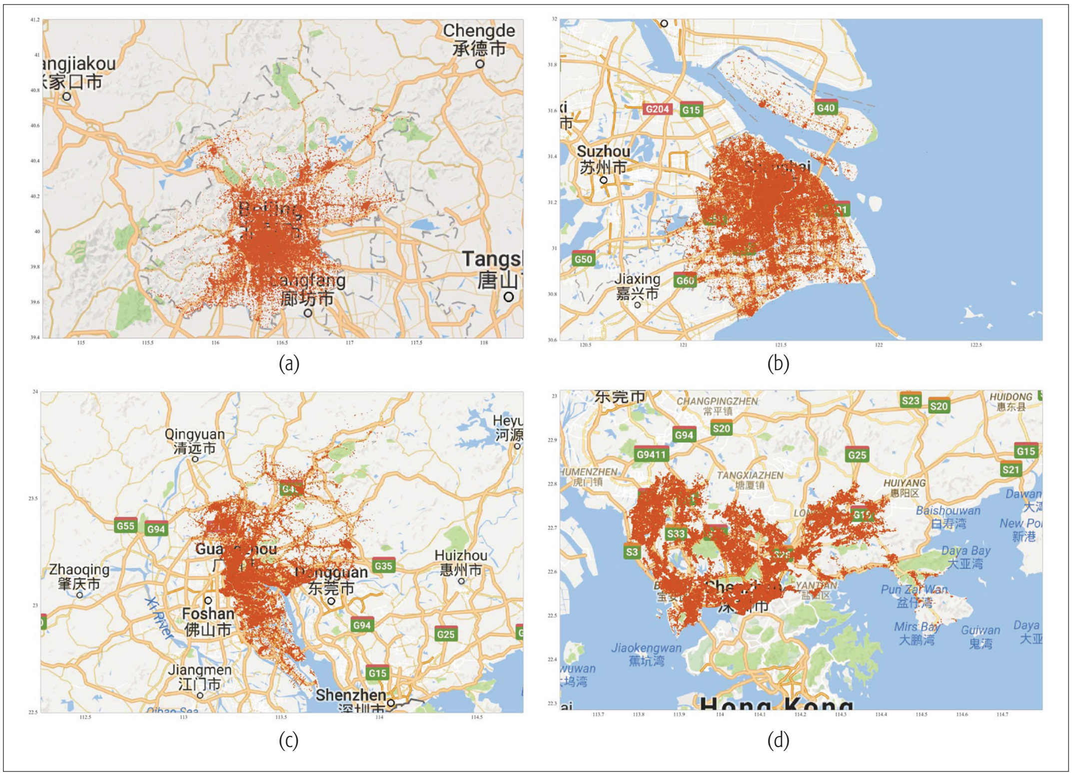

Figure 1. The deployments of wireless APs in four cities: a) Beijing; b) Shanghai; c) Guangzhou; d) Shenzhen.

To ensure secure access to the shared WiFi networks, the users are willing to report their connections to Tencent and have the traffic monitored. It is worth noting that only the related meta data is collected, so that none of the users' personal information was traced during the collection, nor were the conversation data intercepted or stored.

- The WiFi users geographically form a well structured small-world network, which has never been clearly identified by researchers and WiFi users before.

All these observations suggest that WLANs are shifting from a complementary technology for wide area access to a primary accessing technology for today's metropolises, although challenges remain to be addressed.

Data Collection and Description

We have closely collaborated with *Tencent*, one of the largest Internet companies in China, to conduct a nation-wide measurement. Tecent's major products, including QQ, Tencent Weibo, and WeChat, have over 1 billion active subscribers. In a one-month duration (from March 12 to Aril 12, 2015), we collected the location and ownership information of the reported wireless APs, as well as the connection records between WiFi users and those APs. The data were contributed by the users and business partners of *Tencent Mobile Manager* (http://m.qq.com/), a widely used utility software on mobile platforms with over 230 million downloads. A key functionality of the manager is to facilitate a user smartly selecting the best WiFi network among the currently available ones, and for the business partners to share networks with individual users. A considerable portion of the users also volunteer to share known WiFi accesses by providing the passwords to the manager and hence other users of the manager. To ensure secure access to the shared WiFi networks, the users are willing to report their connections to Tencent and have the traffic monitored. It is worth noting that only the related meta data is collected, so none of the users' personal information was traced during the collection, nor were the conversation data intercepted or stored.

Our collected dataset consists of two parts, the data of *8 million wireless APs* and the data of user-AP connection records from *6.4 million daily active users* (around *27 million WiFi accesses per day*). For each reported wireless AP, we have the information of the AP's geographic location (latitude and longitude), its ownership, the Internet service provider (ISP), and the medium access control (MAC) address. With the location information, we are able to analyze the deployments of wireless APs in four major cities in China. Moreover, the ownership information allows us to classify the APs into different categories (i.e., business, residential, and public). Each reported user-AP connection includes the information of the network's BSSID, the user IP (which was hashed so as not to reveal the real user IP), the server IP, the connect time, and the connect duration. By identifying and merging the records that belong to the same user, we are able to analyze the connection activities of individual users.

How Are the WiFi Access Points Deployed?

AP Distribution and Coverage

We first inspect the locations of deployed wireless APs in our dataset, and investigate the current deployments of WiFi networks. Figure 1 visualizes the deployments of wireless APs in the four major cities, Beijing, Shanghai, Guangzhou, and Shenzhen, which are representative of today's metropolises in China (with urban populations of 21, 24, 20, and 12 million, respectively). As shown in Fig. 1, well developed urban regions

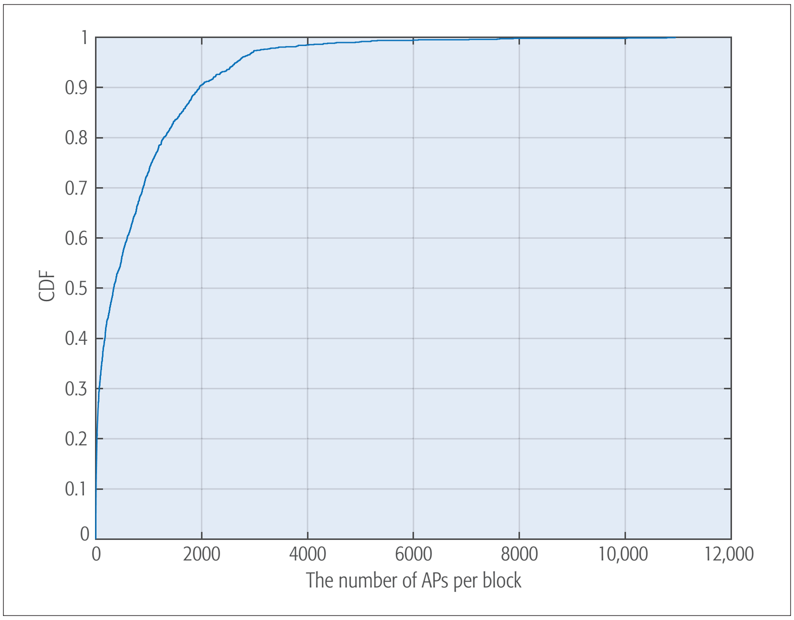

Figure 2. CDF of number of APs per block.

have much denser deployments of wireless APs than less developed suburban regions. It is not surprising that the wireless APs are deployed toward the economically well developed and highly populated regions. The deployment is also confined by natural topography: there is nearly no wireless AP deployed in water, mountains, and forests. In short, the deployments of WiFi networks in today's metropolises are driven by and naturally reflect the intensity of human social and economic activities.

To better understand the WiFi networks' distributions in different areas, we divide the latitude and longitude space into 0.0009° × 0.0009° latitude/longitude grids (approximately 100 m × 100 m). The reason that we adopt this setting is two-fold: first, dividing geographical area into blocks allows us to index locations and better study the relationship between the APs' deployments and locations; second, the APs that are located within this range are likely to have overlapping coverage, and thus the users in the same block probably have opportunities to connect to different WiFi networks, and even to establish device-to-device communications with other users. This will help in explaining and utilizing the potential social relationships between mobile users; for example, early studies discovered that people calling while connected to the same cell tower (colocation) are good proxies for face-to-face meetings [8] .

Applying the above settings, our analysis of the dataset shows that the distribution of the APs is highly skewed: a small number of blocks have deployed very large numbers of APs, while others have far smaller populations of APs. Taking Shenzhen as an example, Fig. 2 shows that about 30 percent of the blocks have over 1000 APs, and around 10 percent of the blocks have more than 2000 APs. Despite the fact that there may be high-rise buildings in developed urban areas, this result suggests that today's metropolises have good coverage of WiFi signals. In fact, many urban areas already have more than enough available APs, and further deployment should be strategically planned and optimized by examining the collective WiFi coverage and possibly removing redundant APs.

In contrast, the statistic of our dataset shows that only 3.2 percent of the blocks in Shenzhen (14.6 percent in Beijing, 9.8 percent in Shanghai, 9.5 percent in Guangzhou) have less than or equal to 1 wireless AP, which we consider insufficient WiFi coverage as few existing commercial APs are powerful enough to cover a 100 m × 100 m block area. After closely examining the blocks that have limited WiFi coverage on the map, we find that these blocks are often in regions that have limited population of residents and rare business activities (e.g., water, mountains, forests, and suburban areas). Our

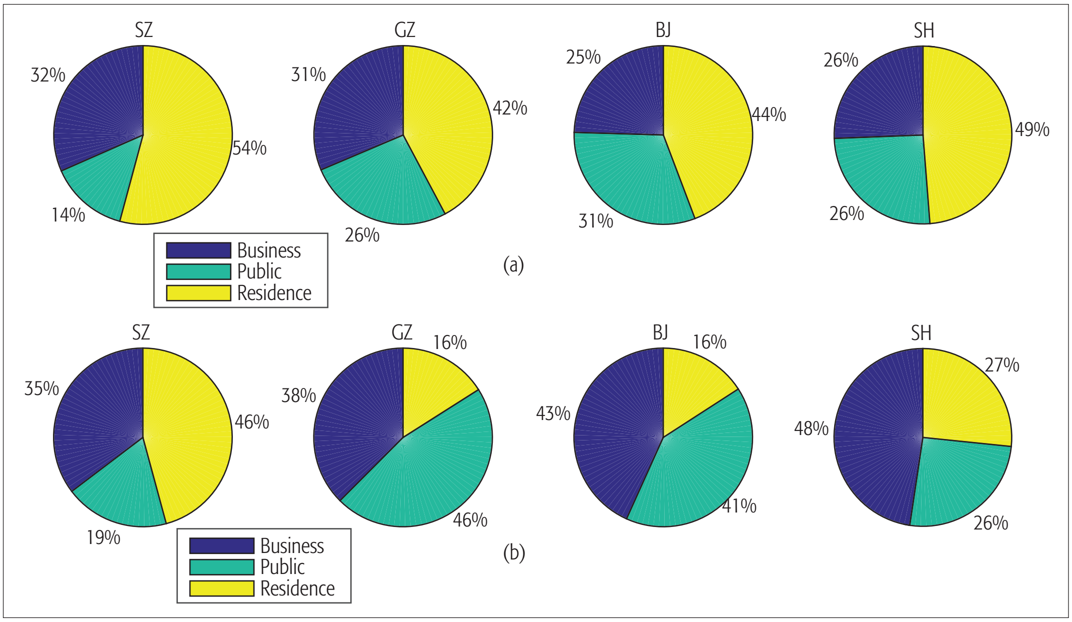

Figure 3. AP/block constitution in four metropolises (SZ stands for Shenzhen, GZ stands for Guangzhou, BJ stands for Beijing, SH stands for Shanghai): a) classification of APs in four metropolises; b) classification of blocks in four metropolises.

measurement shows that today's metropolises already have dense deployments of wireless APs, which can collectively provide near-ubiquitous Internet access. The deployments and applications of wireless mesh networks [9] have been extensively discussed in the literature, which, however, can hardly be realized on a large scale by a sole operator. As more and more wireless APs will be deployed on demand, we can foresee that in the near future wireless mesh networks may be constructed and maintained at the metropolis level with the help of mobile wireless APs that can be carried by vehicles or whose role can be fulfilled by increasingly capable personal mobile devices.

AP Categories

We next analyze the types of wireless APs. Specifically, we classify the APs into three categories according to the collected ownership information: business (e.g., shopping centers, restaurants, hotels, entertainment venues, domestic services, auto services, banks, and other companies), public (e.g., government offices, schools/universities, travel sites, public gyms, cultural venues, and other infrastructures), and residential. We further classify the blocks into these three categories based on the dominant type of wireless APs in this block. Figure 3 shows the constitutions of the APs and location blocks in the four metropolises. In all the cities, residential APs are the largest component, and public APs are the smallest one (except for Beijing, the capital of China). Existing studies on leveraging WiFi networks to offload traffic usually focus on public or business wireless APs. Our results suggest that there remain great opportunities in fully utilizing residential WiFi networks, which can provide substantial bandwidth resources especially during off-peak hours. However, private WiFi network owners may be reluctant to share their paid services from ISPs with strangers, and thus such incentives as monetary reward should be offered, or the social relationship among providers and users should be explored.

According to Fig. 3b, public blocks take up a large portion, especially in Beijing (41 percent) and Guangzhou (46 percent). It makes sense as today's metropolises are well developed in terms of the infrastructure construction. Note that universities play a huge role in modern cities' education systems, and their campus WiFi networks contribute significantly to the public WiFi network coverage. Moreover, we observe that business blocks also account for a considerable percentage in all four cities (35 percent in Shenzhen, 38 percent in Guangzhou, 43 percent in Beijing, 48 percent in Shanghai), implying that today's business operators have considered providing WiFi networks as one of their top priorities to serve customers. Another interesting observation is that residential APs take up a much larger share than residence blocks, while business APs have a much smaller share than business blocks. This can be explained by the fact that business APs usually concentrate on certain locations, and residence APs are more evenly distributed, which again indicates the great potential of exploring residential APs toward ubiquitous coverage.

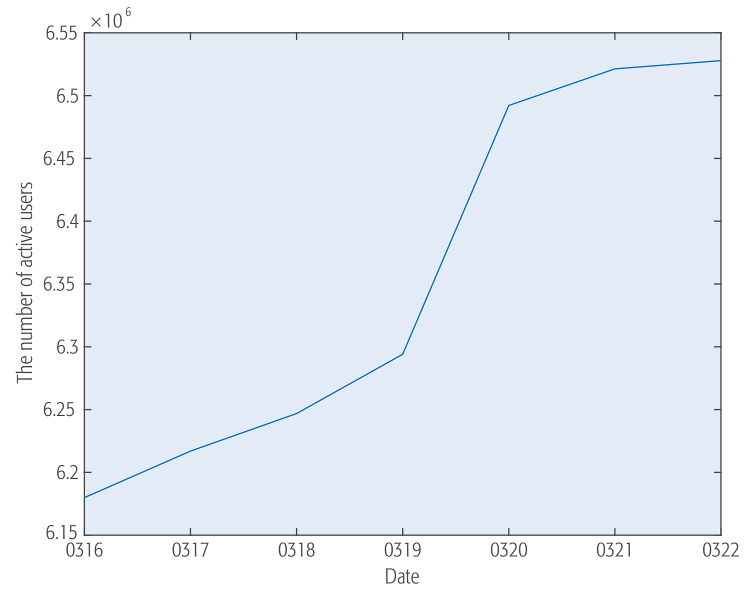

Figure 4. Number of users in each day.

How Does the Individual User Connect to WiFi?

So far we have focused on the information from the wireless APs' perspective. We now examine the data of user-AP connections to study how the users utilize the WiFi networks. Figure 4 shows the number of active WiFi users during the week from March 16 to March 22, 2015. As we can see, there are more users accessing WiFi on weekends (2015, 03, 21–2015, 03, 22) than on weekdays (2015, 03, 16–2015, 03, 20). Our later results also show that, for different scales of WiFi users, the access patterns are quite similar. We next check when each user connects to a WiFi network, and plot the number of user connections in the one-day time span in Fig. 5a. The results are based on the time when each user-AP connection is constructed. It is clear that there are three daily peaks: around 9 a.m., 1 p.m., and between 6 p.m. and 7 p.m. Intuitively, most users arrive at workplaces at 9 a.m., return from lunch breaks at 1 p.m., and get back home before or after dinner between 6 p.m. and 7 p.m. Such mobility patterns (e.g., the home-work/study-home daily routines) for the majority of WiFi users during the weekdays certainly help with developing smart city applications (e.g., smart office/home applications). Knowing the users' mobility routines (even partially) can also inspire the design of content delivery schemes using device-to-device communication with relays of wireless APs.

From the dataset, we also observe that more than 90 percent of users connect to no more than 2 wireless APs in a day. Together with the previous observation, this suggests that for most users, their WiFi accesses are not highly dynamic and may even be predictable. Revealing WiFi access patterns for mobile users can significantly enhance media content delivery. For example, with today's powerful smart routers, the providers can in-network cache the contents close to the APs from which the consumers would demand

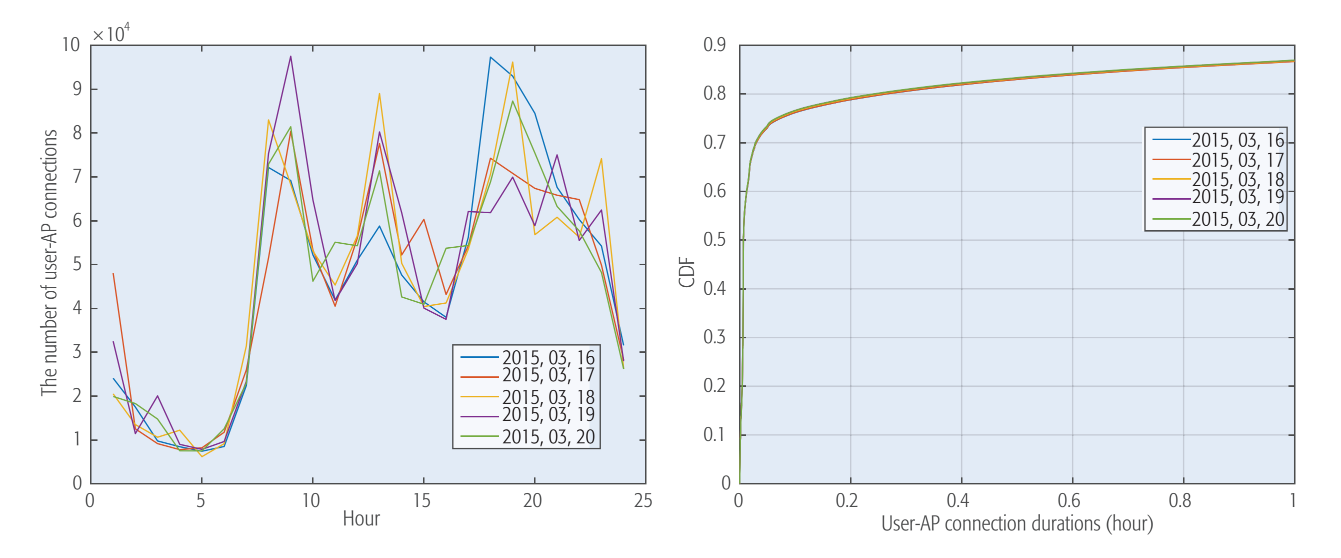

Figure 5. User-AP connection time and duration: a) number of user-AP connections in 24 hours; b) CDF of user-AP connection durations.

the media contents. On the other hand, our observation implies that the wide WiFi coverage in today's metropolises does not directly impact the users' WiFi access behaviors. If ubiquitous WiFi access becomes available in the foreseeable future, mobile users would access WiFi much more frequently and dynamically, which brings both challenges and opportunities in designing smart city applications.

We further investigate how long each time a user connects to a WiFi network during the collection period. Figure 5b shows that the distribution of the connection durations are highly consistent across weekdays. As can be seen, there are mainly two types of user-AP connections: *temporary connections* of short duration and *stable connections* of long duration (together accounting for over 82 percent of the connections), among which about 70 percent last for less than 5 minutes, while over 12 percent are longer than 1 hour. The high percentage of temporary connections implies that smart city applications are demanded to respond to user requests in a timely fashion. On the other hand, the existence of stable connections offer opportunities for long-term use of the corresponding user devices. In particular, once the user-AP connection is sustained for 10 minutes, it probably would last for much longer, in which case one can infer the user is static rather than nomadic, and turn the connected device into a crowd-sensing component in the Internet of Things to collectively contribute to complex tasks, for example, urban noise monitoring and structural health monitoring.

Revealing Implicit Social Relations

A social network is a social structure defined by actors, relationships, and other interactions between actors, which widely exists in both the human world (e.g., family ties between relatives) and the cyber world (e.g., friendships on Facebook). Upon identifying such social structures, social network analysis focuses on patterns of relationships between actors and examines the availability of resources and the exchange of resources between these actors [10]. The resources exchanged can be of many types, for example, information in a communication context. By studying the social network properties, we can understand what kind of information is exchanged, between whom, and to what extent. Today's dense WiFi network deployments in metropolises and the regular accesses, as we have seen, imply that there are certain underlying social structures among the massive number of WiFi users; if identified, they will certainly help us better utilize these WiFi networks. Unlike human or cyber social networks, mobile WiFi users do not have explicit meaningful relations. Our objective in this section is thus to investigate whether WiFi users have implicit social relations of which most of the users are not aware.

Constructing the Social Network

As mentioned, most WiFi users connect to no more than two wireless APs per day, which implies that, in practice, only a very small portion of users connect to the same wireless APs. Defining social relations between users based on sharing common WiFi connections to the same APs will result in a very sparse network that can hardly be interpretable or have practical meanings. Therefore, we assume that there is a social edge between two users if both of them access the WiFi networks in the same location block within a one-day time span. In other words, the social relations between WiFi users are defined on their geographical WiFi access patterns. Given the 100 m × 100 m block size, the WiFi users in the same block can have chances for direct communication or can easily be relayed through certain APs. Accordingly, we have analyzed the 10 million user-AP connection records in Shenzhen, and built the corresponding network of WiFi users with 67,264 nodes and 4,274,997 edges.

The constructed network consists of 245 connected components. We extract the subgraph of the largest connected component (referred to as the *WiFi user network* for the remainder of this article), which contains about 95 percent of the nodes and 99 percent of the edges in the complete graph, and present our analyses based on it hereafter. We summarize the basic statistics of the extracted subgraph in part 1 of Table 1. Note that we do observe some isolated user groups, each of which has a very limited population (from 2 to 100 users). Those isolated WiFi users may live or work in certain suburban areas, which are out of our focus here.

Small-World Network

The small-world network phenomenon is probably the most interesting characteristic of social networks. Watts and Strogatz adopted this concept to describe networks that are neither completely random nor completely regular, but possess characteristics of both [11]. They introduce a measure of one of these characteristics, the cliquishness of a typical neighborhood, as the clustering coefficient of the graph. They define a small-world graph as one in which the clustering coefficient is still large, as in regular graphs, but the measure of the average distance between nodes (the characteristic path length) is small, as in random graphs.

We compute the two small-world metrics for the WiFi user network. The result shows that the WiFi user network has definite small-world characteristics. As shown in Table 1, the average clustering coefficient is extremely high (0.91), which is very close to 1, the clustering coefficient of a regular graph. On the other hand, although the WiFi user network has a diameter of 22, the average shortest path length (the characteristic path length) is only 5.9024, and the 90th percentile diameter is 8, which nearly follows the famous *six degrees of separation* rule. The observation of the small-world network phenomenon confirms that the proposed WiFi user network is indeed a social network, which has never been identified before and is even beyond the awareness of its own members, the WiFi users.

Community Detection

We further apply the *k*-clique clustering algorithm to the WiFi user network. A *k*-clique community is the union of all cliques of size *k* that can be reached through adjacent (sharing *k* – 1 nodes) *k*-cliques. We vary the value of *k* from 10 to 150, and for each *k* present the number of clusters and the average cluster size (the number of nodes in the cluster) in part 2 of Table 1. It is clear that, whatever the value of *k* is, a considerable number of communities can be detected in the WiFi user network. In particular, for *k* = 10, 20, 50, 100, and 150, there are 107.2, 103.8, 92.1, 61.1, and 32.0 percent of the nodes clustered into different communities, respectively. It should be noted that there are cases where over 100 percent of nodes are clustered. This implies that overlapping communities are detected when *k* is small.

The result of *k*-clique clustering indicates that the WiFi user network is well structured and highly connected, which strongly suggests the existence of social communities. Researchers have discussed the construction [12] and the incentive mechanism design [13] of wireless community networks for years. However, no large-scale wireless community network with long-term impact has been built or observed. For the first time, we observe and provide the evidence of such social communities based on the geographical access patterns of WiFi users. This provides great opportunities for various aspects of network optimization. For instance, social media contents can propagate not only in online social networks but also in geographical WiFi user networks through cascaded physical communications relayed by users and APs [14]; the communities can help optimize the crowdsourcer recruitment for mobile crowdsourced sensing [15].

(1) Miscellaneous statistics	
Number of nodes	64,142
Number of edges	4,232,061
Average degree	65.98
Average clustering coefficient	0.91
Diameter (longest shortest path)	22
90-percentile diameter	8
Average shortest path length	5.9024

(2) *k*-clique clustering		
k	Number of clusters	Average cluster size
10	893	77
20	774	86
50	532	111
100	258	152
150	103	199

Table 1. Statistics and clustering results.

Based on our observations, we believe that WLANs are shifting from complementary access technology to primary access technology in today's metropolises and providing rich information for the underlying social relations among WiFi users.

Conclusion

In this article, we have conducted a large-scale measurement study on wireless networks in modern metropolises. Based on an extensive dataset, we have first investigated the coverage of today's WiFi networks, and by classifying wireless APs, we uncovered their rich geographical features. We further studied the access patterns of the WiFi users during weekdays and analyzed the polarization of user-AP connection durations. Finally, and most importantly, we identified the existence of geographical social networks and community structures among WiFi users. To the best of our knowledge, this is the first time such autonomous wireless communities have been observed in large-scale networks. Based on our observations, we believe that WLANs are shifting from complementary access technology to primary access technology in today's metropolises and providing rich information for the underlying social relations among WiFi users.

References

[1] K. Lee *et al.*, "Mobile Data Offloading: How Much Can Wi-Fi Deliver?" *IEEE/ACM Trans. Networking*, vol. 21, no. 2, 2013, pp. 536–50.

[2] Cisco, "Cisco Service Provider Wi-Fi: A Platform for Business Innovation and Revenue Generation," tech. tep. doc. ID 1458684054151755, 2012.
[3] A. R. Ekti *et al.*, "On the Traffic Offloading in Wi-Fi Supported Heterogeneous Wireless Networks," *J. Signal Processing Systems*, Springer, vol. 83, no. 2, 2016, pp. 225–40.
[4] M. Afanasyev *et al.*, "Usage Patterns in an Urban WiFi Network," *IEEE/ACM Trans. Networking*, vol. 18, no. 5, 2010, pp. 1359–72.
[5] X. Chen *et al.*, "Network Performance of Smart Mobile Handhelds in a University Campus WiFi Network," *Proc. 2012 ACM Internet Measurement Conf.*, 2012, pp. 315–28.
[6] J.-P. Laulajainen *et al.*, "Study of Youtube Demand Patterns in Mixed Public and Campus WiFi Network," *Proc. 2014 IEEE Int'l. Wireless Commun. Mobile Computing Conf.*, 2014, pp. 635–41.
[7] S. Biswas *et al.*, "Large-Scale Measurements of Wireless Network Behavior," *Proc. ACM SIGCOMM. ACM*, 2015, pp. 153–65.
[8] L. Wu *et al.*, "Mining Face-to-Face Interaction Networks using Sociometric Badges: Predicting Productivity in an IT Configuration Task," available at SSRN 1130251, 2008.
[9] S. Vural, D. Wei, and K. Moessner, "Survey of Experimental Evaluation Studies for Wireless Mesh Network Deployments in Urban Areas Towards Ubiquitous Internet," *IEEE Commun. Surveys & Tutorials*, vol. 15, no. 1, 2013, pp. 223–39.
[10] C. Haythornthwaite, "Social Network Analysis: An Approach and Technique for the Study of Information Exchange," *Library & Info. Science Research*, vol. 18, no. 4, 1996, pp. 323–42.
[11] D. J. Watts and S. H. Strogatz, "Collective Dynamics of Small-World Networks," *Nature*, vol. 393, no. 6684, 1998, pp. 440–42.
[12] P. A. Frangoudis, G. C. Polyzos, and V. P. Kemerlis, "Wireless Community Networks: An Alternative Approach for Nomadic Broadband Network Access," *IEEE Commun. Mag.*, vol. 49, no. 5, May 2011, pp. 206–13.
[13] M. H. Manshaei *et al.*, "On Wireless Social Community Networks,"*Proc. 27th IEEE INFOCOM*, 2008.
[14] Z. Wang *et al.*, "Propagation-Based Social-Aware Replication for Social Video Contents," *Proc. 20th ACM Int'l. Conf. Multimedia*, 2012, pp. 29–38.
[15] R. K. Ganti, F. Ye, and H. Lei, "Mobile Crowdsensing: Current State and Future Challenges," *IEEE Commun. Mag.*, vol. 49, no. 11, Nov. 2011, pp. 32–9.

Biographies

Lei Zhang [S'12] received his B.E. degree from the Advanced Class of Electronics and Information Engineering, Huazhong University of Science and Technology, Wuhan, China, in 2011, and his M.S. degree from Simon Fraser University, Burnaby, British Columbia, Canada, in 2013. He is currently a Ph.D. student at the School of Computing Science, Simon Fraser University. He is a recipient of the C.D. Nelson Memorial Graduate Scholarship (2013). His research interests include mobile cloud computing, social media, and multimedia systems and networks.

Liting Zhao received her B.E. degree in information security in 2015 from Beijing University of Posts and Telecommunications, China. She is currently a Master's candidate in computer science at Tsinghua University. Her research interests focus on urban data mining and user behavior analysis.

Zhi Wang [S'10, M'14] received his B.E. and Ph.D. degrees in computer science in 2008 and 2014, both from Tsinghua University, Beijing, China. He is currently an assistant professor at Tsinghua University. His research areas include online social networks, mobile cloud computing, and large-scale multimedia systems. He is a recipient of the China Computer Federation Outstanding Doctoral Dissertation Award (2014), ACM Multimedia Best Paper Award (2012), and MMM Best Student Paper Award (2015).

Jiangchuan Liu [S'01, M'03, SM'08, F'17] is a professor in the School of Computing Science, Simon Fraser University, and an NSERC E.W.R. Steacie Memorial Fellow. He was an EMC-Endowed Visiting Chair Professor of Tsinghua University (2013–2016). He is a co-recipient of the inaugural Test of Time Paper Award of IEEE INFOCOM (2015), ACM SIGMM TOMCCAP Nicolas D. Georganas Best Paper Award (2013), and ACM Multimedia Best Paper Award (2012).