# Factors Affecting Performance of Web Flows in Cellular Networks

Ermias Andargie Walelgne*, Setälä Kim*‡, Vaibhav Bajpai†, Stefan Neumeier†, Jukka Manner* and Jörg Ott†
*Aalto University, Finland, ‡ Elisa Oy, Finland
(ermias.walelgne | kim.setala | jukka.manner)@aalto.fi
†Technische Universität München, Germany
(bajpaiv | neumeier | ott)@in.tum.de

***Abstract*—Studies show that more than 95% of the traffic generated by smartphones typically consists of short-lived TCP flows towards websites. The content of such websites often is distributed across multiple servers which requires clients to resolve multiple DNS names and establish multiple TCP connections to fetch the webpage in its entirety. Studies have shown that network latency in a mobile network (attributed to DNS lookup and TCP connect times) contributes heavily to poor experience when browsing such websites. However, there is little understanding of the factors that contribute to high DNS lookup and TCP connect times. In this paper, we take this further by measuring the Domain Name System (DNS) lookup time and the TCP connect time to popular websites from ~25K subscribers of a cellular network operator in Finland. Using a month-long dataset (Oct 2016) of these measurements, we show that LTE offers considerably faster DNS lookup time compared to legacy cellular networks (such as HSPA+ and UMTS). We also show that the model of the device and the proximity of the DNS server to the subscribers impacts the DNS lookup time. Furthermore, the TCP connect time is also affected by the radio technology. We observe that LTE offers a significantly low latency profile such that the TCP connect time to popular websites is reduced by ~80% compared to legacy cellular networks. The presence of ISP caches also considerably improves TCP connect times. Using a `ping` test, we also observe that legacy radio technologies (such as HSPA+ and UMTS) suffer from higher packet loss than LTE.**

## I. Introduction

The trend of users using mobile handheld devices to access the Internet shows a steady increase over the last years. If a user is on the move, these devices commonly use the cellular network to access the Internet. Huang *et al.* [1] show that the majority of network traffic (more than 95%), generated by smartphones typically consists of short-lived TCP flows towards websites. The content of such websites is often distributed across multiple servers, which requires mobile users to resolve multiple DNS names and establish multiple TCP connections to fetch the webpage in its entirety. Internet Service Provider (ISP)s such as T-Mobile [2] have shown that mobile users experience poor web-browsing usually due to high DNS lookup and TCP connect times. Similarly content providers such as Google report [3] that high network latency in a mobile network is contributed by multiple factors such as high DNS lookup, TCP connection and HTTP request times. These latency overheads usually incur before any actual data exchange happens. However, there have been few studies [4],

**Map of Finland**

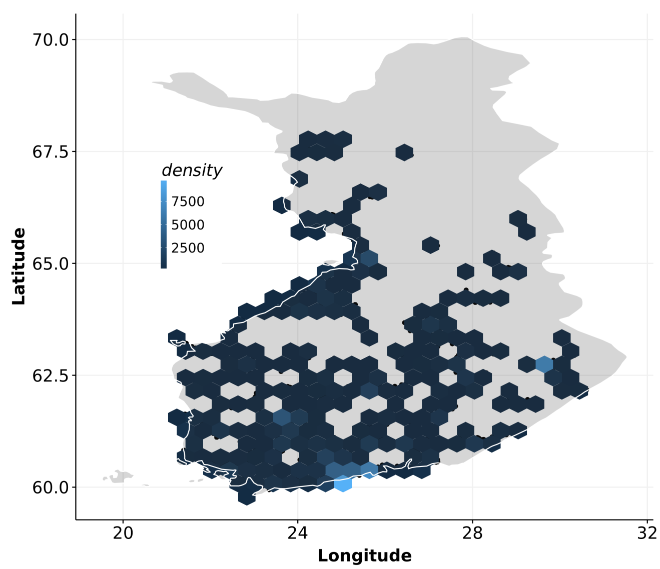


Fig. 1: **The geographical distribution of ~25K subscribers in Finland that participated in this measurement activity.**

[5] that quantify the factors that are responsible for higher DNS lookup and TCP connect times. This is largely because of lack of datasets with rich metadata information (such as the accessed radio technology during the measurement, the device model, *et al.*) that can help to identify those factors for mobile users in a cellular network. Using a month-long (Oct 2016) dataset (see § III) collected by an ISP from subscribers of a cellular network, we profile the performance of ~25K clients distributed across Finland (see Fig. 1) to understand the factors affecting performance in cellular networks. We focus on the performance of short web flows (such as DNS lookup and TCP connect times towards popular websites) that are driven more by latency than by network throughput. We also analyze the packet loss and RTT using more than 2M `ping` measurements towards `www.google.fi`. The performance over the home wireless network is not considered in this work. Towards this end, we provide three main contributions –

– We observe ~2% DNS failures due to `BADVERS` or

ISBN 978-3-903176-08-9 

BADSIG and YXRRSet errors. We show that packet loss can be underestimated in situations where a ping test sends less than 5 ICMP packets (§ IV) per instance. We also show that the legacy (such as HSPA+ and UMTS) technologies suffer from higher packet loss than Long-Term Evolution (LTE).
  - We observe that TCP connect times are affected by radio technology used by the subscriber. LTE offers a significantly low latency profile (§ V) such that TCP connect times to popular websites are reduced by ~80% on LTE compared to legacy networks. Furthermore, we observe that LTE based data subscription plans do not have an impact on TCP connect and DNS lookup times.
  - The device model (§ VI) and the DNS server's proximity to the subscriber has an impact on DNS lookup time. TCP connect times towards popular websites using the same radio technology are comparable, although their DNS lookup times (§ VII) exhibit a difference. ISP caches improve TCP connect times towards www.google.fi and www.youtube.com, while we do not observe cache deployments for www.facebook.com within our dataset. As a result only half of the TCP connections towards www.facebook.com completed within 50 ms, while up to ~80% of TCP connections towards www.google.fi and www.youtube.com completed within 50 ms.

The rest of the paper is structured as follows. We present related work in § II. The methodology for measuring DNS lookup time, TCP connection establishment time, packet loss and the resulting dataset are presented in § III. We then discuss the results in § IV to § VII. We describe the limitations in § VIII and conclude in § IX.

## II. Related Work

We discuss previous work that investigates different factors affecting the performance within cellular networks. For instance, Rodriguez *et al.* [6] identified that multiple DNS lookup operations are one of the major cause for poor network throughput observed in an UMTS network. Xu *et al.* [4] measured DNS lookup time from user's mobile device to evaluate the performance of DNS resolvers and showed that placing the content close to a Gateway GPRS Support Node (GGSN) helps to speed-up content delivery. Jiang *et al.* [7] examined how large buffers in cellular networks contribute to significant TCP queuing delay. Huang *et al.* [1] used data collected from within a carrier's core network to study protocol level interaction of TCP over LTE. They show that TCP underutilizes (by more than 50%) the available bandwidth over LTE. A recent work by Nikravesh *et al.* [8], which is more related to our work, used two mobile apps to measure DNS lookup time to analyze performance variability across carriers and location. Nguyen *et al.* [9] showed how the performance of TCP over an LTE network is affected by handover and sudden load increase on the base station. They use a simulation

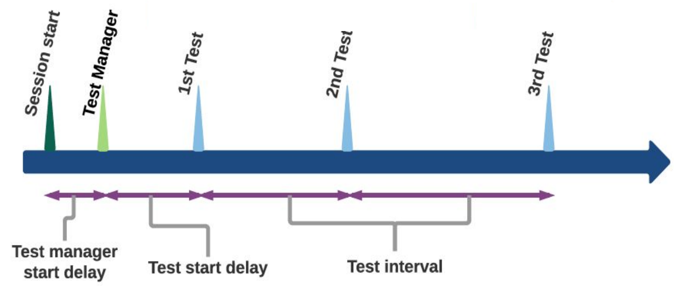


Fig. 2: **Measurement setup session. Immediately after the session starts, the Test Manager module is started. To minimize the possible overload on network or CPU (as many other application might run at the start of the network connection), the Test Manager Start Delay is set to 5 seconds. Test Start Delay is set to few seconds, where at least one measurement is conducted per session. The Test Interval is the interval between subsequent active tests (from ending time of the test until the next test starts). This interval relatively keeps increasing to avoid an excessive data collection.**

environment and show that performance of TCP can degrade as the load on the carrier's cell increases.

Our study covers a broader user-base with ~25K subscribers with ~20M measurements using both, the 3G and LTE network. It combines DNS lookup time, TCP connection establishment time and packet loss metrics that are collected from mobile users' vantage point. In this study, we analyze factors affecting the performance of short web flows in a mobile network which can then be used to identify the potential network performance bottlenecks. Our measurement is conducted towards www.facebook.com, www.google.fi and www.youtube.com. Marquez *et al.* [10] have recently shown that these websites share more than ~30% of traffic in the Orange's mobile network today. We also measure towards www.elisa.net (website of the ISP covered in this study), which covers a substantial fraction of mobile network traffic.

## III. Methodology

The dataset we used for the analysis is collected by a carrier's network measurement platform based in Finland (see § VIII for limitations). It is a mobile application for iOS and Android installed on subscribers mobile devices [11] that conducts measurements tests on DNS lookup, TCP connection establishment and ping.

### A. Measurement Setup

These measurement tests are executed inside a session as shown in Fig. 2. A session starts when a network interface that provides access to the Internet, is available and there is no interface claiming to yield a default route. The session ends in situations when the interface that was providing the default route to the Internet becomes unavailable or the client loses the assigned IP endpoint. Sessions are not periodic, but they are repeated when network conditions change. Such as, when multiple interfaces claim to provide a default route to the Internet, and the 'best interface' [1] changes, the current session is terminated and a new session starts.

### B. Measurement Tests

*1) DNS Lookup Time:* This test measures the time it takes to look up a Fully Qualified Domain Name (FQDN) from a DNS server and resolve it into an IPv4 address (see § VIII for limitations). The test allows one to specify a set of DNS servers and target DNS names. The DNS servers can either be statically configured or automatically assigned by the DHCP server. In our work, we measure the DNS lookup time of four popular websites: `www.google.fi`, `www.youtube.com`, `www.facebook.com` and `www.elisa.net`, as they are commonly known (see § VIII for limitations). Bajpai *et al.* [12] have shown that `www.google.*` websites are served by the same CDN and therefore exhibit similar latency behavior. As such, we use `www.google.fi` for our measurement study.

The test records the resolved DNS name and the IPv4 address of the DNS server. The IPv4 address of the client (majority of which are NATed and consequently receive an IP endpoint from the private [13] address space), the DNS lookup time (in milliseconds), device model type, the radio technology used during the test and the DNS response code indicating the success (or failure) of the test. A timeout of 30 seconds is used in situations where the DNS server is not reachable or the packet is lost. In such a situation, the client does not retry for a failed or timed-out request.

*2) TCP Connect Time:* This test measures the time it takes to establish a TCP connection (over IPv4) to a target website (over port 80) from the client device (see § VIII for limitations). The test starts when the client sends a `SYN` packet to a destination identified by a FQDN. It then subtracts this time value from the time of receiving a `SYN+ACK` packet from the server. This time difference does not include the DNS resolution time.

The test records the starting time of the test, FQDN of the destination host, destination port number, resolved IPv4 address of the destination host, TCP connect time, clients' device model type, the radio technology used during the test, and the success (or failure) of the TCP connection establishment.

*3) RTT and Packet Loss:* This test uses `ping` to measure the RTT and packet loss towards `www.google.fi` (see § VIII for limitations) using ICMP echo request packets. Each `ping` test sends an average of five to nine ICMP echo requests from clients to the target. The payload for each ICMP echo request is configured to be 16 bytes in size.

The test records the DNS name and the resolved IPv4 address of the target, the IPv4 address of the client (majority of which are NATed and consequently receive an IP endpoint from the private [13] address space), total elapsed time of the test, the number of `ping` tests, payload size of the ICMP echo request packet, the minimum, maximum and average RTT, the number of packets sent and received in the test, the response code indicating the success (or failure) of the execution, device model and the radio technology type used during the test.

[1]The interface with the lowest value of the metric attribute

TABLE I: **DNS, TCP and `ping` measurements by website.**

| Website | DNS (#) | TCP (#) | `ping` (#) |
|---|---|---|---|
| `www.facebook.com` | 3,471,440 | 4,572,298 | - |
| `www.google.fi` | 6,981,348 | 4,855,516 | 2,180,700 |
| `www.youtube.com` | 1,628,991 | 4,075,477 | - |
| `www.elisa.net` | 1,821,334 | 5,335,350 | - |

### C. Dataset

The measurements are collected from ~25K subscribers of a cellular network provider based in Finland, geographically distributed as shown in Fig. 1. The dataset consists of ~14M samples of DNS lookup time, ~19M samples of TCP connect time and ~2M samples of `ping` measurements collected in October 2016. Table I provides details of the samples collected towards each target website.

## IV. Failures

### A. DNS Lookup

DNS based redirection techniques are used by content providers (such as Akamai [14]) to determine the location of the end-host and to redirect the contents to the closest content replica [4]. DNS errors may happen for various reasons including poor configuration errors [15], heavy load on the DNS server, and poor network link quality between server and clients. Such errors, if not managed well, could cause drastic damages as it happened in [16], where missing a terminating `'.'` to the DNS records of `.se` zone shutdowns a whole bunch of websites and news outlets in Sweden.

We use the DNS response code to determine the number of successful DNS responses and failures. About 86% of the DNS failures (which is about 2% from the total DNS lookup test) are `BADSIG` or `BADVERS` [17] (Bad OPT Version or TSIG Signature fails), indicating that a responder does not implement the version level of the request [18]. The second most frequent DNS failure code observed is `YXRRSet` [19] which means that the RR Set exists when it should not. Some other DNS failures such as `BADTIME` [20] (out of time windows) and `BADMODE` [21] (Bad TKEY Mode) also rarely happen. One reason for DNS failures to happen is a poorly configured DNS resolver. We noticed that out of all DNS failures that are observed, about 67% of the DNS lookup queries were sent towards the AS790 (Elisa) DNS resolver.

We observe that DNS lookup over LTE experiences about 1.9% of DNS failures, while over UMTS, HSPA and HSPA+ experience 3.4%, 3.9% and 2.7% DNS failures, respectively. Table II shows DNS failures by website. As it can be seen, these failure are almost evenly distributed over the different websites. There is a 2% DNS lookup failure variation between the `www.youtube.com` and `www.google.fi`

TABLE II: DNS Failures per website using the LTE network.

| Website | Failures (%) |
|---|---|
| `www.facebook.com` | 2.16 |
| `www.google.fi` | 0.96 |
| `www.youtube.com` | 2.99 |
| `www.elisa.net` | 2.74 |

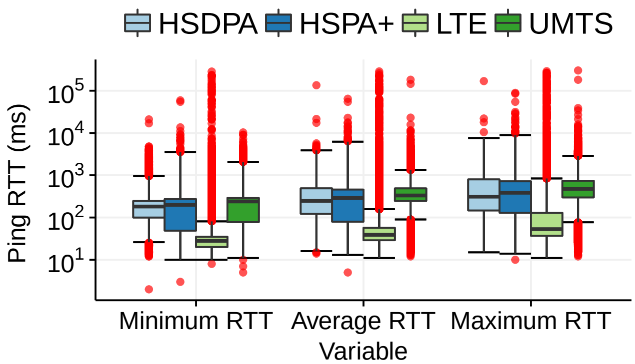


Fig. 3: **Minimum, average and maximum RTT values split by radio technology for a `ping` test towards `www.google.fi`.**

domain names. A variation is also observed on the DNS lookup time (see section V).

### *B. TCP Connect Time*

From the over all TCP connect time measurement dataset a quite small fraction of the dataset (~1.07%) shows a TCP connect time failure. Out of theses errors the majority of them (~0.98%) were raised due to "connection refused" error type. As per RFC-793 [22] this error can happen if the client received a `RST` bit form the server, indicating that the TCP connection must be reset. This could be either due to an intervening firewall that blocks the `SYN` packet or if the destination server is refusing the TCP connection. We observe that out of all measurements conducted using LTE, HSPA+, HSPA and UMTS 1.0%, 1.2%, 2% and 1.5% experience failures, respectively. We observe a similar failure distribution using the same radio technology for each websites.

### *C. Packet Loss Using `Ping` Test Measurement*

Fig. 3 shows that minimum, average and maximum RTT using LTE network are 27, 38, 51 ms, respectively. We observe that 90% of the average `ping` test measurements towards `www.google.fi` using LTE have a RTT time less than 100 ms. While other legacy 3G technologies are quite slow with more than 200 ms RTT.

We observer that ~14.98% of tests in ping measurement (see Table I) have at least one packet loss. Packet loss is calculated from the number of Echo Requests (number of packets sent) and Echo Responses (number of packets received). A packet loss happens if the number of received packets is less than the

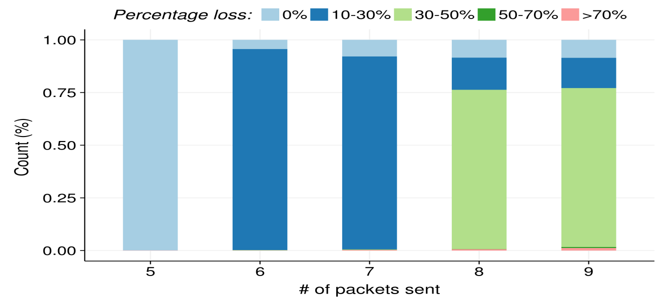


Fig. 4: **Percentage packets loss across the number of packets sent.**

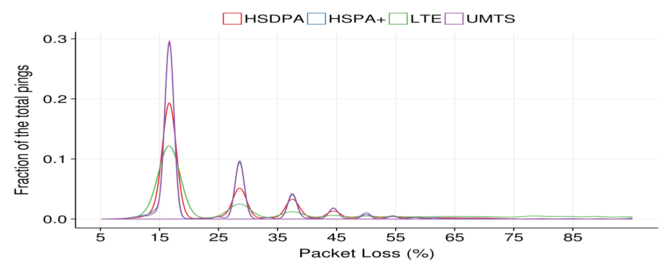


Fig. 5: **Distribution of packet loss as the fraction total `ping` by radio technology type.**

number of packets sent. We observe that most of the packet loss happens, if the number of packets sent at every `ping` test instance is more than five Echo Requests. On the one hand, of about 80% of the `ping` test instances that are measured by sending five individual Echo requests, only 0.33% of the `ping` test instances suffer from packet loss. On the other hand, `ping` test instances that were measured by sending more than five Echo Requests suffer from the highest number of packet loss. For instance, measurements that have 6 or 7 Echo Requests suffer with a higher number of packet loss, covering 57% and 19% of the the total `ping` packet loss, respectively. Fig. 4 shows the percentage of packets loss across the number of packets sent at each instance of the `ping` test.

We noticed that more than 77% of the `ping` test measurements were conducted using the LTE network. Only 2.4% of them lost at least a single packet. The second and the third highest number of `ping` measurements are carried over HSPA+ and UMTS, each sharing 11.18% and 8.28% from total measurement, respectively. More than 50% of the `ping` tests on HSPA+ experience at least one packet loss. The `ping` test using the UMTS network suffers from the highest packet loss, reaching up to 65%.

Fig. 5 shows the distribution of `ping` test fraction that experiences various percentage of packet loss when tested with a different radio technology. For example, with LTE, more than 0.1 fractions of `ping` tests suffer from 15% packet loss. Note that, both Fig. 5 and Fig. 4 show the `ping` test by excluding the cases in which the packet loss was zero percent.

Going forward (sections V - VII-A) we focus on different

TABLE III: DNS & TCP Measurements by radio technology.

| Radio Technology | DNS (%) | TCP (%) |
|---|---|---|
| LTE | 68.94 | 69.59 |
| HSPA+ | 10.59 | 10.23 |
| HSPA | 2.41 | 2.41 |
| UMTS | 14.51 | 14.72 |
| Others | 3.55 | 3.05 |

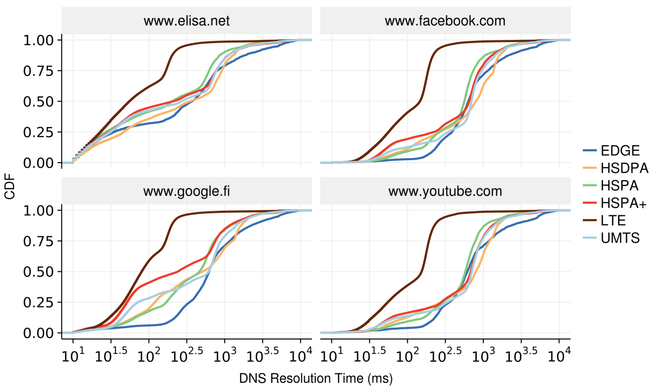


Fig. 6: **DNS response times by radio technology: LTE exhibits significantly lower latency.**

factors impacting DNS lookup and TCP connect times.

## V. Radio Technology

In today's cellular network environment, there is quite a range of radio technologies with different levels of performance. These radio technologies including LTE, HSPA+, HSPA and UMTS have a various range of bandwidth performance. Most of the today's mobile devices are equipped with all of these radio technologies. We analyze how DNS lookup and TCP connect time varies across different radio technologies.

**DNS Lookup Time:** Fig. 6 shows that there is a clear DNS lookup time difference between the radio technologies. There is a fast DNS resolution time in DNS lookup for recent network technologies such as LTE and HSPA+ and a considerably long resolution time for anterior technologies such as HSPA and UMTS.

The LTE network technology consistently shows the best DNS resolution performance on all of the four tested websites. The median difference between LTE and UMTS for resolving `www.google.fi` is 370 ms. Fig. 6 also shows the poor performance of earlier radio technologies such as EDGE, that takes more than half of a second (624 ms) to resolve the IP address of `www.google.fi`. The percentaged difference of DNS response time between LTE and other radio technologies varies among the different domain names. Fig. 6 shows that 50% of the requests send to `www.elisa.net` are resolved in less than 500 ms, irrespective of the radio technology type. Except using the LTE network, only 25% to 30% of the DNS queries send to `www.facebook.com` and `www.youtube.com` are resolved in less than 500 ms. In other words more than 70% of DNS lookup queries send to `www.facebook.com` and `www.youtube.com` took more than 500 ms to get back the resolved IP address.

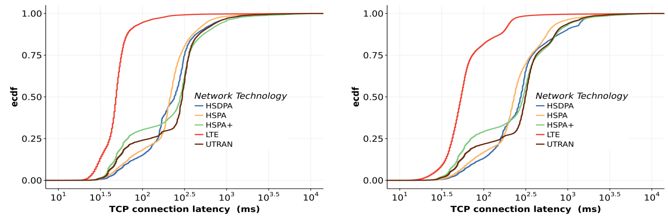


(a) `www.youtube.com` (b) `www.google.fi`

Fig. 7: **TCP connect time towards `www.youtube.com` and `www.google.fi` by radio technology. The distribution exhibits similar pattern for `www.elisa.net` and `www.facebook.com`, too.**

The probability of resolving `www.google.fi` below 100 ms using LTE and HSPA+ radio technology is 65% and 58%, respectively, which is a difference of 7%. Whereas, the probability of resolving `www.youtube.com` below 100 ms using LTE and HSPA+ radio technology is 36% and 23%, respectively, which is a difference of 13%. For most of the 3G and 4G technologies, about 50% of the time, the DNS resolution of `www.google.fi` takes more than 500 ms.

**TCP Connect Time:** We study the performance variation among radio technologies by comparing the latency to reach a give website address through TCP. Fig. 7 shows TCP connect times of `www.youtube.com` and `www.google.fi` using different radio technologies. Similar to the DNS lookup latency, LTE outperforms all other radio technologies. For example, about 92% of the TCP connect time tests using LTE have less than 100 ms latency. Whereas only about 28% of the 3G based TCP connect time tests are below 100 ms. The median TCP connect time of `www.youtube.com` under LTE and Legacy technology is 50 ms and 251 ms, respectively. Thus, LTE reduces the TCP connect latency by 80%. The measurement distribution of TCP connect time and DNS lookup test by radio technology is shown in Table III.

Given that we know the difference between LTE and legacy radio technologies, going forward, we only look at factors affecting performance on the LTE network.

### A. LTE Subscription Plan

We use few randomly selected sample clients' data subscription plan as a reference to study the impact of data subscription plan for DNS lookup and TCP connect time latency. The clients' data plan is classified into 2 packages based on the downlink and uplink speed limits. These are `4G` and `4G-super` for the upper-downlink limit of 25 and 80 Mbits/s, respectively. Note, we only consider LTE in this section.

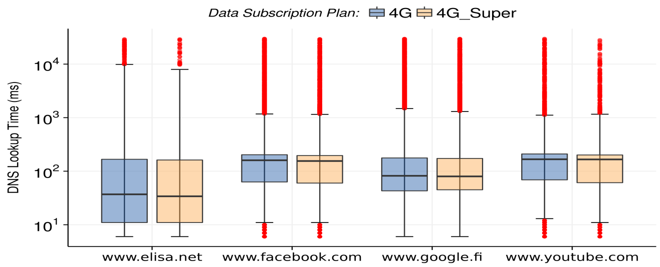


Fig. 8: **DNS response times by subscription plan.**

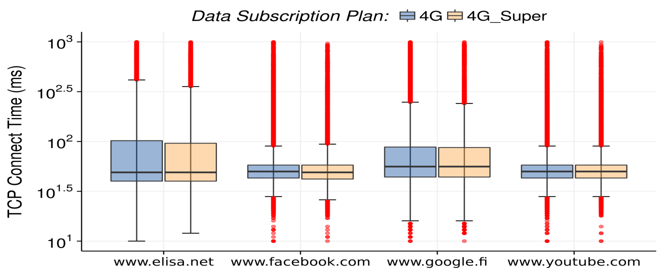


Fig. 9: **TCP connect time by subscription plan.**

**DNS Lookup time:** Fig. 8 depicts the DNS response time per users' data subscription plan for each radio technology. The graph shows that the clients' data subscription plan does not actually contribute to the DNS lookup time performance.

**TCP Connect Time:** As depicted in Fig. 9, the data subscription type has a very small impact on the TCP connection establishment time.

## VI. Device models

We analyzed the impact of different device model types and year of release for both TCP connect time and DNS lookup time performance.

**DNS Lookup Time:** Fig. 10 shows the DNS response time of the top 30 device models, ordered by device models' release year. All the devices are capable of using both, 3G and LTE radio technologies. The selected devices were using the LTE network during the DNS test session. The model names are substituted with the index number to ensure anonymity. The y-axis reflects the DNS lookup time value of at least 10K individual tests for each device that subscribed to a single network operator. We can observe that there is a significant difference in DNS resolution time among device models. For instance, observing the median value of devices released in the year 2015, it appears that the device model *#17* has the highest DNS resolution time, whereas the device model *#22* has a relatively short DNS lookup time for resolving the domain name `www.facebook.com`. The standard deviation (not shown in the plot) of the DNS lookup time across the 30 devices is also highly variable, ranging from 622.45 to 3891.36 ms. The ANOVA [23] F-test for DNS response time is also significant (P-value of 0.0001), asserting that the DNS resolution time is indeed affected by device model type.

To further explore this, we conduct a manual inspection to some of the devices by minimizing the variance such as

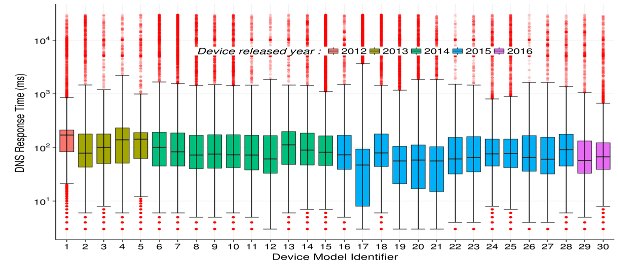


(a) `www.google.fi`

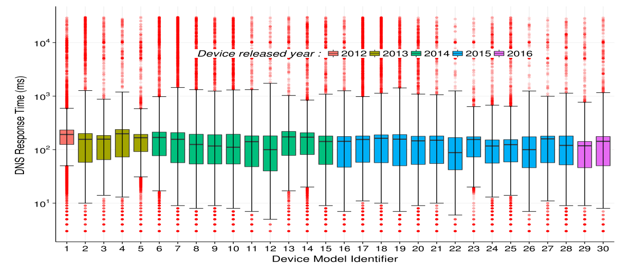


(b) `www.facebook.com`

Fig. 10: **DNS response time of `www.google.fi` (above) and `www.facebook.com` (below) across device models as measured over LTE. Order by device models' release year.**

by fixing the subscribers location and time. From the manual inspection, we observe that few device models consistently show a poor resolution time performance in both LTE and 3G radio technologies. We also observe that devices which have larger internal memory and storage capacity are relatively faster conducting a DNS lookup.

**TCP Connect Time:** The impact of various device model types for TCP connect time latency is very small, especially when it is compared to the DNS lookup time. As shown in Fig. 11, except few devices the median latency among device types when tested towards `www.google.fi` and `www.facebook.com` is less than 100 ms. The device model's release year has also no direct impact on the TCP connection establishment time variation.

## VII. Websites

**DNS Lookup Time:** Fig. 12 shows that the DNS lookup time significantly varies among different websites, using the same radio technology (LTE) that has been accessed during the DNS test. The DNS lookup times of `www.youtube.com` and `www.facebook.com` are significantly slower than the ones of `www.google.fi` and `www.elisa.net`. One cause is that the A entries for `www.google.fi` and `www.elisa.net` (ISP's website) are more likely to be cached by DNS resolvers than `www.youtube.com` and `www.facebook.com`.

**TCP Connect Time:** The TCP connection time is one important measure for websites download time and user satisfaction. Prior work [24] has shown that about 17% of the

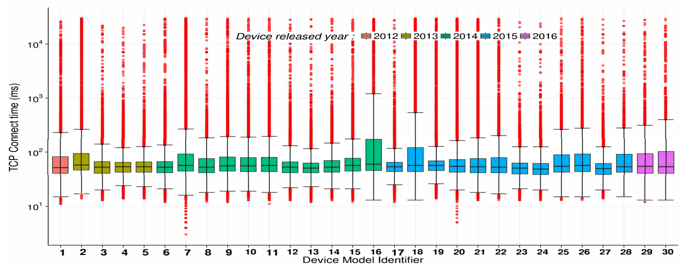


(a) www.google.fi

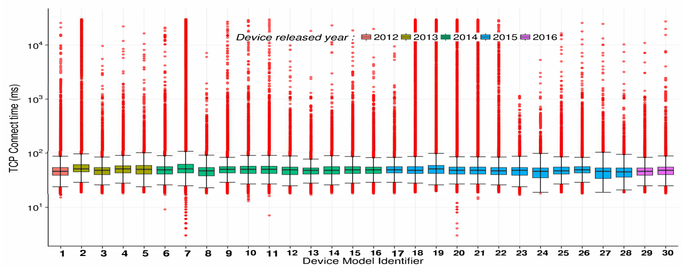


(b) www.facebook.com

Fig. 11: **TCP connect time for www.google.fi (above) and www.facebook.com (below) across device models as measured over LTE.: Order by device models' release year.**

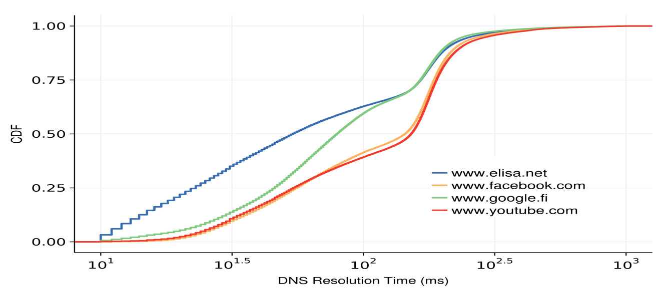


Fig. 12: **DNS response time towards websites using LTE; tested towards different DNS resolvers. Note that the variation almost stays the same if we fix it to a single DNS resolver.**

users are impatient to wait if the response time of a given website is greater than 5 seconds. Thus, we analyze the TCP connectivity time for different websites.

The time elapsed between sending the SYN packet to open the TCP socket and receiving the SYN+ACK response to selected website addresses is shown in Fig. 13. We observe that the majority of TCP connection latencies using LTE range from 20 to 200 ms, irrespective of the website's address. For instance, about 97% of the TCP connections to www.facebook.com are completed in less than 200 ms.

We can see that 90% of the time, www.facebook.com and www.youtube.com can be reached in less than 100 ms from a client's device. Whereas, for www.google.fi and www.elisa.net, only 80% and 76% of the TCP

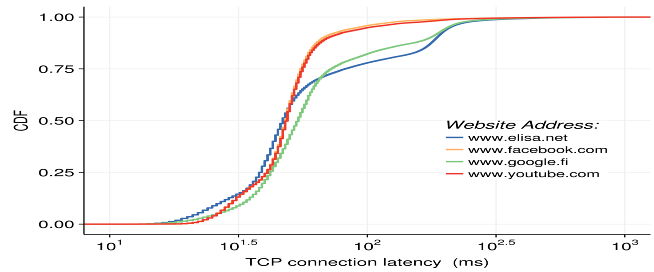


Fig. 13: **TCP connect time towards websites under LTE.**

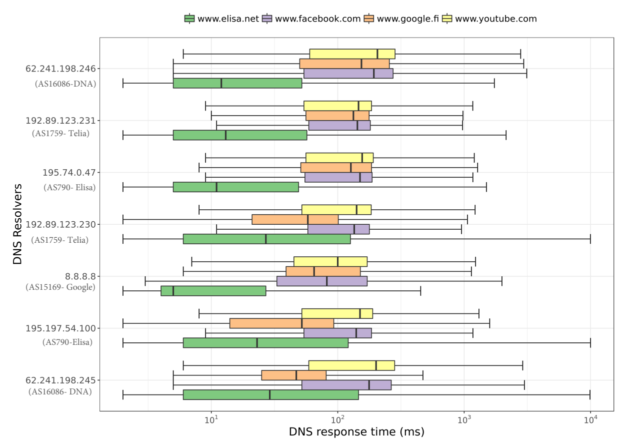


Fig. 14: **DNS response time by resolver IP address using LTE.**

connection test are below 100 ms, respectively.

### A. Destination Autonomous System Number (ASN)

**DNS Lookup Time:** Previous studies show that cellular DNS servers can yield faster DNS lookup time than public DNS resolvers [5]. In light of this, we compare the capability of different DNS servers to resolve a domain name to an IP address. Fig. 14 shows the DNS lookup time of different resolvers per website address. Each of these DNS resolvers IP has more than 10K measurements. We can see that some cellular network DNS servers have a faster DNS lookup time for www.google.fi than Google DNS servers. We also notice that there is a significant variation between DNS resolvers belonging to the same ISP. For instance, two DNS resolvers inside AS790, "EUNET. FI" of two different IP entries 195.74.0.47 & 195.197.54.100, have 133 ms and 51 ms (median) to resolve wwww.google.fi using the LTE network. This variation might happen due to the closer proximity of the DNS resolver to the ISP network [25].

**TCP Connect Time:** For a better network traffic management and performance optimization, network operators may deploy proxy servers between the client and the target destination server [26]. We observe that a proxy or a cache server between the client and the true destination host server may acknowledge the TCP socket request first [27], [28]. This has the advantage of decreasing TCP connection time in the

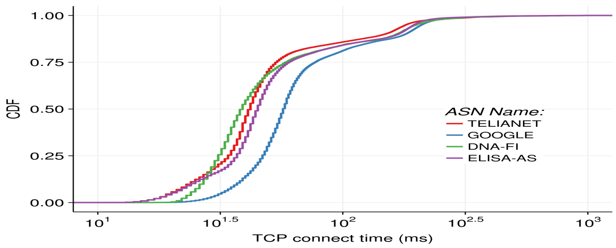


(a) www.google.fi

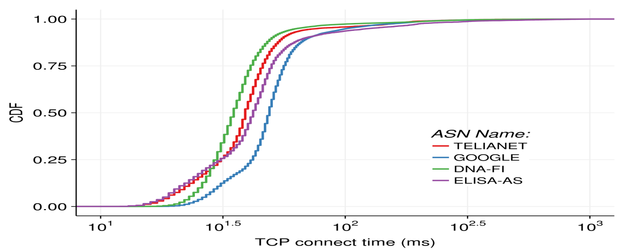


(b) www.youtube.com

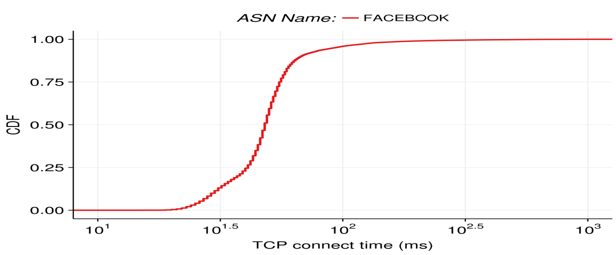


(c) www.facebook.com

Fig. 15: **TCP connect time towards `www.google.fi`, `www.youtube.com` and `www.facebook.com` by destination ASN from LTE networks.**

order of milliseconds (as shown in section VII-A). We use the RIPE [29] service to map the resolved IP address of the websites to an ASN value.

Fig. 15 (a) shows that the latency for a TCP connection to reach the website `www.google.fi` varies based on the ASN number for the same radio technology (LTE). We can see that subscribers served by the ISP network manage to reach the `www.google.fi` website faster than the request sent to Google-owned web-servers. One possible reason of the low TCP connection time for those hosted by ISP would be that web proxies are used to improve browsing performance response [30]. This means, if a TCP connection request is sent to the actual server, the proxy, which is installed between the client and the true destination server, may acknowledge the socket request before passing it to the destination server.

Fig. 15 (b) shows that TCP connect time latency towards `www.youtube.com` varies by the ASN value using the same radio technology (LTE). Clients served by the ISP network have managed to reach the `www.youtube.com` website in short time. This indicates that pushing the content close to the subscriber could potentially reduce the end-to-end latency by more than 20% compared to the requests sent to YouTube-owned web-servers; this is equal to [31], which points that caching improves the fetch time of small files. Fig. 15 (c) shows all TCP connection requests sent to `www.facebook.com` were served by a single ASN. Since `www.facebook.com` does not hit any caches in the ISP network, TCP connect time towards `www.facebook.com` is substantially slower than towards `www.youtube.com` and `www.google.fi`. This is shown in Fig. 15.

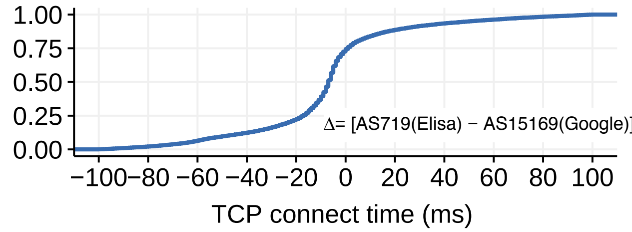


Fig. 16: **TCP connect time towards `www.google.fi` showing the latency difference between ISP cache - Elisa (AS719) and CDN - Google (AS15169) using LTE. Delta is the TCP connect time difference between Elisa and Google when the same user is getting a reply from the two network within a one hour time window.**

$$\Delta t(ct) = isp(ct) - cdn(ct) \quad (1)$$

We use Eq. 1 to calculate the TCP connect time difference between an ISP cache and CDN. For this, we created a pair of CDN and ISP per user within a one-hour time frame. First, we grouped the dataset by the user, ASN and one-hour window. If there is more than one measurement by a given user in a combination, we take the mean value. Then, we keep the ones that have pairs (in this case Google and Elisa). Fig. 16 shows the distribution of difference in TCP connect times between two destinations, where values on the negative scale indicate that ISP cache is faster. We observe that about 70% of TCP connect time towards `www.google.fi` achieve lower latency when they hit ISP cache.

## VIII. Limitations

The measurements only consist of clients based in Finland using IPv4. The only measured services are those that run on port 80. The websites chosen are the most commonly used websites (except `www.elisa.net`) following the Alexa [32] website ranking. The `ping` measurements are conducted only towards `www.google.fi`. As such, it is not known whether and how the observations would differ from a different client base per country and towards different websites or a different services on the Internet. However, these three websites (`www.facebook.com`, `www.google.fi` and `www.youtube.com`) have a high probability of reflecting the majority of the mobile web-user-experience as they generate a considerable size of network traffic in mobile networks [10].

## IX. Conclusion

We presented an analysis on factors that affect DNS lookup time and TCP connect time towards popular websites in cellular networks. We showed that DNS lookup time significantly varies for different websites, even when the same radio technology is accessed during the measurement. We showed that caches closer to the ISP could significantly improve TCP connect time. Also, the proximity of DNS server to the subscriber has a higher impact on DNS lookup time performance. We also observed that LTE offers considerably low latency compared to legacy radio technologies. We show that packet loss can be underestimated in situations where a `ping` test sends less than 5 ICMP packets per instance. Thus, we recommend that a packet loss analysis based on the `ping` test should consider increasing the number of packets per `ping` test instance for better results.

## Acknowledgements

This work is partially funded by the FP7 Marie Curie Initial Training Network (ITN) METRICS project (grant agreement No. 607728).

## References


[1] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, and O. Spatscheck, "An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance," ser. ACM SIGCOMM, 2013, pp. 363–374. [Online]. Available: http://doi.acm.org/10.1145/2486001.2486006

[2] J. Hui and K. Lau, "T-Mobile QoE Lab: Making Mobile Browsing Faster and Open Research Problems," ser. ACM MobiCom, 2013, pp. 239–242. [Online]. Available: http://doi.acm.org/10.1145/2500423.2504585

[3] I. Grigorik. (2013) Performance of Wireless Networks. [Online]. Available: https://hpbn.co

[4] Q. Xu, J. Huang, Z. Wang, F. Qian, A. Gerber, and Z. M. Mao, "Cellular Data Network Infrastructure Characterization and Implication on Mobile Content Placement," ser. ACM SIGMETRICS, 2011, pp. 317–328. [Online]. Available: http://doi.acm.org/10.1145/1993744.1993777

[5] J. P. Rula and F. E. Bustamante, "Behind the Curtain: The Importance of Replica Selection in Next Generation Cellular Networks," ser. SIGCOMM, 2014, pp. 135–136. [Online]. Available: http://doi.acm.org/10.1145/2619239.2631465

[6] P. Rodriguez, S. Mukherjee, and S. Ramgarajan, "Session Level Techniques for Improving Web Browsing Performance on Wireless Links," ser. WWW, 2004, pp. 121–130. [Online]. Available: http://doi.acm.org/10.1145/988672.988690

[7] H. Jiang, Y. Wang, K. Lee, and I. Rhee, "Tackling Bufferbloat in 3G/4G Networks," ser. IMC, 2012, pp. 329–342. [Online]. Available: http://doi.acm.org/10.1145/2398776.2398810

[8] A. Nikravesh, D. R. Choffnes, E. Katz-Bassett, Z. M. Mao, and M. Welsh, "Mobile Network Performance from User Devices: A Longitudinal, Multidimensional Analysis," ser. PAM, 2014, pp. 12–22. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-04918-2_2

[9] B. Nguyen, A. Banerjee, V. Gopalakrishnan, S. Kasera, S. Lee, A. Shaikh, and J. Van der Merwe, "Towards Understanding TCP Performance on LTE/EPC Mobile Networks," ser. AllThingsCellular, 2014. [Online]. Available: http://doi.acm.org/10.1145/2627585.2627594

[10] C. Marquez, M. Gramaglia, M. Fiore, A. Banchs, C. Ziemlicki, and Z. Smoreda, "Not All Apps Are Created Equal: Analysis of Spatiotemporal Heterogeneity in Nationwide Mobile Service Usage," 2017. [Online]. Available: URI:http://eprints.networks.imdea.org/id/eprint/1710

[11] (2017) Mobiili OmaElisa. [Online]. Available: https://verkkoasiointi.elisa.fi/

[12] V. Bajpai and J. Schönwälder, "IPv4 versus IPv6 - who connects faster?" ser. IFIP, 2015, pp. 1–9. [Online]. Available: https://doi.org/10.1109/IFIPNetworking.2015.7145323

[13] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, "Address Allocation for Private Internets," RFC 1918 (Best Current Practice), RFC Editor, Fremont, CA, USA, pp. 1–9, Feb. 1996, updated by RFC 6761. [Online]. Available: https://www.rfc-editor.org/rfc/rfc1918.txt

[14] E. Nygren, R. K. Sitaraman, and J. Sun, "The Akamai Network: A Platform for High-performance Internet Applications," *SIGOPS Oper. Syst. Rev.*, vol. 44, no. 3, pp. 2–19, Aug. 2010. [Online]. Available: http://doi.acm.org/10.1145/1842733.1842736

[15] V. Pappas, Z. Xu, S. Lu, D. Massey, A. Terzis, and L. Zhang, "Impact of Configuration Errors on DNS Robustness," ser. SIGCOMM, 2004. [Online]. Available: http://doi.acm.org/10.1145/1015467.1015503

[16] P. McNamara. (2009, October) Missing dot drops Sweden off the Internet. [Online]. Available: http://www.networkworld.com/article/2251220/lan-wan/missing-dot-drops-sweden-off-the-internet.html

[17] P. Vixie, "Extension Mechanisms for DNS (EDNS0)," RFC 2671 (Proposed Standard), RFC Editor, 1999, obsoleted by RFC 6891. [Online]. Available: https://www.rfc-editor.org/rfc/rfc2671.txt

[18] ——, "Extension Mechanisms for DNS (EDNS0)," pp. 1–7, 1999. [Online]. Available: https://doi.org/10.17487/RFC2671

[19] P. Vixie, S. Thomson, Y. Rekhter, and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)," RFC 2136 (Proposed Standard), pp. 1–26, 1997. [Online]. Available: https://www.rfc-editor.org/rfc/rfc2136.txt

[20] P. Vixie, O. Gudmundsson, D. Eastlake 3rd, and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)," RFC 2845 (Proposed Standard), RFC Editor, 2000, updated by RFCs 3645, 4635, 6895. [Online]. Available: https://www.rfc-editor.org/rfc/rfc2845.txt

[21] D. Eastlake 3rd, "Secret Key Establishment for DNS (TKEY RR)," RFC 2930 (Proposed Standard), RFC Editor, 2000, updated by RFC 6895. [Online]. Available: https://www.rfc-editor.org/rfc/rfc2930.txt

[22] J. Postel, "TRANSMISSION CONTROL PROTOCOL," Internet Requests for Comments, RFC Editor, RFC 793, September 1981. [Online]. Available: https://tools.ietf.org/html/rfc793

[23] J. J. Faraway, "Practical regression and ANOVA using R." 2002. [Online]. Available: https://cran.r-project.org/doc/contrib/Faraway-PRA.pdf

[24] Gomez. (2010) When seconds count. [Online]. Available: htp://ftp.software.ibm.com/software/au/downloads/GomezWebSpeedSurvey.pdf

[25] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig, "Comparing DNS Resolvers in the Wild," ser. ACM IMC, 2010, pp. 15–21. [Online]. Available: http://doi.acm.org/10.1145/1879141.1879144

[26] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson, "Netalyzr: Illuminating the Edge Network," ser. ACM SIGCOMM, 2010, pp. 246–259. [Online]. Available: http://doi.acm.org/10.1145/1879141.1879173

[27] M.-C. Roşu and D. Roşu, "An Evaluation of TCP Splice Benefits in Web Proxy Servers," ser. WWW, 2002, pp. 13–24. [Online]. Available: http://doi.acm.org/10.1145/511446.511449

[28] A. Pathak, A. Wang, C. Huang, A. G. Greenberg, Y. C. Hu, R. Kern, J. Li, and K. W. Ross, "Measuring and Evaluating TCP Splitting for Cloud Services," ser. PAM, 2010, pp. 41–50. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-12334-4_5

[29] RIPE. (2016) RIPE Stat API. [Online]. Available: https://stat.ripe.net/

[30] P. Rodriguez and V. Fridman, "Performance of PEPs in Cellular Wireless Networks," ser. Web Content Caching and Distribution. Springer, 2004, pp. 19–38. [Online]. Available: https://link.springer.com/chapter/10.1007/1-4020-2258-1_2

[31] X. Xu, Y. Jiang, T. Flach, E. Katz-Bassett, D. R. Choffnes, and R. Govindan, "Investigating transparent web proxies in cellular networks," ser. PAM, 2015, pp. 262–276. [Online]. Available: https://doi.org/10.1007/978-3-319-15509-8_20

[32] (2017) Alexa: The top 500 sites on the web. [Online]. Available: https://www.alexa.com/topsites