

Citywide Mobile Internet Access Using Dense Urban WiFi Coverage

Maria Eugenia Berezin, Franck Rousseau, Andrzej Duda
Grenoble Institute of Technology, CNRS Grenoble Informatics Laboratory UMR 5217, France
{berezin, rousseau, duda}@imag.fr

ABSTRACT

We investigate if it is feasible to use the WiFi coverage in urban areas for mobile Internet access and which type of applications can benefit from the Internet access provided by the already deployed WiFi Access Points (APs). Nowadays, most smartphones and other mobile handsets are WiFi-enabled. Moreover, mobile Internet data traffic is expected to grow significantly in the next few years. WiFi access is an interesting alternative to cellular networks, as it is a widespread wireless technology, provides high data rates and has a low deployment cost. In this paper, we analyze the characteristics of WiFi coverage and connectivity of mobile users using different mobility speeds and varying some AP settings. These results allow us to analyze different applications that can be supported and the challenges to face to create a citywide WiFi network given the current infrastructure composed of residential WiFi APs and hotspots.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless Communication

General Terms

Design, Performance

Keywords

Wireless Networks, Mobile Internet Access, WiFi, Mobile Clients

1. INTRODUCTION

With the proliferation of mobile Internet-enabled devices, Internet connection is expected by users anywhere and anytime. Indeed, many forecasts predict an exponential growth for mobile data traffic [8]. Cellular broadband networks are therefore facing the problems of traffic congestion and network capacity. But improvements to these wireless networks

are expensive and new technologies such as 4G still have issues with their service performance.

WiFi networks are a viable solution to reduce the use of 3G networks, offering high bandwidth and a cheap infrastructure [13]. In fact, mobile operators have already started to use WiFi for data offloading. Thus, dense deployment of WiFi Access Points (APs) could ensure Internet connectivity, falling back to 3G where service cannot be offered.

WiFi APs can be found almost everywhere nowadays: municipal wireless networks, cafés, hotels, airports, and private environments at home or work. However, only a fraction of all these APs are open for association to any user. In some other cases, such as residential APs, their owners are members of a "community network", and they can share their Internet bandwidth with other members (e.g. FON [9], FreeWifi [10], CableWifi [7]).

WiFi APs are generally unmanaged with a default configuration. Moreover, they are deployed indoors and in an unplanned manner, which can cause poor reception and interference between neighbor APs, and result in low throughput. They connect to the Internet over a broadband access, such as DSL, which provides high data rates.

In this paper, we consider the use of already deployed WiFi APs to provide urban mobile Internet access. If their coverage is sufficiently dense, users moving with different speeds may benefit from WiFi connectivity. The question is whether such an architecture is feasible, on what parameters depend its performance, and what type of applications can benefit from its services. We investigate these issues by simulating connectivity of mobile users in a city based on detailed traces provided by the Nokia Mobile Data Challenge [15].

2. MOBILE INTERNET ACCESS THROUGH WIFI

WiFi APs are largely deployed in urban areas so we can consider providing a seamless mobile Internet access based on dense WiFi coverage. Unplanned placement of APs may lead to the existence of some areas that lack coverage, however short disconnections may be tolerated if the duration of connectivity is sufficient. The applications that could be supported in a citywide WiFi network, without any modification in the infrastructure or the devices, are delay-tolerant applications, that can either wait between connections or alternate between a 3G and a WiFi link [4, 13]. Short, but frequent connectivity can be used for opportunistic sensing. For instance, traffic and road conditions can be shared by mobile phones and disseminated to vehicles [16]. Due to

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

UrbaNE'12, December 10, 2012, Nice, France.

Copyright 2012 ACM 978-1-4503-1781-8/12/12 ...\$15.00.

the large coverage, location-based services can use the great number of APs for positioning, even in indoor locations, without using a GPS.

When moving among APs, users need to benefit from short handoffs. Depending on the expected handoff delay, different strategies can be implemented to decrease this value: we can use different handoff and association algorithms (e.g. multiple simultaneous associations [18]) or manage mobility in the network [5], which results in instantaneous handoff.

Currently, there is no a homogeneous way of authentication, the selection of APs has to be done manually, and mobility (session transfer) is inexistent. Nevertheless, there are some efforts in this direction: the Hotspot 2.0 and the Next Generation Hotspot initiatives, based on the IEEE 802.11u standard. These specifications allow mobile phones to log into WiFi networks in a seamless way.

3. TRACE BASED EVALUATION

To evaluate the feasibility of such an architecture, we run simulations of mobile user connectivity in a city based on detailed traces provided by the Nokia Mobile Data Challenge (we provide the details of the dataset in the next section).

We have built a model of the city of Lausanne, as it was the city in the data set with the biggest number of APs. We ran several simulations to analyze and understand the characteristics of the WiFi coverage and connectivity provided by the deployed APs.

3.1 Simulation Scenario

We have focused on the central area of the city of Lausanne, where the APs distribution is the most dense. We have developed a simple simulator in which users move along random paths modeled as a Markov chain [19] to describe the probabilities of transition from one direction to another. The model reflects the behavior of users—the dataset contained the recorded GPS coordinates of users.

The simulator computes the position of a mobile user with the granularity of 1 m and verifies if a handoff (or an association, if the device has lost the WiFi connectivity) has to be performed. It computes the distance of the user to the current AP and uses the formula of the Log-distance Path Loss Model [17] to obtain the current AP signal strength. If needed, the simulator performs a handoff and changes the AP.

The simulator takes into account the following parameters:

Range: We used three different coverage ranges for an AP: 20 m, 50 m, and 100 m. The two former values are common ranges for indoor APs, and the latter is a common range for an outdoor AP. As 50 m is the most common range for an AP, we present the results for this value.

Association: As an ideal situation, we assumed that every AP in the city are available for association (we call this parameter “All APs”). We also ran simulations using only “Open APs” (i.e. not protected by any security mechanisms) from the traces (nearly 17% of all APs). Nevertheless, the data from Nokia did not specify if any of the “Open APs” had a captive portal authentication.

Handoff durations: A handoff is a process in which the client scans all channels looking for new APs, authenticates and associates with a chosen AP, and finally obtains an IP

address. During that time, the client cannot send or receive traffic so that the handoff duration has an important impact on the delay and packet loss that applications may tolerate.

We simplified the handoff process by modeling it as a fixed delay. We have assumed the following values for the handoff duration: 0 seconds, 150 ms, 1 s, 2 s, and 5 s. The first value corresponds to an instantaneous handoff, for instance the case of a network supported mobility management. The next value, 150 ms, is the maximum acceptable latency in end-to-end communications for VoIP applications. The latter values are commonly measured handoff durations.

Handoff strategy: When the device is not associated with any AP, it chooses the AP with the strongest signal strength. When the current AP’s signal strength (SS1) is weaker than a fixed Threshold, the client attempts to associate with another AP whose signal strength (SS2) is greater than SS1. We used a -90 dBm threshold, which is a common limit for the lowest data rates: 1 Mb/s.

User speed: We wanted to represent different speeds of a mobile user. We used the following values: walking (1 m/s), by bicycle (5 m/s \cong 18 km/h), by bus (11 m/s \cong 40 km/h), and by car (20 m/s \cong 70 km/h).

Paths: We generated 10,000 different paths using the Markov chain mobility model. Each path has segments of 50 m and the transition probabilities are applied after each segment to change direction. The total distance of each path is 3,600 m. As we wanted to recreate this path in the downtown part of the city of Lausanne, the path was centered in the densest area.

3.2 Results

We define the following performance indices:

Temporal Coverage: period of time during which a user stays in a WiFi coverage area [13].

Connection: period of time during which the user is associated with a same AP.

Disconnection: period of time during which the user is out of the range of any WiFi signal.

Internet access session: period of time during which the user is associated with APs and is able to send and receive Internet traffic (i.e. is not doing a handoff at that time), until a disruption of communications, due to the loss of WiFi coverage.

We have measured the mean temporal coverage for different APs ranges in Figure 1 using “All APs” and “Open APs”. When analyzing the results of “All APs”, a common range of 50 m gave a 70% of mean temporal coverage. We obtained less than 55% of temporal coverage for APs with 20 m of range, which are indoor APs surrounded by obstacles that decrease its range, and for 100 m, we obtained more than 80%. These values (20 m and 100 m) represent the minimum and maximum range for an AP, so we can say that 55% and 85% are the mean values for the minimum and maximum temporal coverage, when every AP is available for association.

When comparing with the results for “Open APs”, we can see that associating with “All APs” provides an improvement of around 10% in the temporal coverage. This means the WiFi coverage is already dense even with only 20% of “Open APs”.

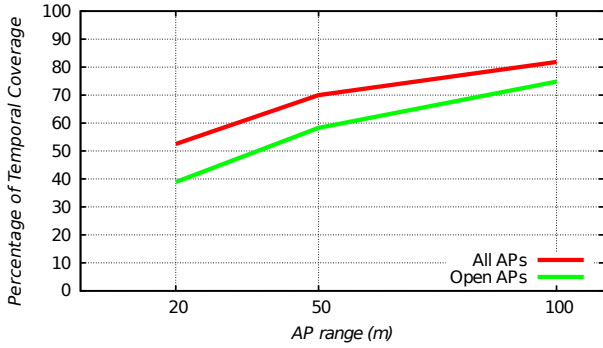


Figure 1: Percentage of temporal coverage in average of a user’s path, with “All APs” and “Open APs” for association, using different AP ranges

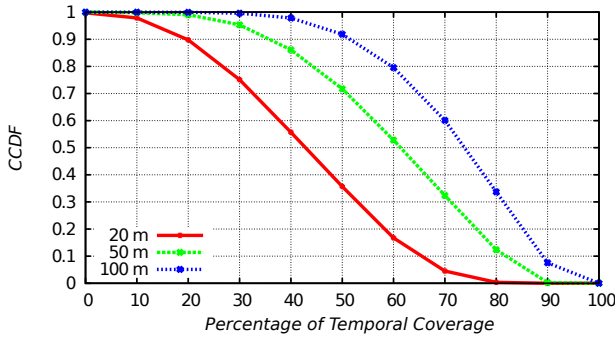


Figure 2: Complementary CDF of the percentage of the temporal coverage of a user’s path, using different AP ranges

Figure 2 plots the complementary cumulative distribution function (CCDF) of the temporal coverage to compare the distribution of temporal coverage of the 10,000 simulated mobile users for each AP range. Clearly, the longer the range, the better the WiFi coverage.

We then performed the simulations with an AP range of 50 m and different speeds for the mobile users. In Figure 3, we present the mean duration of a connection with the same AP while assuming a handoff duration of 0 s. If the user is walking, the mean connection has a duration of 130 s (more than 2 minutes). However, faster speeds have shorter connections of less than 30 s.

In Figure 4, we show the mean duration of an Internet access session for different mobile user speeds and different handoff durations. A user can have a mean session of 572 s (more than 9 minutes) when walking. At bicycle speed, the session lasts, in average, 114 s (nearly 2 minutes). This duration decreases to 52 s (nearly 1 minute) by bus, and to 28 s by car.

As we can observe, if the user is walking, the duration of the handoff does not impact as much as traveling by car. Because the handoff duration is fixed, i.e. does not depends on the speed, the short session by car is dramatically impacted by the handoff duration, as opposed to walking. The Internet access session of a walking user decreases from 572 s to 551 s when the handoff duration increases. In contrast, the total session by car decreases from 28 s to 8 s.




Figure 3: Mean connection duration with a same AP for different user speeds




Figure 4: Mean Internet access session duration (logarithmic scale) until disconnection for different handoff durations and user speeds

We present the total Internet access session in Figure 5. By car, the time of the Internet access session can decrease dramatically from 70% to 20% when the handoff duration is longer. As stated before, longer handoff delays reduce the duration of the useful connectivity.

Figure 6 and Figure 7 present the disconnection results. The first figure shows the mean disconnection duration. We can observe that faster speeds have shorter disconnection durations. On the other hand, walking has a long duration, of almost 3 minutes, without WiFi connection. Figure 7 shows the Cumulative CDF of the disconnection durations of all the simulations, for each user speed. For walking users, 40% of the disconnections have more than 100 s (around 2 minutes) of disconnection.

4. DATA ANALYSIS

In this section, we describe the data used in our simulations. We characterize some APs properties and we detail the corrections to the APs positions. Finally, we describe a mobility pattern generated from the GPS traces of the mobile users.

4.1 Data Description

We obtained the data from the Mobile Data Challenge organized by Nokia [15]. The provided data was collected from the smartphones of 38 users during more than one year.




Figure 5: Percentage of total Internet access session in average during a user’s path for different handoff durations and user speeds




Figure 6: Mean disconnection duration for different user speeds

We used the GPS, the WLAN, and the WLAN location data sets [12]. The GPS traces contained the GPS information of the mobile phone. Due to energy constraints, GPS readings were mostly available during outdoor movements, with a sampling period of 10 seconds.

The WLAN traces contained the APs discovered during a WiFi scan performed by the device, and the WLAN location traces contained the GPS location of some APs. These traces had a sampling period between 60 and 900 seconds. Some of the WLAN fields were anonymized, such as the MAC address and the SSID, to ensure privacy of the participants.

4.2 APs characteristics

We have extracted the WLAN devices in infrastructure operational mode from the WLAN traces, and we found nearly 127,000 unique APs distributed over Switzerland. The following information was present in the current data set: time, anonymized MAC address of WLAN device, anonymized SSID, signal level, channel, encryption type, operational mode.

When inspecting the traces, several APs changed the values of their settings over time, such as the channel number and the type of security, along the whole data set. Consequently, for the following plots, we considered all the different values of channels and types of security that an AP could have had during the data collection.




Figure 7: Complementary CDF of the disconnection duration for different user speeds

We analyzed the following characteristics:

Signal strength: In Figure 8(a), nearly 40% of the records have a signal strength stronger than -70 dBm, a common value for data rates of 54 Mb/s¹. In 70% of the total records, the device could connect with a signal strength stronger than -85 dBm, having a data rate of at least 11 Mb/s.

Channels: The smartphones scanned in the 2.4 Ghz band, so the discovered APs are listening to a channel between 1 and 13. In Figure 8(b), we observe that 65% of the APs use the channels 1, 6 and 11. The use of orthogonal channels decreases the interference among neighbor APs, but if some of the neighbors are in the same frequency, partially overlapping channels [14] can be a better solution.

Security: In Figure 8(c), less than 20% of APs do not have any protection, i.e., they are open for association. Open APs already share their bandwidth and provide Internet access for any guest. Community Networks, such as FreeWifi, appear as “Open” in the scan list, but require valid credentials through the use of a web portal, to identify the guest.

4.3 APs distribution

The WLAN location data set contained the coordinates of nearly 2,500 APs. The following information was present in the current data set: time, anonymized MAC address, longitude, latitude.

The location of an AP appeared several times in the traces, and sometimes these values had a difference of a few meters. Therefore, we calculated the centroid (arithmetic mean) of these coordinates to obtain an approximate position.

The geographic coordinates of the APs were expressed in decimal degrees, e.g. 45.2002°, 5.7222°. Most of the coordinates had between 7 and 11 decimal places. But some of the coordinates appeared truncated, with less than 3 decimal places. We removed these coordinates as they had a difference of more than 1 km with other coordinates with more decimal places. The result of applying this heuristic can be seen in Figure 9. As we can see on the left map, all the APs were positioned over a diagonal.

Finally, as the WLAN location data set contained only the position of 2,500 APs (less than 2% of total discovered APs), we used the coordinates obtained previously from the

¹Measured experimentally, see <http://www.tp-link.com/en/products/details/?model=TL-WN721N#spec>




Figure 8: APs characteristics



Figure 9: APs distribution: difference between average coordinates (left) and removing the truncated coordinates (right)

WLAN location data set in order to deduce the coordinates of the other APs that appear only in the WLAN traces. As a result, we increased the total number of known locations to nearly 33,500 APs (26% of total APs).

We assumed that, when a device performs a scan, all discovered APs are near that location within a radius of 100 m (common coverage range for an outdoor AP). To create different values for each of the nearby APs, we took the first 3 digits of the known location and we added 5 random last digits. This heuristic obtained a new different location less than 30 m away from the original coordinates, e.g.:

known coordinates: [45.2002 °, 5.7222 °]
 new coordinates: [45.20040053 °, 5.72220417 °]
 distance: 47.3 m.

4.4 Mobility pattern

The GPS data set provided the coordinates of the user when the device was outdoors. The following information was present in the current data set: time, geo-location (altitude, longitude, latitude), speed, heading, accuracy and Dilution of precision (DOP), time since GPS system started.

We wanted to extract a mobility pattern from the data set and to create a more realistic mobility pattern than the Random Walk Model. We used a Markov chain [19] to describe the probabilities of transition from one direction to another.

We processed all the heading fields that appeared in the GPS data set. Heading is the angle in degrees where the device is pointing to compared to the North. We transformed the degrees in 8 directions: north, north-east, east, south-east, south, south-west, west, north-west. We used

three directions in a chain to obtain finer-grained mobility characterization, e.g.:

(previous, current, next) = (south, south, east)

We computed the transition probabilities using the following formula:

$$P(next|current, previous) = \frac{\sum P(curr, prev, next)}{\sum P(curr, prev)}$$

We obtained that in more than 75% of the cases the user moves in the same direction as in the two previous steps. We can see two paths in Figure 10: one created using the Random Walk model (left) and the other using the transition probabilities from the GPS data (right). There are longer straight segments in the latter than in the former.




Figure 10: Paths created using the Random Walk model (left) and using the transition probabilities from the GPS data (right).

5. RELATED WORK

Several works have considered the deployment and usage of WiFi APs in cities. We present an overview of the different aspects of the subject:

APs Characterization: The first study examined statistics of over 5 million APs collected through wardriving. It characterized the default settings (e.g. SSID), location, and density of APs found in several metropolitan cities [11]. Similarly, other works quantified the impact of interference on end-client performance [2], analyzed the coverage and duration of connectivity, and measured the performance of TCP uploads [6].

Large-scale WiFi networks: Inspired by the popularity of WiFi technology, municipalities, non-profit and private entities have deployed free WiFi hotspots in cities. Some examples of these open urban WiFi networks are MIT Roofnet, Google WiFi Network, Madison MadMesh. Despite the efforts, these networks suffered mainly from packet loss [1], or

coverage holes [3]. These wireless networks need more investment in infrastructure, and are far from having a good performance.

WiFi offloading: Prior works have studied mobile data offloading as an alternative to cellular network upgrades, as WiFi hotspots are cheaper and easier to deploy. Balasubramanian et al. studied the availability of WiFi and its performance and proposed a predictor of offload capability [4]. Lee et al. analyzed the traces of 100 mobile phones during 2 and 1/2 weeks to predict the time of transfer based on mobility patterns [13].

6. CONCLUSIONS

The questions we answered in this paper are the following: can we use the already deployed WiFi APs for a city-wide mobile Internet access? What are the characteristics of the WiFi connection? What type of applications can benefit from this wireless technology? And finally, what are the challenges to face if we want to take advantage of this infrastructure?

We used the data collected from smartphones in the city of Lausanne. We analyzed the WiFi APs distribution and their characteristics. We ran several simulations, varying mobile user speeds and the APs properties, such as range, association, and handoff duration. We measured the WiFi coverage and connection, Internet access session, and the disconnection duration.

The results showed that WiFi coverage is large and connectivity has different characteristics to exploit. We proposed several applications that could be used in this wireless network, and the challenges that should be tackled in order to run these applications in an effective way.

7. ACKNOWLEDGEMENTS

This work was supported by the French National Research Agency (ANR) project ELAN under contract ANR-08-VERS-008.

8. REFERENCES

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '04, pages 121–132, New York, NY, USA, 2004. ACM.
- [2] A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self-management in chaotic wireless deployments. In *Proceedings of the 11th annual international conference on Mobile computing and networking*, MobiCom '05, pages 185–199, New York, NY, USA, 2005. ACM.
- [3] A. Arjona and S. Takala. The Google Muni Wifi Network—Can it Compete with Cellular Voice? In *Telecommunications, 2007. AICT 2007. The Third Advanced International Conference on*, page 11, may 2007.
- [4] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. N. Levine, and J. Zahorjan. Interactive WiFi connectivity for moving vehicles. In *Proceedings of the ACM SIGCOMM 2008 conference on Data communication*, SIGCOMM '08, pages 427–438, New York, NY, USA, 2008. ACM.
- [5] M. Berezin, F. Rousseau, and A. Duda. Multichannel Virtual Access Points for Seamless Handoffs in IEEE 802.11 Wireless Networks. In *Vehicular Technology Conference (VTC Spring), 2011 IEEE 73rd*, pages 1–5, may 2011.
- [6] V. Bychkovsky, B. Hull, A. Miu, H. Balakrishnan, and S. Madden. A measurement study of vehicular internet access using in situ Wi-Fi networks. In *Proceedings of the 12th annual international conference on Mobile computing and networking*, MobiCom '06, pages 50–61, New York, NY, USA, 2006. ACM.
- [7] CableWifi. <http://www.cablewifi.com/>.
- [8] Cisco White Paper. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2010-2015, February 2011.
- [9] FON. <http://www.fon.com/>.
- [10] FreeWifi. <http://www.free.fr/>.
- [11] K. Jones and L. Liu. What Where Wi: An Analysis of Millions of Wi-Fi Access Points. In *Portable Information Devices, 2007. PORTABLE07. IEEE International Conference on*, pages 1–4, may 2007.
- [12] J. K. Laurila, D. Gatica-Perez, I. Aad, J. Blom, O. Bornet, T.-M.-T. Do, O. Dousse, J. Eberle, and M. Miettinen. The Mobile Data Challenge: Big Data for Mobile Computing Research. In *Proc. Mobile Data Challenge by Nokia Workshop, in conjunction with International Conference on Pervasive Computing, Pervasive*, 2012.
- [13] K. Lee, J. Lee, Y. Yi, I. Rhee, and S. Chong. Mobile data offloading: how much can WiFi deliver? In *Proceedings of the 6th International Conference, Co-NEXT '10*, pages 26:1–26:12, New York, NY, USA, 2010. ACM.
- [14] A. Mishra, E. Rozner, S. Banerjee, and W. Arbaugh. Exploiting partially overlapping channels in wireless networks: turning a peril into an advantage. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, IMC '05, pages 29–29, Berkeley, CA, USA, 2005. USENIX Association.
- [15] Mobile Data Challenge by Nokia Workshop. <http://research.nokia.com/mdc>.
- [16] P. Mohan, V. N. Padmanabhan, and R. Ramjee. Nericell: rich monitoring of road and traffic conditions using mobile smartphones. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, SenSys '08, pages 323–336, New York, NY, USA, 2008. ACM.
- [17] T. Rappaport. *Wireless Communications: Principles and practice*. Prentice Hall, 2002.
- [18] H. Soroush, P. Gilbert, N. Banerjee, M. D. Corner, B. N. Levine, and L. Cox. Spider: improving mobile networking with concurrent wi-fi connections. In *Proceedings of the ACM SIGCOMM 2011 conference*, SIGCOMM '11, pages 402–403, New York, NY, USA, 2011. ACM.
- [19] J. Yoon, B. D. Noble, M. Liu, and M. Kim. Building realistic mobility models from coarse-grained traces. In *Proceedings of the 4th international conference on Mobile systems, applications and services*, MobiSys '06, pages 177–190, New York, NY, USA, 2006. ACM.