# Revisiting the State of Cellular Data Connectivity in India

Asheesh Sharma
asheesh.csstaff@cse.iitd.ernet.in

Manveen Kaur
manveen.csstaff@cse.iitd.ernet.in

Zahir Koradia
zahir.koradia@gmail.com

Rahul Nishant
mcs122818@cse.iitd.ac.in

Sameer Pandit
mcs122821@cse.iitd.ac.in

Aravindh Raman
aravindh.csstaff@cse.iitd.ernet.in

Aaditeshwar Seth
aseth@cse.iitd.ernet.in
IIT Delhi
New Delhi, India

## ABSTRACT

The count of mobile Internet users in India has been growing at a rate of 27% annually and is expected to reach 300M by 2017. There is however limited understanding of whether this rapid growth is happening while also ensuring that good quality of service is provided to users. To find out, building upon our earlier work [17] we deployed a measurement framework in 20 rural, semi-urban, and urban sites in North India and probed four leading 2G and 3G telecom providers to measure performance metrics such as availability, throughput and latency. We also observed some design and configuration aspects of cellular networks that affect the quality of service perceived by users. Our results point to many instances where misconfigurations or inadequate provisioning or poor monitoring of cellular networks led to significantly low performance provided to users. We are now using these results to argue for more robust QoS regulation in the country, and show how the current regulations for 2G and 3G services are not sufficient to hold providers accountable for the quality of service provided by them.

## Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Reliability, availability, and serviceability

## General Terms

Measurement, Performance, Reliability

## Keywords

Cellular data networks


*ACM DEV 2015,* Dec 1 – 2, 2015, London, United Kingdom
 ISBN 978-1-4503-3490-7/15/12 ...$15.00.
DOI: http://dx.doi.org/10.1145/2830629.2830649.

## 1. INTRODUCTION

The rapid proliferation of mobile phones seen around the world, and optimistic projections for the growth of mobile Internet in developing regions in the next few years, have indeed become legendary tales in the information technology revolution of the 21st century. In fact, persuading voice users to convert to using data services has spawned an industry in itself with initiatives such as Internet.org and startups such as Apps Daily and Jana, using a variety of techniques to influence users to come online. A recent study by McKinsey [13] projects India as one of the largest markets for mobile Internet, but also highlights challenges of infrastructure, user capability, affordability, and availability of relevant content and services, which if not enabled well can become impediments to the adoption of data services. In this paper, we focus on the first challenge, that of infrastructure quality, and investigate the quality of service currently provided on 2G and 3G networks in India.

While acquiring new users, telecom providers need to expand their infrastructure in tandem to cater to the increased demand. However, consumer complaints are noted commonly in India that providers are not able to meet the quality of service committed by them [7]. The regulatory environment in India is quite proactive, and TRAI (Telecom Regulatory Authority of India) has taken several measures to hold providers accountable for meeting minimum QoS requirements. In the voice calling space, the providers submit quarterly QoS reports on call drops, base station failures, etc which are also audited by third party agencies appointed by TRAI. In fact, very recently the issue of call drops has received significant attention and has become a tussle between the government and telecom providers both pushing back on each other [20]. The telecom providers claim that too little spectrum is available to them and the government also does not help subsidize or acquire sites for the providers to put up new cellphone towers, while the government claims that they have been quite responsive to telco requests and the telecom providers need to utilize their infrastructure more efficiently. Similar regulations on mobile data services are not as mature as yet – a test methodology to measure key QoS parameters has been prescribed by TRAI [26], but the quarterly data is not audited as in the case of voice calling, and penalties or other mechanisms are not in place if the providers do not meet some minimum QoS standards.

A consultation was organized by TRAI recently to discuss whether minimum QoS guarantees for mobile data should be mandated or not [27], and the providers reverted with an argument that any guarantees for 2G or 3G would be hard to ensure because of user mobility, and also linked it with the net neutrality argument in some manner [1].

In this paper, we make an effort to strengthen the argument for having more robust QoS regulations for 2G and 3G services in India. We do this by collecting performance data on throughput, latency, and availability from 20 sites in rural, semi-urban, and urban areas in North India, and give concrete evidence that the QoS provided differs considerably from advertised values, and is widely different in different locations. We then compare these measured values with data reported by the providers to TRAI, and show that the values differ substantially and that more realistic test environments should be mandated by TRAI to mimic the actual user experience with cellular data services. We also show through two examples that better infrastructure monitoring and more careful network configurations can help improve the QoS even with the current infrastructure – providers who have configured smaller buffer sizes in their network elements are able to provide better latency, and similarly providers who have configured less reactive switching between 2G and 3G are able to avoid certain detrimental interactions with the higher layer TCP protocol that impacts the final throughput achieved. Overall, we highlight the need for telecom providers to manage their networks more efficiently and provide better QoS to consumers, who currently do not seem to be getting the service to which they are entitled. As part of a larger effort in this project, we have partnered with a consumer rights organization to now take these findings to TRAI and other regulatory bodies, and push for urgent revisions on QoS regulation in the country.

## 2. RELATED WORK

In terms of methodology, our work comes closest to that of tools like Netalyzer [18] and research on largescale measurement of broadband networks [23, 19, 2, 8, 10], all of which have highlighted the need to understand network conditions from the end user's point of view. Users are impacted by policies and capacity provisioning of edge ISPs with aspects like buffer sizes and traffic shaping having a significant impact on the QoS experienced by the users. Our measurement architecture is similar to Netalyzer, and we have also borrowed several techniques from these papers to probe firewalls, caching, etc reported in our earlier research [17]. Similar work of largescale probing of the network edge is however not common in the cellular data connectivity space. Cellular network measurement research has actively covered TCP behavior on cellular networks [14, 12, 4, 5], and has also seen interest in understanding hardware characteristics such as the radio wake up latency and scheduling policies [3, 9, 21, 25] with the broad objective to tweak protocols so as to obtain better performance on cellular data networks. Our own focus in this paper is not on network protocol modifications to get better performance, but to conduct measurements on 2G and 3G networks in the developing region context of India to understand the current state of QoS experienced by the users. Our goal is to specifically use the measurements to reason about appropriate regulatory mechanisms that can help ensure that telecom providers will work towards better quality of service provisioning for the users.

A similar focus on using measurements to push for stronger regulatory mechanisms is also seen in [24] where it is proposed that providers should publish a "nutrition label" with details about the QoS provided by them, just like how it is mandatory for packaged food items to carry information about their ingredients and nutrition levels. Although telecom providers in India do disclose this information to TRAI which is published on the TRAI website, it is arguable that many consumers may not know about it or consult the website, and therefore a nutrition label like disclosure will help improve consumer awareness. We have included this proposal in our recommendations to TRAI. Another example in more mature regulatory environments as in the United States, is the federal appointed agency Sam Knows [16] which maintains thousands of broadband gateways installed in consumer homes around the country, and collects measurements from these devices to get controlled yet unbiased measurements directly from the end-users. This is another recommendation we have made to TRAI. Our work is therefore taking network performance research and situating it in the context of consumer rights, to lobby for stronger quality of service regulatory mechanisms.

Much of the works cited above conducted measurements in the developed parts of the world. Developing regions are likely to exhibit interesting characteristics though, and several researchers have worked in this context. [15] profiled Internet usage in a small community in Zambia which got its backhaul connectivity through a low bandwidth satellite link, but with fast wireless meshes within the community. It revealed the need for efficient caching and peer to peer solutions to keep the traffic local. [29] conducted measurements in Ghana and found that due to a lack of server infrastructure locally, complex web pages which involved multiple redirections and DNS resolutions led to very high latencies for web browsing. [6] measured broadband and mobile Internet performance in South Africa, and arrived at a similar conclusion that the lack of inter-connectivity between ISPs and distance of the content from the users, led to poor performance and reliability. Our work is along similar directions, of conducting measurements in a developing region. We too draw attention to the degree of connectivity between ISPs and peering with CDNs, along with keeping a focus also on the quality of network configuration and monitoring by different telecom providers in India.

## 3. METHODOLOGY

As reported in an earlier paper [17], during 2013 we collected 2G and 3G measurements from 7 rural and urban locations, having probed each location for a period of at least 3 months. We have since then repeated the measurement exercise in more locations, bringing the count of the total number of rural and semi-urban sites probed to 15, and urban sites to 5. The methodology followed was the same as in our earlier work. We wrote a measurement suite on Linux based Netbooks which were placed at these sites, and were configured to run tests to measure the throughput, latency, availability, etc of 2G and 3G connections provided by different telecom providers. On each Netbook, we attached three Huawei USB modems to be able to probe three different telecom providers simultaneously.

We leveraged our existing relationship with several social enterprises and NGOs, to identify sites where we could safely place our equipment for a long stretch of time over several months, and also refer to local staff members working out of these locations to check or restart the Netbooks if required. With help from the organizations PRADAN and Vikas Samvad, we identified 11 sites in the state of Madhya Pradesh, which were either local offices of these organizations or the homes of their staff members and volunteers. Similarly, with help from the social enterprises Air Jaldi and Gram Vaani, we identified 5 sites in the state of Jharkhand. 3 sites were used in the state of Rajasthan which were homes of family members of some of the authors of this paper. Finally, 1 site was in Delhi out of our lab itself. For all sites, the SIM cards for the 2G and 3G connections were procured locally, and only those providers were probed which steadily gave a high signal strength of at least 20 ASU at the sites. Wherever 3G services were available, we probed both 3G and 2G performance. Note that 3G services are however available only in urban and semi-urban areas, therefore 2G measurements prevail in our dataset. This is reflective of the adoption of 3G services in the country, which was approximately half of 2G adoption during the time we conducted the tests.

The choice of using Netbooks was primarily motivated because of the long battery life of commodity Netbooks, so that the deployments would not require any complex setup with UPSes or solar powered units to manage power failures. We however had to handle several other challenges, which often required strong support from the local staff or friends and family at these locations. One such challenge was that the 2G and 3G connections we bought locally were prepaid connections and hence had to be recharged periodically. We did build watchdogs on the Netbooks which would use AT commands to query the remaining usage quota on the connections, and according send us alert emails so that we could add money through APIs provided by prepaid recharge vendors. There however were several instances when the SIMs lost their validity and had to be replaced. Another challenge was that the USB modems would sometimes hang. Despite watchdogs which would attempt to first re-mount that particular USB modem, failing which the Netbook would be rebooted automatically once tests on other connections had completed, the only failsafe was to unplug and plug the modem for which we had to seek help of the local staff. Finally, due to disk failures or misconfigurations, and massive floods in one site in Madhya Pradesh, the Netbooks themselves had to be replaced at a few locations and required us to travel to the sites with replacements. For these reasons we could not use all the measurement data produced, and had to select for our analysis only those sites and providers for which we got long stretches of good quality data.

Table 1 summarizes the providers and access technologies we were able to probe successfully at the different locations. EDGE and UMTS are 2.5G and 3G technologies respectively belonging to the GSM family, and 1xRTT and EvDO are 2.5G and 3G technologies belonging to the CDMA family. For ease of exposition, we refer to EDGE and 1xRTT as 2G, and UMTS and EvDO as 3G technologies. Among the GSM based providers, we chose BSNL, Airtel and Idea which are among the largest providers in the country, and refer to them as $G_1$, $G_2$ and $G_3$ respectively with the $G$ meant to denote GSM. Reliance is the only operator providing CDMA based services, and we refer to it as $C_1$ with the $C$ meant to denote CDMA. Labels $R_1$ to $R_{11}$ refer to rural locations, $S_1$ to $S_4$ are semi-urban locations, and $U_1$ to $U_5$ are urban locations.

**Table 1: Measurement locations/service providers**

| | $G_1$ (BSNL) | $G_2$ (Airtel) | $G_3$ (Idea) | $C_1$ (Reliance) |
|---|---|---|---|---|
| $R_1$ (Lamta, MP) | EDGE | EDGE | - | 1xRTT |
| $R_2$ (Paraswada, MP) | EDGE | EDGE | EDGE | - |
| $R_3$ (Ukwa, MP) | EDGE | EDGE | EDGE | - |
| $R_4$ (Amarpur, MP) | EDGE | EDGE | EDGE | - |
| $R_5$ (Samnapur, MP) | EDGE | - | - | 1xRTT |
| $R_6$ (Janakpur, MP) | EDGE | - | - | - |
| $R_7$ (Manjha, MP) | - | EDGE | - | 1xRTT |
| $R_8$ (Hanumanpura, RJ) | EDGE | EDGE | EDGE | - |
| $R_9$ (Kulhi, JH) | EDGE | UMTS | - | - |
| $R_{10}$ (Gondlipokhar, JH) | EDGE | UMTS | - | - |
| $R_{11}$ (Getalsud, JH) | UMTS | EDGE | - | 1xRTT |
| $S_1$ (Dindori, MP) | EDGE | - | - | - |
| $S_2$ (Panna, MP) | EDGE | UMTS | UMTS | - |
| $S_3$ (Jatadungari, MP) | - | - | EDGE | - |
| $S_4$ (Ormanjhi, JH) | - | EDGE | UMTS | - |
| $U_1$ (Pondi, MP) | EDGE | - | - | 1xRTT |
| $U_2$ (Jaipur, RJ) | UMTS | UMTS | EDGE | - |
| $U_3$ (Sikar, RJ) | UMTS | UMTS | EDGE | - |
| $U_4$ (Ranchi, JH) | EDGE | UMTS | - | - |
| $U_5$ (Delhi) | UMTS<br>EDGE | UMTS<br>EDGE | UMTS<br>EDGE | EVDO |

Figure 1 shows the key components of the measurement architecture we deployed. The Netbooks consult a control server to get a list of tests and test parameters to execute, and then run these tests such as upload and download iperfs to a different measurement server. Yet another data server is used to collect data from the Netbooks when they are not running the measurement tests. We used virtual machines hosted on Linode for the measurement and control servers, and a server at IIT Delhi as the data server. Interested readers are referred to our earlier paper [17] for more details and reasons behind our choice of this architecture.

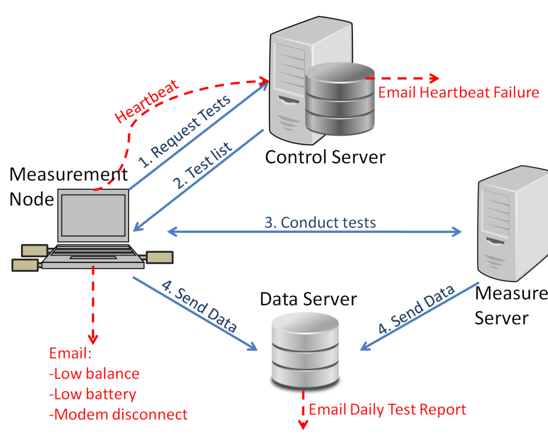


**Figure 1: Measurement architecture. Solid lines test execution steps. Dashed lines show suite health monitoring components**

# 4. QOS MEASUREMENT RESULTS

With the objective to understand the QoS provided by telecom providers, we probed three key metrics: availability, upload and download throughput, and latency, at all the measurement sites. We then compared this data with the values reported by telecom providers to TRAI.

## 4.1 Availability

For each service provider at each location, we evaluated the fraction of time for which connectivity was available. To do so we timestamped network connection and disconnection events reported by the USB modems during the time when the upload/download/latency experiments ran on the modems, and also noted any modem down times during this period when the modem was not responding and remounting attempts were being made by watchdog scripts. Using this we calculated the availability as:

$$availability = \frac{connected\_time}{(measurement\_duration - down\_time)}$$

where *connected_time* is the duration in seconds for which connectivity was available, *measurement_duration* is the time for which experiments ran on the modems, and *down_time* is the duration for which the USB modem may have been in a hung state. Figure 2 shows the availability of service providers across different locations. It is alarming that in some cases the availability is as low as 35%, which means that the modem was able to successfully remain connected to a base station for only 35% of the time for which it tried. Further, we notice that with the exception of $C_1$, availability is typically lower in rural and semi-urban locations than in urban locations. This is probably an artifact of misconfigurations or insufficient monitoring of rural cellular sites, some of which we bring to notice again in subsequent sections.

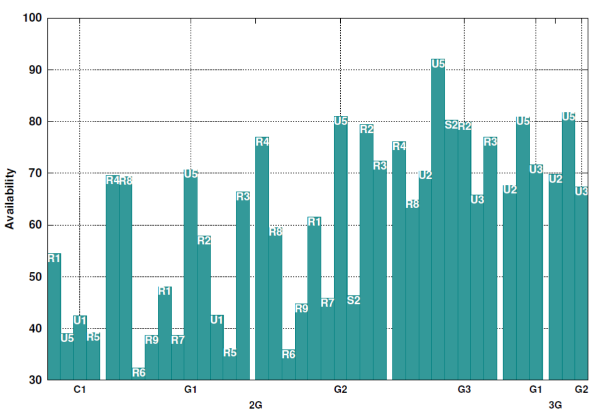


**Figure 2: Availability of service providers across rural and urban locations for 2G and 3G access**

## 4.2 Throughput

We used long duration single threaded TCP tests with iperf on the uplink and curl on the downlink to measure the throughput. Figure 3 shows the 2G and 3G uplink/downlink throughput values. The observations are similar as for the availability data, that some providers like $G_3$ provide consistent performance, but there is a high degree of variability with other providers. We will also bring to notice apparent evidence of misconfiguration in subsequent sections where we show that TCP connections can actually get stalled for tens of seconds at times, thus affecting the throughput, and pointing towards the need for better monitoring and configuration of cellular sites.

What is also alarming is the extent to which the obtained throughput is often much lower than the values advertised by the service providers. Table 2 shows these advertised values. In fact, misleading advertising which promise speeds of "up to 14.4 Mbps" are common sites on wall paintings and billboards all across India, and the Supreme Court even challenged such advertisements [22]. TRAI needs to take note of misleading advertisements which are not only uncompetitive but also hinder the awareness of the consumers in terms of knowing what QoS are they actually entitled.

## 4.3 Latency

We used ping and traceroute measurements to find the latency values to the first IP hop in the network, arguably the GGSN or the PDSN gateway in the cellular network. We also compared this to the end to end latency to the Linode measurement server, to understand what proportion of the latency is spent in the radio access network.

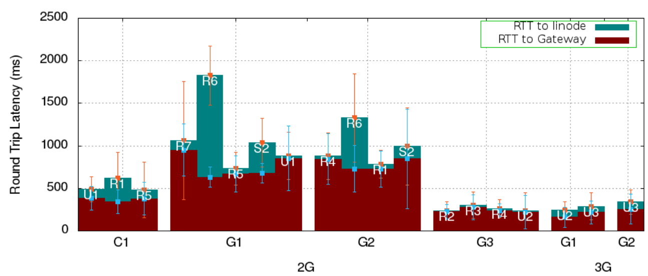


**Figure 4: RTT to the measurement server and cellular gateways**

Figure 4 shows the round trip latencies to the measurement server, and its sub-component to the gateway node in the network. It is interesting to note that providers like $G_3$ are able to provide almost 3G like latencies on 2G connections, but other providers have much higher latencies and also show wide variations across different locations, again pointing towards different network configurations which probably cause these variations. We confirm this by measuring the buffer sizes on the downlink and uplink by sending a train of UDP packets and spotting the first packet which was lost. Figure 5 shows the buffer sizes, and we can observe that providers $G_3$ and $C_1$ with the smallest buffers are also the ones with the smallest latencies. Large buffers are known to lead to high latencies and the bufferbloat problem, where interactive flows suffer when co-existing with long flows [10]. We therefore argue that just more careful network configurations can alleviate several QoS problems, without any need to provision additional infrastructure.

Another observation from Figure 4 is that the latency beyond the gateway is also lower for $G_3$ and $C_1$, indicating that these ISPs are likely to have better connectivity with

**Table 2: Throughput values advertised by service providers**

| | 2G Uplink | 2G Downlink | 3G Uplink | 3G Downlink |
|---|---|---|---|---|
| C1 | - | 20 Kbps | - | 256 Kbps |
| G1 | 144 Kbps | 144 Kbps | 14 Mbps | 14 Mbps |
| G2 | 256 Kbps | 256 Kbps | 21.1 Mbps | 21.1 Mbps |
| G3 | 236 Kbps | 236 Kbps | 5.7 Mbps | 21 Mbps |

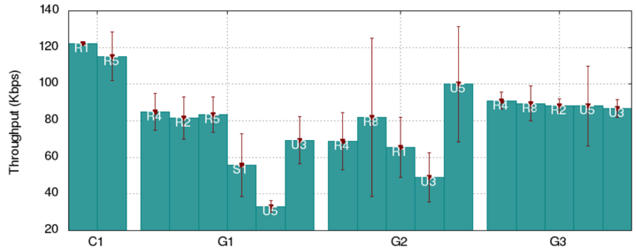


(a) 2G Uplink Throughput

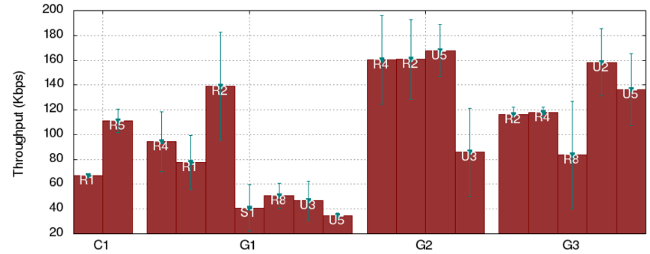


(a) 2G Downlink Throughput

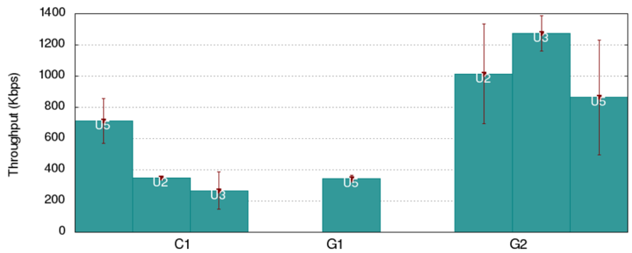


(c) 3G Uplink Throughput

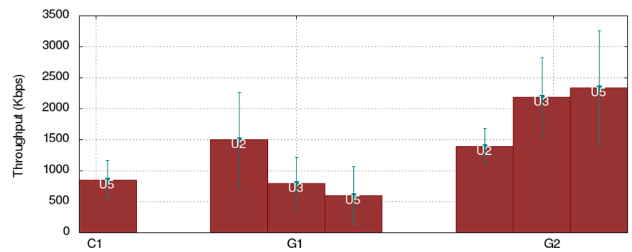


(d) 3G Downlink Throughput

**Figure 3: Measured throughput values across various locations and service providers**

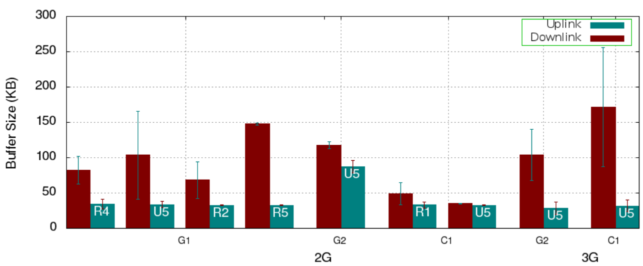


**Figure 5: Uplink and downlink buffer sizes (KB)**

the rest of the Internet. We explore this in greater detail in subsequent sections when we look at inter ISP connectivity.

## 4.4 Comparison with TRAI reported values

TRAI releases a quarterly report on the quality of service provided in different states for cellular data services [28]. These are self reported values by the telecom providers, based on a test methodology specified by TRAI [26]. On the surface the test methodology looks similar to our own, where TRAI specifies the file size and number of tries for upload and download to measure throughput, and the use of ping to find latencies. The tests however are actually conducted in a controlled environment most likely from network elements located deeper inside the radio access network which does not accurately mimic the real world environment that users actually experience. Table 3 shows the data from the TRAI report for the service provider Airtel ($G_2$), and compares it with the parameters as measured by us during the same period. Broadly, the throughput values are of the same order, the latency values measured by us are quite higher, but the greatest difference is in the availability values where the provider actually reports 100% availability in most cases!

The takeaways therefore from a QoS regulation standpoint is for TRAI to mandate more realistic environments for providers to conduct network tests, or to depute third party agencies such as Sam Knows in the United States to report measurements collected from actual user devices [16]. These measures would bring more data to the public domain and thus draw attention to the widely different QoS provided in different locations by different providers. Proposals such as nutrition labels [24] could also help empower consumers by making them more aware of the QoS to expect and then use the information to make better choices when buying data plans. However, unless TRAI does not mandate some minimum QoS standards to which providers can be held accountable, or the published information is not made available to consumers easily to be able to exercise their choice in selecting providers, even these stronger regulatory measures may arguably not yield much benefits. We therefore believe that TRAI should continue to mount pressure on the providers to manage their networks better since our data indicates that just careful network configurations alone can help to a significant extent.

# 5. NETWORK CONNECTIVITY

We next explore inter-ISP connectivity and CDN linkages of cellular providers because these aspects influence the quality of service perceived by users [29, 6].

We start with listing some common CDN providers, and then use traceroute and the RIPE database to chalk out AS hops to these CDN networks from the four cellular providers probed by us. Figure 6 shows the consolidated AS map. We can see from the map that Bharti Airtel and Tata Communications (which acquired a controlling stake in the state owned VSNL network in 2002 [11]) hold dominant positions in providing backbone connectivity for India with the rest of the world. We also interestingly find that providers like $G_3$

**Table 3: Comparison of performance metrics observed by us with TRAI reported values by service providers**

| Location | | Observed Values | | | TRAI Values | | |
|---|---|---|---|---|---|---|---|
| | | Availability (%) | Throughput (kbps) | Latency (ms) | Availability (%) | Throughput (kbps) | Latency (ms) |
| RJ | 2G | 59.49 | 131.02 | 308.45 | 100 | 116.5 | 206.3 |
| | 3G | 69.73 - 71.58 | 1394.03- 2187.13 | 128.84 - 256.89 | 100 | 2476.3 | 191.8 |
| JH | 2G | 41.97 - 57.64 | | 453.85 - 1040.75 | 100 | 199.7 | 244.8 |
| | 3G | | | | 100 | 3485.2 | 64 |
| MP | 2G | 36.37 - 80.17 | 85.89 - 160.72 | 405.16 - 936.45 | 99.83 | 131.8 | 119 |
| | 3G | 73.66 | 907.80 | 476.57 | | | |
| Delhi | 2G | 81.44 | 167.82 | 310.58 | 100 | 130.2 | 173.3 |
| | 3G | 91.69 | 2332.25 | 540.41 | | 2369.4 | 80.4 |

(Idea Cellular) do not own any backbone links of their own, but use service from $G_2$ (Bharti Airtel) and other providers, and are still able to provide better latency values.

It it also clear that the telecom providers are well networked with CDNs. Akamai has its servers located in all the provider networks except Idea. Most providers also peer directly with many CDN providers, showing that they are cognizant of the benefits of CDNs. To study this further, and to see implications of the AS connectivity on the quality of experience for web browsing, we measure the page load times for several webpages accessed from different providers.

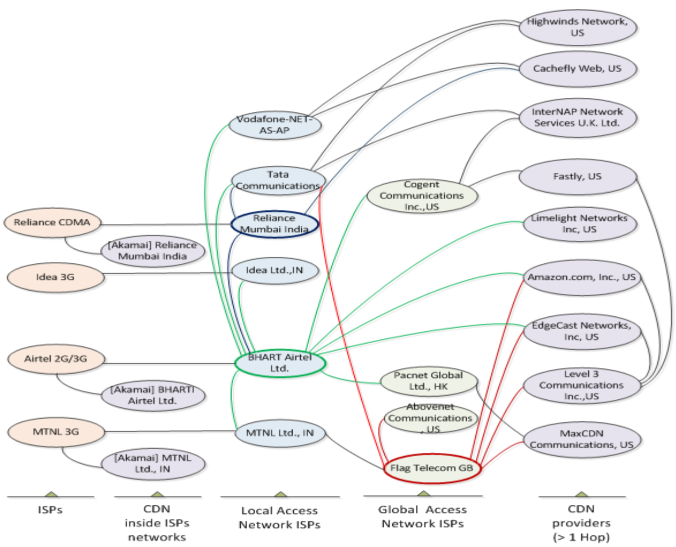


**Figure 6: AS map for inter-ISP connectivity**

Page load time is certainly not the best or the only metric to assess the quality of experience for web browsing, but we chose it as a first step which could be executed easily. We used the Selenium browser automation framework with the Chrome driver on the Netbooks, to probe a number of different websites and collect HAR traces which could be analyzed to find the page load times and other HTTP dynamics. We ran this test in New Delhi for different categories of websites including news websites, government websites, and e-commerce websites. The websites were chosen carefully as those which use a CDN provider to serve at least 80% of their content, and those which do not use CDNs at all. We also ensured that the websites we selected had similar cumulative page sizes of between 2-3MB.

Figures 7 (a) and (b) show page load times for webpages that do not use CDN providers to host their content, and Figures 7 (c) and (d) show page load times for webpages that do use CDNs. Tables 4 and 5 gives the list of webpages probed. A comparison for webpages with and without CDNs shows that with 2G access technology, the page load times are better by almost 50% when using CDNs, and with 3G access technology they are almost 25% better. The second observation is that irrespective of whether or not CDNs are used, $G_2$ (Airtel) consistently gives the best page load times followed by $G_3$ (Idea), and $G_1$ (BSNL) gives the worst. Since most of the CDN using websites we probed use Akamai, Airtel was able to do quite well thanks to its close linkage with Akamai. BSNL however despite its close linkage has a poorly configured network, and is not able to utilize its CDN linkages. This clearly shows that network configuration and connectivity both have important implications on the quality of experience that providers can ensure to their users.

We also advocate that a crowdsourcing platform which collects such QoE parameters for different websites and telecom providers, across locations, can help consumers make better choices when buying data plans. This would in turn usher more competition between the providers. TRAI itself could host such a platform, or help create a platform run by neutral agencies.

**Table 4: Pageload time for websites without CDN**

| Website alias | Page alias |
|---|---|
| http://www.ibps.in/ | W1 |
| http://www.incredibleindia.org/en/ | W2 |
| http://india.gov.in/ | W3 |
| https://www.irctc.co.in/eticketing/loginHome.jsf/ | W4 |
| http://www.kerala.gov.in/ | W5 |
| https://www.maharashtra.gov.in/ | W6 |
| https://morth.eproc.in/ProductMORTH/publicDas/ | W7 |
| http://www.odisha.gov.in/portal/default.asp/ | W8 |
| https://www.sbi.co.in/ | W9 |
| http://www.tn.gov.in/ | W10 |

## 6. CONNECTION STALLS

In our earlier work [17] we reported a curious observation that large file downloads on TCP would sometimes stall for long periods of time of up to tens of seconds. During these stall events, a lost packet would result in multiple timeouts and get retransmitted repeatedly. However, unlike regular

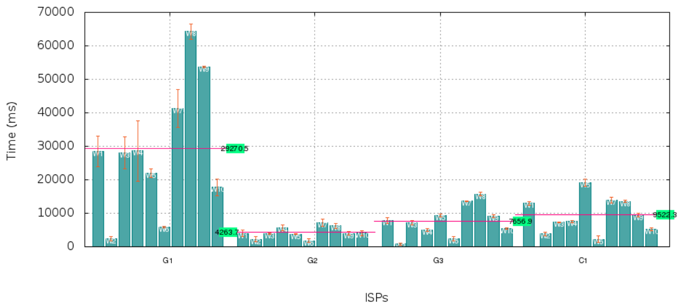


(a) 2G : without CDN

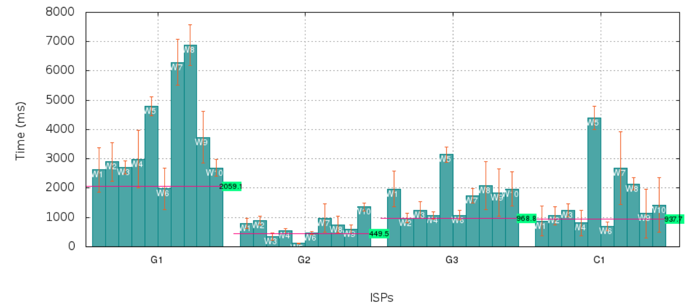


(b) 3G: without CDN

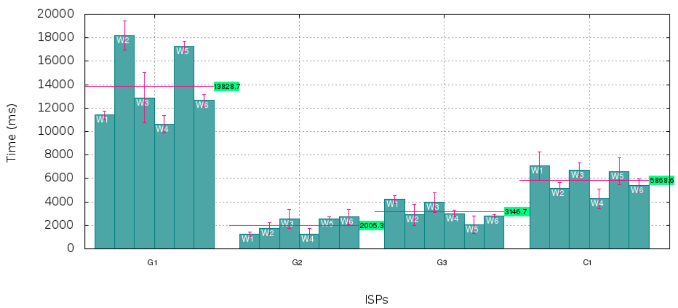


(c) 2G : with CDN

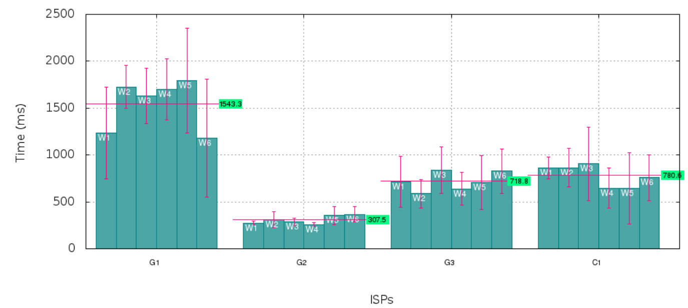


(d) 3G : with CDN

**Figure 7: Comparison of Page Load times for websites using and non using CDN services**

**Table 5: Pageload time for websites with CDN**

| Website alias | Page alias | CDN provider |
|---|---|---|
| http://timesofindia.indiatimes.com/ | W1 | Akamai |
| http://www.espncricinfo.com/ | W2 | Akamai |
| http://www.bbc.com/bbc/ | W3 | Fastly |
| http://www.bostonmagazine.com/ | W4 | Internap |
| http://www.news.com.au/news/ | W5 | Akamai |
| http://www.yatra.com/yatra/ | W6 | Akamai |

timeout situations where all packets following the lost packet are also lost and later retransmitted, in this case the packets were not lost but would sit in some buffers along the way and eventually reach the receiver without requiring any retransmissions. We found that the occurrence of these stalls was quite frequent in some locations and detrimental for performance because they would cause the entire TCP connection to pause and later initiate a slow start. We validated earlier that the stalls were not an artifact of the modem, or the server, or due to any middleboxes that might be buffering out of order packets, and we had not been able to satisfactorily explain the phenomenon. We now have an answer and it has to do with device and network configurations which cause the USB modems to search for other networks periodically, and thereby temporarily suspend data transfer – the signal strength drops to zero and a new connection has to be initiated after the search procedure is over. Since this happens more in some sites than others, it again points towards configuration problems that lead to such events.

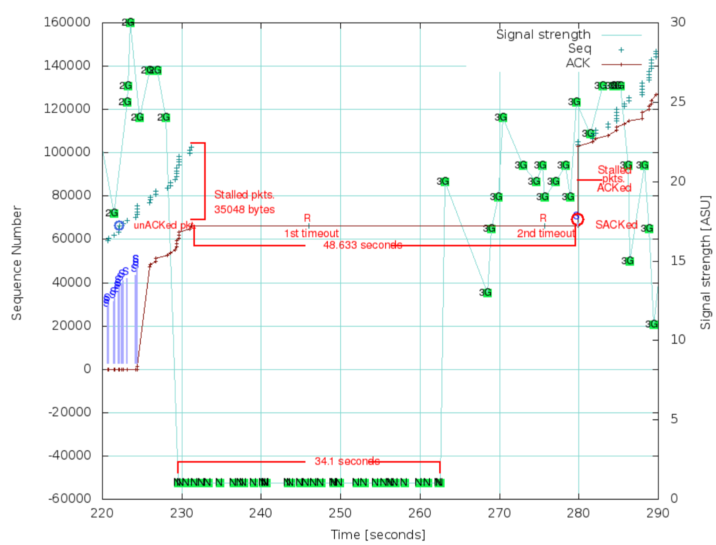


**Figure 8: Example of a stall event**

Figure 8 shows such a stall event on the sender side trace. The signal strength at the receiver is also plotted on the same graph. We can see that when acknowledgements are not received after time 232, the first unacked packet is retransmitted twice during timeout events, and finally when the acknowledgement is received at time 240, it is soon after followed by an acknowledgement for the stalled packets without requiring any retransmissions for these packets. This shows that the stalled packets were actually transmitted just once before time 232; these packets were then waiting in some buffers along the way and reached the receiver together with the retransmitted packet. The signal strength trace shows that during this stall event, the receiver actually switched from a 2G to a 3G connection, and registered a signal strength of zero when the switching was in progress. A similar analysis of the receive side trace showed that the seemingly lost packet was actually received during its first transmission itself but the ack could not be transmitted because the device lost its signal. Therefore, due to the switching not only did the connection stall for more than 50 seconds, but TCP also went into repeated timeouts.

We investigate this in more detail. Figure 9 (a) shows the latency during different kinds of switching events that happen. Locations such as $R_5$ and $R_6$ on connection $G_1$ only had 2G available and see switching from 2G to 2G with latencies as high as 20 seconds. Thus, even in the case of zero mobility, the modems go into a search phase like in a handoff. Other locations which had 3G available, see switching from 3G to 2G, 2G to 3G, and 3G to 3G, pointing to even more drastic outcomes from the search process. We do not know whether the search process was initiated by the mobile or the network, since both GSM and CDMA technologies allow both kinds of handovers, but the essential point this observation highlights is to reduce unnecessary events either through better configuration settings or more intelligent algorithms.

Figure 9 (b) shows the frequency of switching, with many locations witnessing a switching event every 40 seconds on average. The negative axis on this figure shows how often these switches led to stall events where there was at least one timeout, and many locations have such stalls once every minute. There are however also combinations of providers and locations with low switching frequencies, but such cases of good configuration settings are seemingly rare.

Figure 9 (c) plots the duration of these stall events. As in the example shown earlier, these stall durations run into tens of seconds. The events are therefore detrimental in the case of any long transfers since almost every minute some tens of seconds of connectivity goes unutilized, and in the case of TCP this further sends the connection into timeouts.

A deeper analysis of some locations reveals that 3G-2G switches are the most frequent, and in almost 50% of the cases these switches lead to stall events. This is shown in Figure 9 (d) for location $R_7$ with $G_1$ (BSNL). Other kinds of switches such as 3G-3G and 2G-3G are less frequent, and also lead to fewer stalls. A possible reason behind this observation that 3G-2G switches have the worst outcome, is that when a connection is in 3G before the switch then the latencies are lower and therefore there is a higher chance of timeouts occurring because of the switch.

Overall, we found that long duration downloads in most locations had stalls more than 40% of the time, and in some cases almost 90% of the downloads had stalls. Downloads without stalls gave a throughput 25% higher than downloads with stalls in the case of 2G, and 65% higher in the case of 3G connections. This seems like it could be an avoidable penalty: In the case of sites with only 2G access, repeated searching could be disabled, or at least made less frequent. With sites having 3G and 2G access, a deeper analysis of the provider logs should be done to check if 3G access is being deliberately downgraded to 2G by the providers because their networks are underprovisioned, or the networks are just misconfigured and cause unnecessary switches.

When considered in perspective with the latency and throughput measurements in the earlier sections, it is clear that some providers are able to provide more consistent performance than others, and some sites are better configured than others. These observations therefore point towards the need for providers to be more careful in managing their networks, which can either be ensured through stronger and more appropriate regulations or through greater consumer awareness, so that the providers can be pushed to work harder at delivering better performance.

# 7. CONCLUSIONS

We showed through our measurements that the quality of service obtained by users differs considerably from advertised values by the telecom providers, and from values reported by them to the telecom regulatory authority in India. We also showed that in many cases just more careful configurations of the cellular networks could lead to better performance. We are now working together with a consumer rights organization to take up this evidence to the government regulators and argue for stronger QoS regulations in the country.

# 8. ACKNOWLEDGEMENTS

This work was supported with funding from the Ford Foundation. We are deeply grateful to our field partners, PRADAN, Vikas Samvad, Air Jaldi, and Gram Vaani, who provided invaluable support to run the experiments. We also acknowledge the strong support that CUTS (Consumer Unity and Trust Society) has provided to help us understand the QoS regulatory framework in India, and have partnered with us to work towards stronger regulations in the country.

# 9. REFERENCES

[1] Personal communication with TRAI officials, June 2015.
[2] A. Akella, S. Seshan, and A. Shaikh. An Empirical Evaluation of Wide-area Internet Bottlenecks. In *Proc. IMC*, 2003.
[3] J. J. Alcaraz, F. Cerdan, and J. Garcia-Haro. Optimizing TCP and RLC Interaction in the UMTS Radio Access Network. *Network Magazine of Global Internetworking*, 2006.
[4] P. Benko, G. Malicsko, and A. Veres. A Large-scale, Passive Analysis of End-to-End TCP Performance Over GPRS. In *Proc. INFOCOM 2004*, 2004.
[5] R. Chakravorty, A. Clark, and I. Pratt. GPRSWeb: Optimizing the Web for GPRS Links. In *Proc. International conference on Mobile Systems, Applications and Services*, 2003.
[6] M. Chetty, S. Sundaresan, S. Muckaden, N. Feamster, and E. Callandro. Measuring Broadband Performance in South Africa. In *Proc. ACM DEV*, 2014.
[7] 3G consumer. Pathetic Internet Speeds on 3G, http://bit.ly/1KH1JT5, http://bit.ly/1KH1JT5, Sept 2015.
[8] M. Dischinger, A. Haeberlen, K. Gummadi, and S. Saroiu. Characterizing Residential Broadband Networks. In *IMC*, 2007.
[9] A. Elmokashfi, A. Kvalbein, J. Xiang, and K. Evensen. Characterizing Delays in Norwegian 3G Networks. In *Proc. PAM*, 2012.
[10] J. Gettys. Bufferbloat: Dark Buffers in the Internet. *Internet Computing, IEEE*, 15(3):96–96, 2011.
[11] Tata Group. Tatas Vault Ahead in Telecom, Acquire 25% Stake in VSNL, http://bit.ly/1V5WxyN, February 2002.
[12] H. Inamura, G. Montenegro, R. Ludwig, A. Gurtov, and F. Khafizov. TCP over Second (2.5G) and Third (3G)

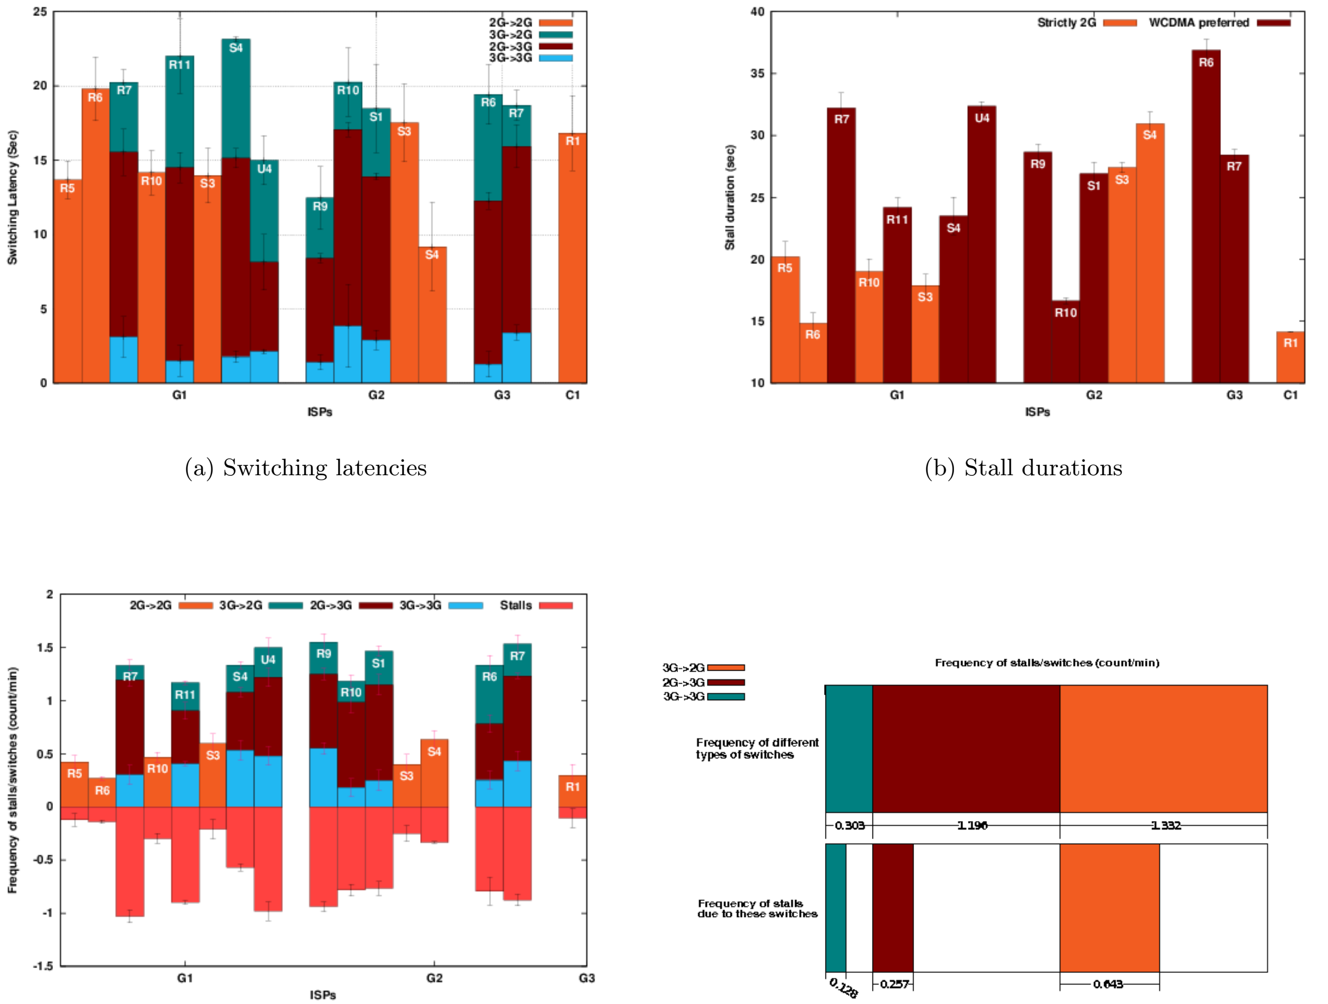


(a) Switching latencies (b) Stall durations

(c) Frequency of stalls and switches (d) Proportion of switches leading to stalls

**Figure 9: Switches and stalls at different locations**


Generation Wireless Networks. RFC 3481 (Best Current Practice), February 2003.

[13] McKinsey Inc. Offline and Falling Behind: Barriers to Internet Adoption, http://bit.ly/Ztobeu, September 2014.

[14] H. Jiang, Y. Wang, K. Lee, and I. Rhee. Tackling Bufferbloat in 3G/4G Networks. In *Proc. IMC*, 2012.

[15] D. Johnson, E. Belding, K. Almeroth, and G. Stam. Internet Usage and Performance Analysis of a Rural Wireless Network in Macha, Zambia. In *Proc. NSDR*, 2010.

[16] Sam Knows. https://www.samknows.com/, September 2015.

[17] Z. Koradia, G. Mannava, A. Raman, G. Aggarwal, V. Ribeiro, A. Seth, S. Ardon, A. Mahanti, and S. Triukose. First Impressions on the State of Cellular Data Connectivity in India. In *ACM DEV*, December 2013.

[18] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Netalyzr: Illuminating the Edge Network. In *Proc. IMC*, 2010.

[19] K. Lakshminarayanan, V. Padmanabhan, and J. Padhye. Bandwidth Estimation in Broadband Access Networks. Proc. IMC, 2004.

[20] N. Narajayan and Scroll.in. Tower Shortage, Spectrum Crunch and an X Factor - Why Call Drops are Common in India, http://bit.ly/1N5ltAx, August 2015.

[21] J. Prokkola, P.H.J. Perala, M. Hanski, and E. Piri. 3G/HSPA Performance in Live Networks from the End User Perspective. In *IEEE International Communications Conference*, 2009.

[22] PTI and Times of India. Delhi Highcourt to Test MTS's 9.8 Mbps Internet Speed Claims, http://bit.ly/1QTUGY1, 2014.

[23] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Broadband Internet Performance: A View from the Gateway. In *Proc. ACM SIGCOMM*, 2011.

[24] S. Sundaresan, N. Feamster, R. Teixeira, A. Tang, K. Edwards, R. Grinter, M. Chetty, and W. Donato. Helping Users Shop for ISPs with Internet Nutrition Labels. In *Proc. ACM SIGCOMM Workshop on Home Networks*, 2011.

[25] W. Tan, F. Lam, and W. Lau. An Empirical Study of 3G Network Capacity and Performance. In *Proc. IEEE INFOCOM*, 2007.

[26] TRAI. The Standards of Quality of Service for Wireless Data Services Regulations, 2012, http://bit.ly/1V5UlY9, 2012.

[27] TRAI. Consultation Paper on Amendments to the Standards of Quality of Service for Wireless Data Services Regulation, 2012, http://bit.ly/1QTUGY1, April 2014.

[28] TRAI. Quality Performance Monitoring Report (PMR) on Quality of Service of Wireless Data Services, http://bit.ly/1FY1MV6, March 2015.

[29] Y. Zaki, J. Chen, T. Potsch, T. Ahmad, and L. Subramanian. Dissecting Web Latency in Ghana. In *Proc. IMC*, 2014.