

Comparison of User Traffic Characteristics on Mobile-Access versus Fixed-Access Networks

Mikko V.J. Heikkinen^{1,2,3} and Arthur W. Berger^{3,4}

¹ Aalto University School of Electrical Engineering
mikko.heikkinen@aalto.fi

² Helsinki Institute for Information Technology HIIT

³ Massachusetts Institute of Technology

⁴ Akamai Technologies
awberger@mit.edu

Abstract. We compare Web traffic characteristics of mobile- versus fixed-access end-hosts, where herein the term “mobile” refers to access via cell towers, using for example the 3G/UMTS standard, and the term “fixed” includes Wi-Fi access. It is well-known that connection speeds are in general slower over mobile-access networks, and also that often there is higher packet loss. We were curious whether this leads mobile-access users to have smaller connections. We examined the bytes-per-connection and packet loss based on packet retransmissions from a sampling of logs from servers of Akamai Technologies. We obtained 149 million connections, across 51 countries. The mean bytes-per-connection was typically larger for fixed-access: for two-thirds of the countries, it was at least one-third larger. Regarding distributions, we found that the difference between the bytes-per-connection for mobile- versus fixed-access was statistically significant for each of the countries, and likewise for packet loss. However, the difference is typically small. For some countries, mobile-access had the larger connections. As expected, mobile-access often had higher packet loss than fixed-access, but the reverse pertained for some countries. Typically packet loss increased during the busy period of the day, when mobile-access had a larger increase.

1 Introduction

Mobile broadband has become a significant factor in the Internet communications market, and it continues to grow: Cisco [5] forecast that global mobile IP data traffic will double every year through 2014. Informa [7] estimated that there would be globally 670 million mobile broadband subscribers in 2011.

We are interested in comparing Web traffic characteristics of mobile- versus fixed-access end-hosts, where herein the term “mobile” refers to access via cell towers, using for example the 3G/UMTS standard, and the term “fixed” includes Wi-Fi access. Whereas prior work has compared the applications used by mobile- and fixed-access devices [10], here we are interested in the network level, and comparing the size of connections (i.e., number of bytes per connection) for mobile versus fixed devices that are accessing the Web.¹

¹ An extended version of this paper is available as a technical report [6].

It is well-known that connection speeds are in general slower over mobile-access networks [2]. Also, and as reported herein, often, though not always, there is higher packet loss with mobile-access, see §4.2. We are curious whether this leads users to have smaller connections, or would they persevere, so to speak, through the more adverse network conditions.

From a sampling of logs from July 2010 from servers of Akamai Technologies, we examine the number of bytes-per-connection, and packet loss based on packet retransmissions. The data for this study were collected prior to the deployment of 4G/LTE, and thus form a baseline for comparison for when 4G/LTE is broadly in use.

The contributions of this study are:

1. The first reported comparison of mobile- versus fixed-access connection-size and packet-loss
2. The comparison of the daily variation in bytes-per-connection and packet loss, for mobile- and fixed-access
3. Results spanning 51 countries

This paper is structured as follows: §2 reviews the related work. §3 describes our data set. §4 contains our results. §5 summarizes and discusses our results.

2 Related Work

Regarding side-by-side comparison of mobile- and fixed-access traffic, Akamai’s [2] quarterly “State of Internet” reports connection speeds for both fixed- and mobile-access. Also, Sandvine [10] reported traffic profiles for both fixed- and mobile-access. Complementing these two studies, the present work also makes such a comparison, though in contrast to the prior work, we examine different attributes: the number of bytes per connection and packet loss. Using a data set from the Akamai content distribution servers, our study is global in scope and presents results for over 50 different countries.

In addition to the two previous side-by-side fixed- and mobile-access traffic comparisons at the network level, at least two studies have compared fixed- and mobile-access traffic at the application level: Hossfeld et al. [8] compared the performance of a peer-to-peer file sharing application in both fixed- and mobile-access networks, whereas Svoboda [11] compared the session lengths of online gamers in fixed- and mobile-access networks. Furthermore, Kalden & Ekström [9] compared (non-side-by-side) the results from their analysis of GPRS mobile-access traffic to studies of fixed-access traffic by other researchers.

3 Data Set and Methodology

We used data from log files of Akamai Technologies that contained information that enabled a comparison of mobile- and fixed-access traffic on a per country basis. The data consisted of a global sub-sampling of TCP connections between

clients and Akamai servers. These connections are for data that originated at Web sites of medium to large organizations and where the aggregate set of clients is globally distributed. The log lines included the Unix time, the source address of incoming packets seen by the Akamai server, the number of bytes per connection, and the number of retransmitted packets, given Selective Acknowledgment (SACK), which we use as an estimate of packet loss. For the present analysis we used logs from the week of July 25 through 31, 2010.

Our methodology for associating a log line with mobile- or fixed-access is as follows. As the log lines do not contain this association, we started with the source address, which could be the address of the end-device itself, or the public facing address of a mobile access point, etc. We associated this address with an Autonomous System (AS) based on Border Gateway Protocol (BGP) feeds collected by Akamai. We then used the fact that some network operators that offer mobile service dedicate an AS for that purpose, and likewise for fixed-access services; though a complicating factor is that this is not a universal practice, and some ASs contain both mobile and fixed access. For its quarterly “State of Internet” reports [2], Akamai had already identified ASs dedicated for mobile-access, as well as for fixed-access. Some of these ASs had been identified based on prior knowledge, some were identified by their name, and some by contacting the network operator. Additional mobile ASs were found by an initial discriminator of the ASs having a relatively low average connection speed, and then from this pool of ASs, further inquiries were made to confirm whether they were mobile ASs. Given countries in which mobile ASs had been identified, a sampling of fixed-access ASs were also selected to provide a comparison. ASs that could not be identified with high confidence as being in one of these categories were excluded.

We then selected TCP connections in the log files where the client IP address was in one of the selected mobile or fixed ASs. We used the Akamai geo-location service EdgeScope [1] to identify the country in which the client IP address was located. (EdgeScope provides a service-level agreement (SLA) that the association of address to country has 99% accuracy.) We obtained 149 million connections, across 51 countries, where we excluded countries for which the dataset contained less than 1,000 mobile or fixed connections (and some countries with least data to save space, see [6] for results on 57 countries). The median number of mobile-access connections per-country was 48,000, and for fixed-access it was 650,000. As we were interested in comparing the mobile and fixed daily demand, we again used EdgeScope to obtain the latitude and longitude of the client IP address, from which we obtained the local time-zone relative to GMT. This enabled daily demand plots where hour “0” corresponds to midnight for the given client.

4 Results

4.1 Number of Bytes per Connection

Summary Statistics. Table 1 shows the 3rd quartile and mean of the number of KiloBytes-per-connection, partitioned by country and by fixed-access versus

mobile-access. (See [6] for an extended version of the table.) The rows are arranged in increasing order of the mean number of KBytes-per-connection for fixed-access. Note that the mean is larger than the 3rd quartile as the distribution of KBytes-per-connection tends to have a small percentage of large connections. There is clear variation across countries in KBytes-per-connection: The 3rd quartile (75% quantile) varies from 6 to 21 for fixed-access, and 7 to 27 for mobile-access; and the mean varies from 38 to 152 for fixed-access, and 19 to 178 for mobile-access. The 3rd quartile of all countries is 17 for fixed-access, and 14 for mobile-access; the mean of all countries is 113 for fixed-access and 47 for mobile-access.

There are also clear differences between fixed-access and mobile-access. For 65% of the countries, the mean bytes-per-connection was at least one-third larger for fixed-access. For most of the countries ($\sim 75\%$ based on the 3rd quartile), the fixed-access connections have more KBytes than mobile-access; and thus there is a minority where the reverse pertains. As an example, in the USA, the 3rd quartile of KBytes-per-connection for fixed-access is larger than mobile-access, being respectively 16 and 12; while in the South Korea the order is reversed, having values 6 and 27, respectively. In the USA, the mean is again greater for fixed-access, 152 versus 44 KBytes-per-connection, while in France the corresponding values are 102 and 178.

Distributions. We also examined the cumulative distribution function (CDF) of KBytes-per-connection for fixed-access and mobile-access. As one might expect for such a large data set as ours, using the non-parametric Kolmogorov-Smirnov test, we found that the null hypothesis that the two sample distributions (fixed- and mobile-access) come from the same population distribution is rejected with high confidence, for all of the countries, typically with p-values much less than 0.01. Although the two sample distributions are statistically distinct, for many of the countries the visual difference in a plot is rather slight. Though, for a minority of countries, the difference is dramatic.

Figure 1 is a sampling of four countries. For both country pairs, we show two plots: a CDF with a linear scale on the axes, and a complementary cumulative distribution function (CCDF) with logarithmic scales. The former is useful for seeing the bulk 90% of the connections, and the latter for the minority of large-size connections, which impact the means reported in Table 1.

For the USA and France, the distributions given fixed-access versus mobile-access are rather similar, at least for the bulk of the connections. The medians are essentially the same. By the 3rd quartile the difference is more noticeable. An interesting contrast between the USA and France is that in the former the distribution given fixed-access connections has a heavier tail, whereas in the latter the mobile-access connections do, up to 10 MBytes. For example, in the USA, 0.7% of the fixed-access connections are at least 1 MBytes, which is greater than the 0.4% of mobile-access connections. In France, again 0.7% of the fixed-access connections are at least 1 MByte, whereas 2.9% of the mobile-access connections are. For connections of 100 MBytes or more, fixed-access dominates (though of course the percentage of connections is quite small).

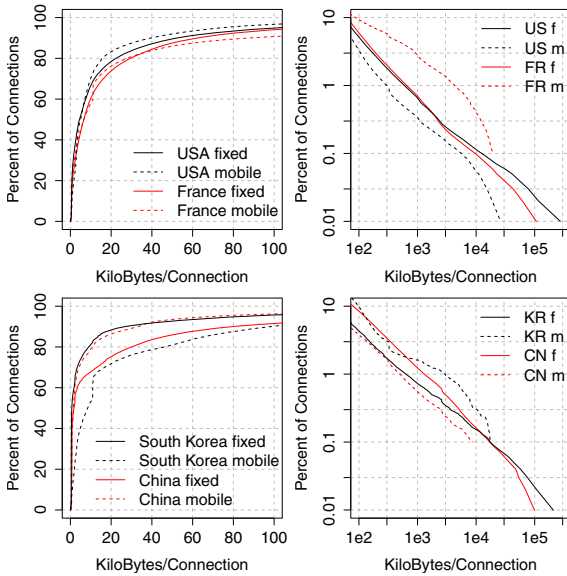


Fig. 1. Distribution for size of connections (left) and complementary distribution for size of connections (right) for the USA and France (top) and South Korea and China (bottom)

In contrast to the USA and France, China and South Korea are two of the minority of countries in Table 1 where the difference in distributions is visually quite evident even for the non-tail portion. Also, as a contrast between South Korea and China, in the former the mobile-access connections are larger (in the sense of the CDF), and in the latter the reverse pertains, even up to connections of 10 MBytes. For example, in South Korea, 22% of the mobile-access connections are greater than 40 KBytes, while fewer (8%) of fixed-access are. In China, 8% of the mobile-access connections are greater than 40 KBytes, whereas more (15%) of fixed-access are.

4.2 Packet Loss Based on Packet Retransmissions

Packet loss on the connection is one of the performance measures of the quality of service provided by the network operator. We were curious how packet loss compared on fixed-access versus mobile-access. Note that this comparison of mobile- and fixed-access packet loss is for clients accessing Web content from Akamai servers. The Akamai server is typically in the same AS as the client, in which case the loss, when it occurs, is within client's AS. And if the server is not in the same AS, then in all likelihood it is in a nearby upstream AS. Thus, in general, the number of packet retransmissions tends to be less as compared to the client accessing Web content directly from an origin site, as then the path is longer, with greater opportunity for experiencing congestion. For each




Fig. 2. Mean connection size versus packet loss, per country, for the subset of countries whose mean packet loss is no more than 3%

connection, we compute the percent of packets sent from the server to the client that are duplicate packets, given SACK TCP, which we use as an estimate of lost packets.

Table 2 reports mean per-connection percent packet loss, partitioned by country and by fixed-access versus mobile-access connections. (See [6] for an extended version of the table.) Typically, there is no packet loss: The median per-connection packet loss was 0.0% in all countries; and the 3rd quartile was also 0.0% in many, but not all cases (see [6]). Note that the mean per-connection packet loss gives equal weight to large and small connections. As a comparison, Table 2 also reports the overall, or access-network, packet loss, defined as the total number of duplicate packets, summed across the set of connections, divided by the total number of packets sent. The rows are ordered in increasing value of fixed-access minus mobile-access mean per-connection packet loss. Countries where this difference is positive, the last 13 rows of the table, had higher packet loss on the fixed-access connections. Note that the mean per-connection packet loss is more often greater than the overall access-network packet loss, which indicates that smaller connections tend to have higher loss. The mean per-connection packet loss of all countries is 2.2% for fixed and 2.5% for mobile, whereas the mean access-network packet loss of all countries is 1.1% for fixed and 1.8% for mobile.

4.3 Dependence of Connection Size on Packet Loss

We wanted to examine the heuristic notion that higher packet loss leads to smaller connection sizes. From Table 2 and Table 1, we have the average network packet loss and mean connection size, respectively, per country. Figure 2 displays a scatter plot of these values and a fitted linear regression line. Since the few countries with the upper-end packet loss can be considered atypical, Figure 2 shows the subset of countries whose mean packet loss is no more than 3%. The plot is truly scattered; though, by eye one can sense a downward trend, i.e. smaller mean connection size with higher packet loss, for both fixed- and mobile-access. The regression lines are included not because a linear model is a good fit,




Fig. 3. Daily distribution for size of connections and packet loss, all countries

but to indicate the downward trend. The statistical test on the regression lines having zero slope is rejected with high confidence ($p < 0.002$).

Although the data support the notion that higher packet loss leads to smaller connection size, one’s viewpoint of the strength of the trend (-20 KBytes-percentage-packet-loss for mobile in Figure 2) is obviously colored by one’s prior expectation. Regardless, the plot also clearly shows the great variability. Of particular note are the two countries whose fixed access has high packet loss of at least 4% and relatively high mean connection size of at least 125 KBytes: China and Russia (see Tables 1 and 2).

4.4 Daily Traffic Pattern

We compared the daily demand for both fixed- and mobile-access. Figure 3 shows the fraction of bytes sent in each of the 24 hours of a day, on fixed-access and on mobile-access networks across all countries. Figures for some individual countries are in [6]. We determined the latitude and longitude of the client address, from which we obtained the local time zone relative to GMT, see §3, and thus we could bin the traffic such that the hour “0” corresponds to midnight for the given client. Also shown on the plots is the mean per-connection packet loss in the respective hour. The daily demand pattern for fixed- and mobile-access are very much alike—fixed-access has a slightly higher afternoon peak, and mobile-access has a bit higher proportion the post-midnight early-morning hours. Packet loss on mobile-access is higher than on fixed-access and has a larger increase during the heavy traffic period of the day. As one would expect, the packet loss for both mobile- and fixed-access increases during the busy period of the day. If one considered the regression of the per-hour packet loss on the per-hour fraction of bytes sent, then the hypothesis of zero slope would be rejected with extremely

Table 1. KiloBytes-per-Connection Statistics

Country	3rd Quartile		Mean	
	Fixed	Mobile	Fixed	Mobile
Uruguay	17	16	38	53
Peru	19	19	45	27
Venezuela	18	17	47	39
Sri Lanka	18	16	47	45
El Salvador	18	20	50	44
Pakistan	19	17	52	59
Croatia	18	9	55	19
Malaysia	20	18	55	43
Israel	16	17	57	49
Greece	19	15	59	41
Paraguay	19	17	59	44
Colombia	19	14	59	52
Slovenia	20	9	60	38
Italy	18	17	60	44
Thailand	18	14	62	31
Slovakia	18	17	62	71
Chile	19	17	64	43
Czech Rep.	18	17	67	43
Australia	13	11	68	37
Morocco	20	18	70	50
Hungary	19	16	71	47
South Africa	12	12	72	123
Kuwait	18	20	72	153
Puerto Rico	18	17	74	72
Lithuania	19	18	77	64
Brazil	19	20	78	58
Mexico	18	21	79	36
Spain	16	16	80	49
Singapore	16	15	82	46
New Zealand	16	16	83	40
Hong Kong	18	14	87	65
Portugal	20	19	88	53
Poland	17	16	91	56
Taiwan	13	12	91	77
Ireland	20	17	92	101
Belgium	19	15	95	93
Ukraine	19	14	96	70
France	21	18	102	178
Austria	18	15	104	54
UK	18	16	104	69
Estonia	18	17	107	56
Norway	16	13	111	38
Canada	18	11	113	38
Germany	19	12	122	71
Switzerland	17	11	123	54
Russia	21	16	125	92
Moldova	18	19	129	54
China	19	7	130	64
Netherlands	14	11	135	79
South Korea	6	27	135	131
USA	16	12	152	44

Table 2. Percent Packet Loss Statistics

Country	Mean Packet Loss			
	Per-Connection		Access-Network	
	Fixed	Mobile	Fixed	Mobile
Peru	2.0	5.8	2.0	5.1
Czech Rep.	1.3	4.5	1.0	3.2
Chile	2.0	5.1	1.4	2.7
Moldova	1.1	3.7	0.4	4.7
Colombia	2.3	4.8	1.9	2.2
Poland	1.8	4.2	1.1	1.9
Morocco	3.2	5.4	2.1	3.0
Portugal	1.6	3.7	0.8	3.4
Estonia	1.6	3.7	1.0	2.1
Ukraine	1.6	3.6	1.2	2.7
Hungary	1.8	3.7	1.0	2.0
Brazil	2.3	4.1	1.6	2.4
Croatia	1.4	3.1	1.0	3.8
Norway	1.5	3.2	0.5	3.1
Belgium	1.4	3.1	0.7	0.8
Spain	1.6	3.3	0.9	1.8
Lithuania	1.6	3.3	1.1	2.0
Venezuela	3.6	5.0	2.9	3.7
Greece	1.5	2.9	1.0	2.1
France	1.5	2.9	0.9	0.6
Puerto Rico	3.0	4.3	1.6	4.2
Italy	2.6	3.9	1.5	2.0
Israel	2.1	3.3	1.6	6.8
New Zealand	2.6	3.6	1.5	2.9
Ireland	2.5	3.3	0.9	1.0
Slovakia	2.1	2.9	1.6	1.3
Russia	4.0	4.7	5.2	2.9
South Korea	0.7	1.4	0.4	0.6
Austria	1.4	2.0	0.6	0.9
Canada	1.3	1.9	0.6	1.1
Taiwan	2.1	2.6	1.5	1.0
Hong Kong	1.9	2.4	1.1	1.1
Pakistan	4.8	5.2	4.2	3.5
Germany	1.2	1.6	0.5	4.2
Uruguay	2.9	3.3	2.4	2.3
Australia	2.5	2.9	1.3	1.6
Mexico	2.6	2.8	1.3	2.4
USA	1.9	1.9	0.9	1.4
Sri Lanka	3.5	3.1	3.2	3.3
Switzerland	1.5	1.0	0.6	0.6
Singapore	3.8	3.2	4.8	2.6
Netherlands	1.2	0.6	0.5	0.2
Slovenia	1.7	0.8	1.1	0.8
UK	2.8	1.7	0.9	1.0
Paraguay	5.4	4.0	5.0	3.6
El Salvador	5.1	3.1	5.5	2.6
Kuwait	3.0	0.7	2.3	0.2
South Africa	3.4	1.0	3.5	0.6
Malaysia	6.0	3.5	6.5	4.2
Thailand	3.9	0.3	3.3	0.5
China	7.4	3.8	4.0	3.5

high confidence ($p < 1e-6$) for both mobile- and fixed-access. The correlation of packet loss with the fraction of bytes sent is 0.94 for mobile-access and 0.95 for fixed-access.

5 Conclusions

We examined the number of bytes-per-connection, and packet loss based on packet retransmissions from a sampling of logs from servers of Akamai Technologies. Regarding to the original question of whether the more adverse conditions on mobile-access networks leads to shorter connections, the rough, first-order answer is “yes”; though, a fuller answer is much more nuanced. One caution to keep in mind is that the statistical analysis of the present study does not prove a causal relationship. Tables 1 and 2 do show that on a per-country basis, packet loss is higher and the mean connection size is smaller on mobile-access for most countries; but there are noted exceptions. From the viewpoint of distributions, we found that the difference between the bytes-per-connection for mobile-access versus fixed-access, as well as the packet-loss distributions, was statistically significant, for all countries we analyzed. However, when plotted, the visual difference is typically small. Also, for some countries, the mobile-access had the larger connections.

In a scatter plot of per-country mean connection size versus mean packet loss, there is a statistically significant trend of smaller connections having higher packet loss for both mobile- and fixed-access, though there is great variability. Aggregating across counties, we found that the daily demand variation is about the same for mobile- and fixed-access, and for both, the packet loss does increase during the busy period of the day, though the increase of mobile-access is greater, suggesting greater sensitivity to the increased demand, that is, a greater likelihood of constrained capacity. Though, as reported in [6], some countries have little to no daily variation in packet loss. We also found per-connection packet loss for both fixed- and mobile-access is often greater than the overall access-network packet loss, indicating smaller connections having higher loss.

Similar to our result of negative correlation between connection size and packet loss rate, bit rate had a consistent positive correlation and jitter above certain limit had a negative correlation with the length of Skype VoIP calls [4]. Also, within certain bounds, round-trip time (RTT), queueing delay, and packet loss rate had negative correlations with session lengths of an online multiplayer game [3]. In general, adverse effects on the quality of service (QoS) seem to decrease the intensity of usage of that service.

Self-selection explains partially why mobile-access connections are typically smaller than fixed-access connections. Users may prefer not to stream long video clips, engage in large downloads, or do other high volume transactions over mobile-access connections. Prevalence of high-speed mobile-access connections could lead to more traffic over mobile-access. In other words, a complementary effect would occur, where higher speed of mobile-access would enable more large-volume transactions.

Analyzing the root causes for our results would require investigating in detail the status of network infrastructure together with device and service implementation and usage in each country, which are out of the scope of this study. One could speculate there are significant differences among those factors across countries. The disparity between fixed-access and mobile-access across countries could be used in simulations or modeling of network performance.

Acknowledgments. During this work, MH was a Visiting Student at MIT CSAIL. MH's work was supported by FICNIA and TEKES. We would like to thank Steven Bauer, David Clark, Rubén García, Tuomo Komulainen, William Lehr, Antti Riikonen, Jesse Sowell, Stephen Woodrow, and anonymous reviewers for their assistance and comments. Any opinions expressed, and any errors are solely the responsibility of the authors.

References

1. Akamai: Akamai's EdgescapE Geo-Location Service. Tech. rep. (2010)
2. Akamai: State of the Internet 2Q/2010. Tech. rep. (2010)
3. Chen, K.T., Huang, P., Wang, G.S., Huang, C.Y., Lei, C.L.: On the Sensitivity of Online Game Playing Time to Network QoS. In: IEEE INFOCOM (2006)
4. Chen, K.T., Huang, C.Y., Huang, P., Lei, C.L.: Quantifying Skype User Satisfaction. SIGCOMM Comput. Commun. Rev. 36(4), 399–410 (2006)
5. Cisco: Cisco Visual Networking Index: Forecast and Methodology, 2009-2014. Tech. rep. (2010)
6. Heikkinen, M.V.J., Berger, A.W.: Comparison of User Traffic Characteristics on Mobile-Access versus Fixed-Access Networks. Tech. Rep. MIT-CSAIL-TR-2011-028, MIT CSAIL (2011), <http://hdl.handle.net/1721.1/62579>
7. Hibberd, M.: Mobile Data Traffic Almost Triples Year on Year. Informa (2010)
8. Hossfeld, T., Tutschku, K., Andersen, F.U.: Mapping of File-Sharing onto Mobile Environments: Feasibility and Performance of eDonkey with GPRS. In: IEEE WCNC, pp. 2453–2458 (2005)
9. Kalden, R., Ekström, H.: Searching for Mobile Mice and Elephants in GPRS Networks. SIGMOBILE Mob. Comput. Commun. Rev. 8(4), 37–46 (2004)
10. Sandvine: Fall 2010 Global Internet Phenomena. Tech. rep. (2010)
11. Svoboda, P.: Measurement and Modelling of Internet Traffic over 2.5 and 3G Cellular Core Networks. Ph.D. thesis, Vienna University of Technology (2008)