StreetSense: Effect of Bus Wi-Fi APs on Pedestrian Smartphone

Sehyun Bae
Seoul National University[1]
kbird@popeye.snu.ac.kr

Daehyun Ban
Samsung Electronics
daehyun.ban@samsung.com

Dahyeon Han
Seoul National University
hann862@popeye.snu.ac.kr

Jiyoung Kim
Seoul National University
beyondj@popeye.snu.ac.kr

Kyu-haeng Lee
Seoul National University
khlee@popeye.snu.ac.kr

Sangsoon Lim
Samsung Electronics
ssoon.lim@samsung.com

Woojin Park
Samsung Electronics
woojin1.park@samsung.com

Chong-kwon Kim
Seoul National University[2]
ckim@snu.ac.kr

ABSTRACT

Recently, we have received a growing number of reports that complain about poor and unstable internet connections at bus stops in metro Seoul. Careful analyses led us to conclude that Wi-Fi APs equipped on buses instigate the trouble. According to the ambitious free Wi-Fi expansion plan by the city of Seoul, public buses started to equip Wi-Fi APs. As buses with APs stop and go, they actualize intermittent connection opportunities to riders waiting at the bus stops. However, the connection durations are too short such that bus APs are a nuisance rather than a convenience. We collected the basic statistics such as AP inter-arrival and sojourn times and measured link level performance metrics. We observed the effect of frequent frame losses on the TCP congestion control and eventually on the TCP throughput. We also measured the performance of applications such as PLT (Page Load Time). The measurement results showed that passing APs are useful only for some applications in very limited situations while they are virtually useless and just irritations in many cases. We also discovered that poor Wi-Fi connections pervert MPTCP; MPTCP performs worse than the generic single path TCP over the LTE network. We expect that our results will be used as the reference data in redesigning Wi-Fi offloading mechanisms as well as in planning and deploying urban Wi-Fi networks.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Reliability, availability, and serviceability; Performance attributes; C2.3 [**Computer-Communication Networks**]: Network Operations-Public networks

Keywords

AP, WiFi, LTE, MPTCP, Smartphone, Bus

1. INTRODUCTION

Aided by rapid technology advancements, nomadic convenience, and repeated cost reductions, we have observed widespread diffusion of Wi-Fi networks as one of most popular Internet access techniques. At the time of writing, more than 180 Million Wi-Fi APs are deployed worldwide [1]. Started as an indoor novelty that replaces unwieldy cables, Wi-Fi has expanded its horizon as outdoor amenities that transform urban landscape. Numerous grassroots movements to offer private Wi-Fi APs to the public as well as citywide or even nationwide campaigns to install free Wi-Fi hotspots have been reported. For example, Google funded the municipal Wi-Fi network in Mountain View, CA [2] and NY City announced an ambitious plan to construct free Wi-Fi cities [3].

In addition to open and free Internet access, Wi-Fi networks have been used by cellular network companies as a traffic offloading system. The demands for mobile data traffic have increased explosively and this tendency is expected to continue in the future [4]. Even though cellular operators keep expanding the network infrastructure, the capacity increments failed to overtake the demands that are estimated to explode at a 69% annual growth rate. Because Wi-Fi network construction is much more economical than cellular network and Wi-Fi can be easily set up without careful planning, operators multiply Wi-Fi hotspots to alleviate the data overloading conditions on the cellular infrastructure. Many operators designate Wi-Fi as the default connection; when both cellular and Wi-Fi are available, smartphones prefer to connect to Wi-Fi. Aggregated with the pricing strategy that Wi-Fi is complimentary while cellular is still prohibitively expensive, default use of Wi-Fi appeases subscribers also. In addition, the

operators installed Wi-Fi hotspots at locations such as cafés, libraries, subway stations and other gatherings where relative location of APs and subscribers are static. In static environments, Wi-Fi hotspots of constrained coverage render satisfactory performances.

Wi-Fi hotspots are now being deployed to moving mass transportation including subway trains and public buses. For example, the city of Seoul recently announced a plan that they will install Wi-Fi APs to every bus in Seoul such that passengers can access the Internet while on board [5]. Even before the ambitious plan rolled out, several metro Seoul buses have already started to be geared with APs. While APs on bus provide convenience to passengers, they inflict unforeseen side effects to riders and pedestrians at bus stops. As buses with APs make stops for a short duration to let passengers to get on/off, prioritized Wi-Fi provokes smartphones to switch connections from cellular networks to the APs. However the connection durations generally are less than 30 seconds which is too short for finishing transactions. The passengers started to report poor Internet connections in 2014 when buses started to equip APs and the number of complaints to Samsung Electronics, a major smartphone supplier, has increased as the number of AP installed buses increase. Our quality assurance team investigated the problem and discovered that the APs on passing buses are the root of the problem.

Like many large cities in Asia, public transportation is the most popular means of mobility in Seoul. More than 5.8 million citizens ride buses daily and busy bus stops serve tens of scheduled regular bus routes. Some busy bus stops have fixed Wi-Fi APs but usually bus stops do not have fixed APs. While waiting for their buses, many riders (See Figure 1 that shows riders at a busy bus stop in Gangnam) use their smartphones surfing the Web, watching streaming videos, etc. At a busy bus stop, several buses arrive every minute and it is not unusual to observe more than ten buses are queued for getting passengers on/off board. Normally, buses stay at a bus stop for 30~90 seconds and this triggers a connection switch to the APs on passing buses. After a few, and sometimes no frame transmissions, currently available APs disappear and new APs turn up. This pattern repeats and causes serious irritations to riders some of who are running time-critical jobs such as searching restaurants or reading electronic maps.

This paper presents the measurement results that we have collected for three weeks in April 2015 at bus stops in Seoul. Busy bus stops serve more than 20 bus routes each of which has inter-service time between 5 to 10 minutes during rush hour. We measured basic statistics such as AP (= bus) inter-arrival time, sojourn time, signal strength changes in the time domain. We also measured TCP and UDP throughputs and analyzed how intermittent connections affect the rate adaptation at the MAC layer and congestion control at the TCP layer. Because user perceived QoEs (Quality of Experience) such as PLT (Page Load Time) and buffering rate might be more important performance metric than network level metrics, we thoroughly investigate QoE metrics also. Finally, we investigated the performance of MPTCP that use both cellular network and intermittent Wi-Fi connections.

Our measurement results illuminated the adverse effects of poor and intermittent Wi-Fi connections. In most cases, the connection duration over bus Wi-Fi is less than 30 seconds. Only a handful of packets can be transferred during the short connection interval. Because TCP congestion control prevents the use of available bandwidth instantly, the effect of the short connection time amplifies at the TCP layer. Poor TCP performance is again directly translated to poor QoE (Quality of Experience) at the application domain. The PLT and download speed are two or three times worse than the LTE connections. The performance of MPTCP that spans over both LTE and Wi-Fi is particularly discouraging; its performance is about three times worse than the generic single path TCP over LTE networks.

2. RELATED WORK

Wi-Fi networks have been widely deployed as the tether-less last mile as well as supplementary networks offloading cellular data traffic. Several cities and municipals provided free Wi-Fi access in their communities [2, 3, 6]. [11] observed the Wi-Fi usage patterns in the Google's Mountain View community network and identified three usage patterns according to device types [11]. The usage pattern of smartphones is characterized by the shortest session length and the least bandwidth use. However, as smartphones are rapidly replacing PCs[7], the usage pattern of smartphones may change also.

[24, 19] examined the feasibility of accessing Wi-Fi APs from moving vehicles. These early studies – even though performed in synthetic environments with only one or a few APs – confirmed the possibility of dynamic connections to Wi-Fi APs from moving vehicles and triggered numerous further investigations. Early research efforts focused on measuring the performance of Wi-Fi connections from moving vehicles in real-world environments. Later the research is extended to Wi-Fi offloading. Particularly, an early study [13], which measured the performance of Wi-Fi connections from moving vehicles in a metropolitan area, revealed both the feasibility and the limitation of Wi-Fi accesses from moving vehicles: the average connection time is 13 seconds and the bandwidth is 30 KB/s.

The explosive growth of mobile data demands superimposed with the confined radio spectrum and the high CAPEX of cellular infrastructure prompted cellular companies to employ Wi-Fi networks as a mobile traffic offloading system. The 3GPP standard includes architectures I-WLAN [8] and IFOM [9] that integrate cellular networks and Wi-Fi networks seamlessly. Wi-Fi offloading has been one of the most intensively studied topics handling both non-vehicle environments [20, 21] and vehicle moving cases[12]. Usually Wi-Fi offloading attains less performance gains in vehicular mobility scenarios than in stationary or slow moving scenarios due to more dynamic changes in channel states and shorter connection durations. [15, 25] provided technical overview and challenges of vehicular Wi-Fi offloading.

[20] dealt with stationary and walking mobility scenarios with 100 smartphone users in metropolitan areas. They showed that data offloading yields greater performance enhancement in static environments than vehicle moving cases. However, another study[21] with 200 students in a University campus well equipped with Wi-Fi hotspots showed the limited applicability of Wi-Fi offloading due to the fact that Wi-Fi hotspots are optimized for laptops rather than smartphones.

[12] conducted an extensive performance study of Wi-Fi offloading in three cities in the USA. Their results showed the poor performance of Wi-Fi networks mainly due to limited availability. [17] performed a direct head-to-head comparison between a 3G network and Wi-Fi network in a vehicular communication scenario in Long Island, NY. They discovered that 3G and Wi-Fi networks present different and complementary characteristics; Cellular networks provide wide and stable coverage while Wi-Fi networks furnish intermittent but higher bandwidth. However, the improved

cellular network infrastructure and the introduction of new techniques such as LTE and LTE-A might change the technical landscape. Several studies indicate that cellular networks offer throughput comparable to or even larger than Wi-Fi while Wi-Fi still maintains the shorter latency [15, 16].

Many researchers proposed mechanisms to enhance the performance of Wi-Fi offloading for fast moving vehicles. [12] devised a fast switching mechanism called Wiffler that enables significant performance gains for delay tolerant traffic. [26] proposed a mechanism that predicts the mobility of vehicles and performs data prefetching to projected hotspots to maximize Wi-Fi traffic offloading for both delay sensitive and delay tolerant applications. ATOM [22] selects Wi-Fi or cellular interfaces intelligently to optimize user QoE. Cedos [23] is another mechanism that maximizes Wi-Fi use for delay tolerant traffic.

Generalization and proliferation of Wi-Fi networks naturally triggered the adoption of MPTCP [18] in recent smartphones such as the Samsung Galaxy S6 [10]. [14] explored the performance of MPTCP with smartphones focusing on the impact of several factors such as flow size, rate/route control algorithms and path characteristics on the performance. Their results showed that MPTCP generally improved the application level performance. Recently, extensive experiments [16] carried out in 167 countries over a 6-month period showed that MPTCP is good for elephants but not very effective for short flows necessitating clever mechanisms that adaptively select the best interface and optimal application of MPTCP.

Unlike the previous studies that handled scenarios where vehicles move around stationary Wi-Fi APs, we study the performance of Wi-Fi offloading in an environment where stationary subscriber accesses moving APs mounted on buses. Initially, we thought that these two environments are similar and expected similar results. However, further investigations showed that there are significant differences. Passengers in moving vehicles can use relatively long and stable Wi-Fi connections when the vehicles stop at traffic signals. On the other hand, subscribers waiting at bus stops are encountered with APs that dynamically appear and are available only for short durations that are too short for performing meaningful transactions.

3. EXPERIMENT ENVIRONMENT

Figure 1 shows passengers waiting at a busy bus stop in Gangnam, Seoul. Because the BIS (Bus Information System) informs the expected bus arrival times, riders know their waiting times and many run various smart phone applications such as web browsing, mobile massaging, video streaming and etc. Several large bus stops have fixed APs but many more do not have fixed APs. Even at bus stops with fixed APs, riders far from the fixed AP may find the APs on the bus near them provide stronger signals and they connect to a passing bus APs dynamically.

Figure 2 shows a typical network model employed by cellular operators. Cellular networks are carefully planned and deployed such that most subscribers at any locations receive adequate performance. Operators also deployed expansive networks of fixed Wi-Fi hotspots, but this paper concentrates on Wi-Fi APs on moving buses. KT (Korea Telecom), a major cellular operator in Korea, connects bus Wi-Fi APs to a P-GW via an LTE connection to provision large bandwidth and quick handovers between heterogeneous networks.

As mentioned earlier, we are focusing on how the Wi-Fi on moving buses affects the performance of riders waiting for buses. To do this, we have conducted extensive measurements at tens of different bus stops located in Seoul such as Gangnam, Yangjae, and Banpo for three weeks. We use Samsung Galaxy Note 3 smartphones and laptops to capture the traffic. Note that Samsung Galaxy Note 3 has an 802.11n Wi-Fi interface. We collect BSSIDs of connected APs, RSSI, data rate, and throughput with measured time from both smartphone and laptop while downloading files. We use the Wireshark packet analyzer on laptop computers for a detailed inspection of packet contents. In addition to basic performance metrics such as L2 level transmission rates, frame loss probabilities and L4 level throughput, we measure page loading time of some popular web sites with chrome browser on the smartphones to examine the QoE(Quality of Experience).

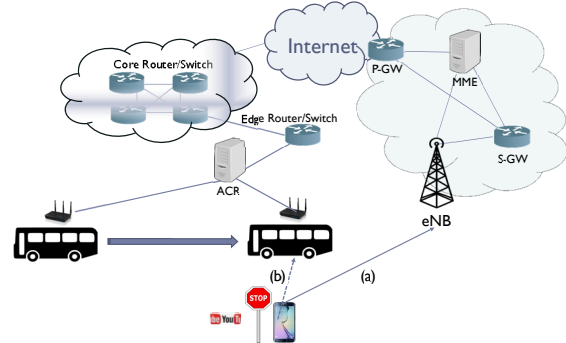**Figure 1. Passengers at a bus stop in Gangnam, Seoul**

Figure 2. A simplified model of the network. Users can access the Internet via both LTE (connection (a)) and Wi-Fi LTE (connection (b)). Wi-Fi is the default network and subscribers prefer to connect to Wi-Fi if possible.

4. MEASUREMENT RESULTS

Before delving into performance measurements, we first display some basic statistics to illustrate the dynamicity of bus arrivals and departures at bus stops in Seoul. We then examine the performance measures at various protocol layers. Particularly, we analyzed the effect of unstable Wi-Fi connections on the performance of rate adaptation at the MAC layer and congestion control at the TCP layer.

4.1 Basic Statistics

The important factors that affect the performance of waiting passengers include the arrival frequency of APs, their signal strength, sojourn time, and the number of simultaneous users. Among these statistics, the inter-arrival time and sojourn time determine the dynamicity of Wi-Fi AP connection opportunities. Figure 3 illustrates the inter-arrival time and sojourn time at a bus stop in Gangnam. We only show the result measured at one stop because inter-arrival times can be different depending on the number of bus routes at the stop. The arrival time of an AP is the time when we first receive a beacon from the AP and sojourn time is the time between the first beacon and the last beacon. Note that buses frequently queued up for a long line (frequently over 50 meters) and some buses passed the observation point without stopping. In this case, we may fail to detect the APs. We identify APs by their MAC addresses.
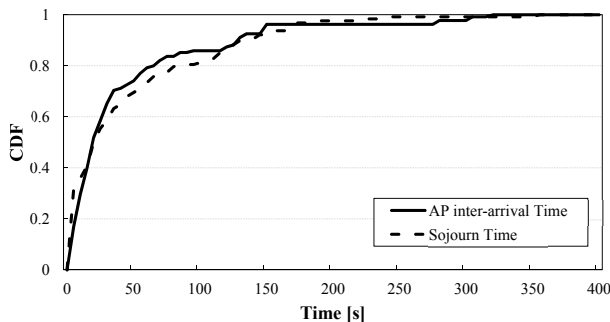

Figure 3. Inter-arrival time and sojourn time of APs on buses at a bus stop.

We can observe from the figure that the AP(=bus) inter-arrival time and the sojourn time are quite short: mean inter-arrival time and sojourn time are 22 seconds and 23 seconds, respectively.

Figure 4 shows one exemplary measurement result displaying the RSSI of signals from bus APs. Signals from different APs are shown in different colors. Short inter-arrival time and sojourn time define the high dynamicity of Wi-Fi AP availability. First, short inter-arrival and sojourn times force frequent switches between APs as well as between AP and LTE. Frequent switching incurs the switching overhead on mobile devices and also increases latency and energy consumption. The short sojourn times can inflict significant adverse effects also; TCP connections are set up and disconnected frequently, and short-lived TCP connections fail to fully utilize the available bandwidth due to the slow start characteristics of the TCP congestion control mechanism.

We argue that the Wi-Fi availability pattern at a bus stop is different from what we observed from moving vehicles. Moving vehicles usually pass APs on the streets unnoticed and access Wi-Fi APs mostly when they stopped for a traffic signal or at bus stations. Therefore, moving vehicles can realize the better utilization of Wi-Fi networks than riders at bus stops.

4.2 Link Layer and TCP Performance

Figure 5 shows the changes in transmission rates at the data link layer (Figure 5.(a) and (b)) and the TCP goodput (Figure 5.(c)) during an interval of (50, 230) of Figure 4. We use the same colors in Figures 4 and 5 to identify APs. During this period, the smartphone made connections to four different APs. The connection times are 77, 20, 22, and 25 seconds, respectively.

Many smartphones use the de factor standard rate adaptation scheme called ARF (Automatic Rate Fallback) or its variations. Except short durations in (80, 100), transmission rates change very dynamically. Note that ARF was not optimized for dynamic signal changes and short ephemeral connections that are the norm at bus stops. We believe this is not a problem particular to ARF; most of current rate adaptation schemes may not perform well in the bus stop environment.


Figure 4. RSSI of signals from bus APs. We can observe that 28 buses with AP arrived at the bus stop during the 30 minutes-long observation period. Typically, a bus stops for less than 30 seconds. Occasionally, when stuck by a traffic stop, a bus stays at more than 60 seconds (Observed at time points of 40 second, 700 second, 1300 second and 1600 second).

Figure 5 (c) and (d) show the throughput and congestion window size of a TCP connection. To measure TCP performance, we let the smartphone to download a movie file from our server. The average RTT of the connection is 564 msec. The first connection during (55, 80) achieves the most TCP throughput while the last connection contributes virtually none. Let us examine the TCP throughput in an interval (50, 100). Most TCP throughput is achieved before the 80 second mark. During this period, congestion window size increases intermittently but steadily from the minimum to 250 KB. Even though the link layer tries to transmit frames after the 80 second mark, most of them fail with virtually no TCP throughput. Note the stable uplink transmission rate in the (80, 100) interval does not guarantee good TCP performance. We scrutinized the packets captured by the Wireshark during the observation period. The associated Wi-Fi connection is disconnected at 80 second due to the bad channel quality and TCP finishes its session via RST after the timeout at 224 second.

Two short connections at 150 second and 180 second are the examples that manifest the importance of connection length on TCP performance. Due to the short connection times, congestion window size, even though it increases up to 50KB from the minimum, it fails to reach the maximum and bandwidth available for short durations is wasted. We also investigated the connection at the 200 second mark. Over this connection, TCP tries to continue the previous session with a new AP via TCP keep-alive, but the new connection is too unstable to re-initiate the TCP connection.


(a) TX Rate (Uplink)

(b) TX Rate (Downlink)

(c) TCP Throughput

(d) Congestion Window Size

Figure 5. One snapshot of wireless traces

4.3 Comparison of TCP and UDP Performance

In this section, we compare the throughputs of UDP and TCP. Since the IP addresses of the mobile devices and laptops change as their Wi-Fi connections are switched, it is very hard to measure the UDP throughput on a downlink; the server cannot comprehend the client's address. For this reason, we measured the uplink throughput of UDP and compare it with that of TCP. Unlike the previous set of measurements, we used Iperf to generate uplink traffic for both UDP and TCP. In particular, we mimicked CBR-like traffic that transmits 1470 byte long packets to the server every 0.5s and measured throughput at the server every one minute.

Figure 6 shows the throughputs of UDP and TCP. We repeated the same experiment of uploading for 100 times. The throughput of UDP, as expected, is greater than that of TCP; 1.7Mbps for UDP, 1.0Mbps for TCP on average. However, both TCP and UDP fail to utilize the Wi-Fi bandwidth fully. The maximum throughputs are only about 7Mbps and 5Mbps for UDP and TCP, respectively.
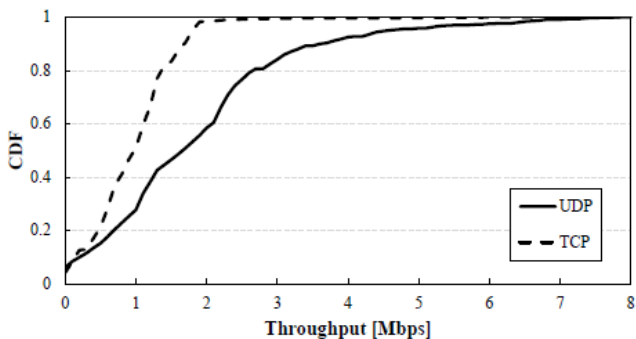

Figure 6. Throughputs of UDP and TCP
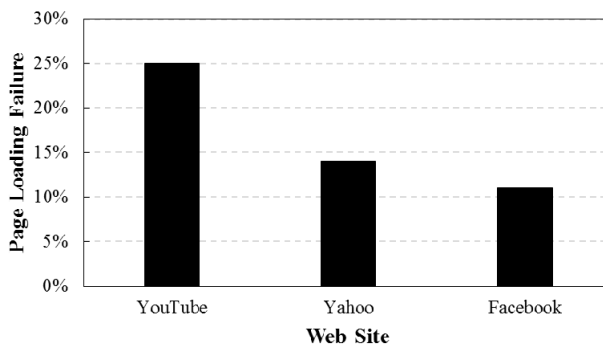
4.4 Application Layer Performance

User perceived performance called QoE(Quality of Experience) may be the more important performance measure than network level performance. Important application level quality metrics include PLT (Page Load Time), page loading failure rate and latencies.

We measured the time taken to load the YouTube home page (https://www.youtube.com/). The PLT is defined to be the time difference from when the request is submitted to finish time. Figure 7 (a) compares the PLTs of two connections; one over the LTE network and another over the Wi-Fi network. Over the LTE connection, web pages are loaded in less than 4 seconds in all trials and the variation is very small. However, over the Wi-Fi network, 4 seconds of PLT is achieved only by 8% of trials. Because some trials require unusually long time, we define a trial as a failure if the web page loading does not finish in 75 seconds. The failure rate is 25%, and the median time of a successful page load is about 15 seconds.

Figure 7 (b) illustrates the page loading failure rates when Wi-Fi connections are used to download homepages of YouTube, Yahoo and Facebook, respectively. Again, we define a trial as a failure if it waits more than 75 seconds to load the homepage. The failure rates over the LTE networks are all zero. The YouTube homepage whose size is largest among the three experienced the largest failure rate.

(a) PLT

(b) Page Loading failure Rate

Figure 7. Application layer performance (a) PLTs of the YouTube.com homepage over LTE and Wi-Fi connections, (b) Page loading failure rate

4.5 MPTCP Performance

MPTCP has been recently introduced to expedite TCP throughput using multiple connections. Ubiquitous availability of both the cellular network and Wi-Fi network triggered the adoption of MPTCP in recent smartphones. Prior performance studies of MPTCP over LTE and Wi-Fi networks confirmed that MPTCP improved TCP throughput, especially for large files. The prior work dealt with stable Wi-Fi networks. We posit that poor quality Wi-Fi, the norm at Seoul bus stops, may assert negative effects on MPTCP unless intelligent load balance mechanisms are accompanied.

Figure 8 shows that the apprehension might indeed be true. We compared the times to finish downloading a file of size 37.8MB in three different scenarios; a) a generic single path TCP over the LTE network, b) MPTCP over both LTE and Wi-Fi APs on buses, and c) MPTCP over LTE and stable Wi-Fi networks. We included the third case to highlight the effect of poor Wi-Fi connection quality on the performance of MPTCP. The download time of MPTCP over stable APs is shortest, re-confirming the prior work that MPTCP is effective for long files over stable networks. However, the advantage of MPTCP over the generic TCP is not as great as the previous studies. Note that Korean cellular network operator have invested heavily on LTE and LTE-A network constructions and their bandwidth is two or three times larger than that of stable Wi-Fi networks. This large performance discrepancy might be the reason why the advantage of MPTCP over generic TCP is reduced.
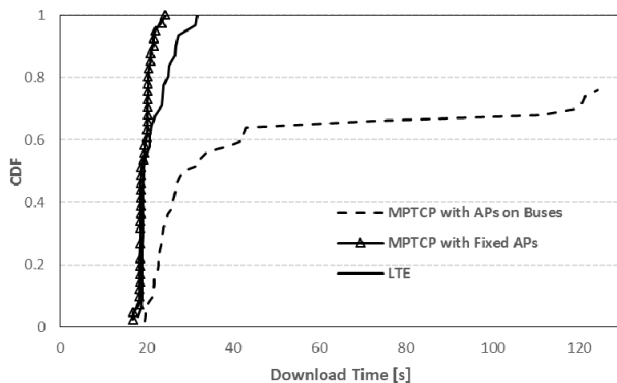

Figure 8. The comparison of time to take download 37.8 MB file through LTE, MPTCP with Aps on buses, and MPTCP with fixed APs

Poor Wi-Fi connections exert a significant negative effect on MPTCP. We define a trial as a failure if it takes more than 130 seconds to finish. The failure rate is 24%; this is particularly unacceptable considering that the failure rate of generic TCP is zero. Also, the mean download time of successful trials is around two times larger than that of generic TCP.

5. CONCLUSION and FUTURE WORK

While APs installed on buses provide convenient Internet accesses to the passengers on board, they can be a nuisance to riders and pedestrians at bus stops. The impact of this side effect has not been measured before and as far as we know this study is the first attempt for thorough investigation of the problem.

Our measurement study carried out at bus stops in Seoul showed that the side effect of bus Wi-Fi APs is quite critical. APs on buses not only fail to provide convenience but also deteriorate the performance of waiting riders significantly. The performance degradation seems to be amplified at upper layers such as TCP and application layers.

The results may indicate the necessity to re-address the prior principle of mobile data offloading; use Wi-Fi as the default network. Instead of connecting to Wi-Fi networks blindly, a careful survey of the quality of the Wi-Fi connection should be performed. Also, to guarantee the proper performance of MPTCP, intelligent and dynamic load balancing to sub-flows might be required.

6. ACKNOWLEDGEMETNS

7. REFERENCES

[1] WiGLE: Wireless Network Mapping. Available: https://wigle.net/.

[2] Google Wi-Fi: Available: http://wifi.google.com/

[3] ABC News. New York City Pay Phone Booths Noew Free WiFi Hotspots, 2012. http://abcnews.go.com/Technology/yorkcity-pay-phone-booths-now-freewifi/story?id=16756016.

[4] Cisco Annual Reports. Available: http://www.cisco.com/

[5] Money Today Korea. Available: http://news.mt.co.kr/.

[6] Tim Carmody, Wi-Not? South Korea's Seoul To Blanket The City With Free Wi-Fi, Fastcompany. Available: http://www.fastcompany.com/1760834/wi-not-south-koreas-seoul-blanket-city-free-wi-fi.

[7] A Future Fueled by Phablets – Worldwide Phablet Shipments to Surpass Portable PCs in 2014 and Tablets by 2015, According to IDC, Sep 2014, http://www.idc.com/getdoc.jsp?containerId=prUS25077914

[8] Mobility between 3GPP Wireless Local Area Network (WLAN) interworking (I-WLAN) and 3GPP systems, 3GPP TS 24.327.

Available: http://www.3gpp.org/dynareport/24327.htm

[9] IP flow mobility and seamless Wireless Local Area Network (WLAN) offload, 3GPP TS 23.261. Available: http://www.3gpp.org/dynareport/23261.htm

[10] Samsung Galaxy S6. Available: http://www.samsung.com/us/explore/galaxy-s-6-features-and-specs/.

[11] M. Afanasyev, T. Chen, G. M. Voelker, and A. C. Snoeren. Usage patterns in an urban wifi network. *IEEE/ACM Trans. Netw.*, 18(5):1359–1372, Oct. 2010.

[12] A. Balasubramanian, R. Mahajan, and A. Venkataramani. Augmenting mobile 3g using wifi. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services*, MobiSys '10, pages 209–222, New York, NY, USA, 2010. ACM.

[13] 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.

[14] Y.-C. Chen, Y.-s. Lim, R. J. Gibbens, E. M. Nahum, R. Khalili, and D. Towsley. A measurement-based study of multipath tcp performance over wireless networks. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, IMC '13, pages 455-468, New York, NY, USA, 2013. ACM.

[15] N. Cheng, N. Lu, N. Zhang, X. S. Shen, and J. W. Mark. Vehicular wifi offloading: Challenges and solutions. Vehicular Communications, 1(1):13-21, 2014.

[16] S. Deng, R. Netravali, A. Sivaraman, and H. Balakrishnan. Wifi, lte, or both?: Measuring multi-homed wireless internet performance. In Proceedings of the 2014 Conference on Internet Measurement Conference, IMC '14, pages 181-194, New York, NY, USA, 2014. ACM.

[17] P. Deshpande, X. Hou, and S. R. Das. Performance comparison of 3g and metro-scale wifi for vehicular network access. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, IMC '10, pages 301–307, New York, NY, USA, 2010. ACM.

[18] A. Ford, C. Raiciu, M. Handley, S. Barre, J. Iyengar, et al. Architectural guidelines for multipath tcp development. IETF, Informational RFC, 6182:2070{1721, 2011.

[19] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket switched networks and human mobility in conference environments. In *Proceedings of the 2005 ACM SIGCOMM Workshop on Delay-tolerant Networking*, WDTN '05, pages 244–251, New York, NY, USA, 2005. ACM.

[20] 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.

[21] S. Liu and A. Striegel. Casting doubts on the viability of wifi offloading. In *Proceedings of the 2012 ACM SIGCOMM Workshop on Cellular Networks: Operations, Challenges, and Future Design*, CellNet '12, pages 25–30, New York, NY, USA, 2012. ACM.

[22] R. Mahindra, H. Viswanathan, K. Sundaresan, M. Y. Arslan, and S. Rangarajan. A practical traffic management system for integrated lte-wifi networks. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, MobiCom '14, pages 189–200, New York, NY, USA, 2014. ACM.

[23] Y. Moon, D. Kim, Y. Go, Y. Kim, Y. Yi, S. Chong, and K. Park. Practicalizing delay-tolerant mobile apps with cedos. In MobiSys. ACM, 2015

[24] J. Ott and D. Kutscher. Drive-thru internet: IEEE 802.11b for "automobile" users. In *INFOCOM 2004. Twenty-third AnnualJoint Conference of the IEEE Computer and Communications Societies*, volume 1, pages –373, March 2004.

[25] A. Sathiaseelan and J. Crowcroft. Internet on the move: Challenges and solutions. *SIGCOMM Comput. Commun. Rev.*, 43(1):51–55, Jan. 2012.

[26] V. A. Siris and D. Kalyvas. Enhancing mobile data offloading with mobility prediction and prefetching. In *Proceedings of the Seventh ACM International Workshop on Mobility in the Evolving Internet Architecture*, MobiArch '12, pages 17–22, New York, NY, USA, 2012. ACM.