

Wireless Broadband Measurement in California

YoungJoon Byun, Sathya Narayanan, Stephanie Mott
California State University, Monterey Bay
Seaside, CA, USA
ybyun@csumb.edu, snarayanan@csumb.edu,
smott@csumb.edu

Jason Schwenkler
California State University, Chico
Chico, CA, USA
schwenkl@gic.csuchico.edu

Ken Biba
Novarum, Inc.
ken@novarum.com

Robert Osborn, Michael Morris
California Public Utilities Commission
San Francisco, CA, USA
robert.osborn@cpuc.ca.gov, michael.morris@cpuc.ca.gov

Abstract—An integrated software tool using the open source applications iPerf and ping was developed to measure broadband wireless performance in California. Statistics on TCP download/upload throughput, latency, UDP packet jitter and loss, network type, and network provider were recorded. The software can be installed on a Window-based laptop and Android-based smartphone. An analysis script was also developed to analyze measurement results on the database server. Using the software, testers conducted field test at 1,200 different locations in California, and the test result will be accessible to the general public using a viewer application. In addition, selected locations ran Glasnost, a traffic shaping detection tool, to determine if network providers altered performance of flows based on type or port number.

Keywords-component; mobile performance; iPerf; field test; Android smartphone; State of California

I. INTRODUCTION

The California Public Utilities Commission (CPUC) wanted to conduct vendor-independent wireless network performance testing to measure the performance of major mobile wireless networks in California as an honest broker. As a part of these efforts, we have developed and integrated software tools to measure wireless bandwidth and performance. The software is based on open-source applications iPerf [1] and ping. Test results are displayed as the test is executing, and then uploaded to a database server for analysis. Through the project, we would like to objectively evaluate major providers of mobile wireless broadband service across the state of California and provide Californians with additional information about their mobile broadband connection. Furthermore, we hope that we are able to build awareness about the importance of mobile broadband in getting information and services from the Internet.

Fig. 1 describes an overall architecture for our wireless broadband measurement. We set up two test servers, one on the East coast in Northern Virginia and one on the West coast in Northern California. On the Windows-based laptop, broadband data was collected for four providers, Verizon,

AT&T, T-Mobile, and Sprint, using four wireless datacards or modems. Same set of data was also collected using four Android smartphones for the same four providers. Test results were uploaded and analyzed on a database server where eventually they will be made available to the public via a website and smartphone viewer applications.

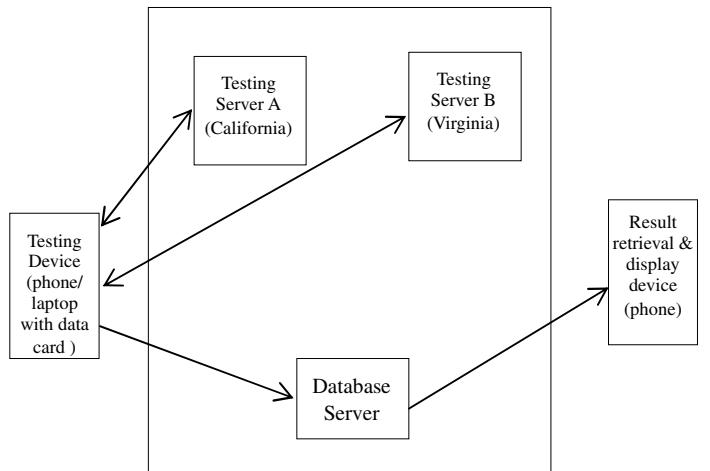


Figure 1. Overall architecture for wireless broadband measurement.

II. MEASUREMENT TOOLS

For each test, we record the time, location (i.e., latitude and longitude), round trip time (RTT), TCP download and upload throughput, UDP jitter and datagram loss, network type (i.e., EDGE, EVDO, GPRS, HSPA, LTE, UMTS, etc), network provider (i.e., AT&T, Sprint, T-Mobile, Version), and traffic shaping information (in a subset of the locations). To determine the location, the Android phone uses the onboard satellite GPS while the laptop uses a GlobalSat BU 353 USB GPS receiver. The iPerf tool measures throughput and bandwidth of TCP protocol between a tester device and an iPerf server. It also measures the data loss and jitter of UDP protocol. The ping tool provides the minimum, maximum, and average RTT by sending several short

packets to a server. It can also indicate packet loss during the packet traversal. The glasnost tool checks if a network has traffic shaping [2]. In case of a weak or low signal, a timeout of 2 minutes was added for each test.

When a tester conducts a measurement in a location for a specific wireless provider, our integrated software implicitly runs the tools of Fig. 2 in the following sequence:

- (1) GPS capture: A GPS device reads the latitude and longitude information from satellites to identify the accurate location of the tester.
- (2) Connectivity check: Using ping, we check if a minimum bandwidth is available to the server to continue with the test at the location. For this purpose, the software sends several ICMP ping packets for four seconds. If the software doesn't get any response from the server during the period, we conclude that the area has very weak signal or no signal for the provider and stop the test sequence. If ping responses are received, our software moves on to the next step in the test sequence.
- (3) iPerf TCP test: The iPerf client software conducts TCP upload and download measurements to the East and West coast servers for 10 seconds, respectively. Even though we request the measurement for only 10 seconds, it could take much longer in a weak signal area. To address such situations we added a 2 minute timeout to each of the individual tests in the sequence.
- (4) Ping to West server: The ping tool measures the RTT to the West coast server for 10 seconds.
- (5) iPerf TCP test: We conduct the TCP upload and download measurements one more time.
- (6) Ping to East server: The ping tool measures the RTT to the East coast server for 10 seconds.
- (7) iPerf UDP test: iPerf tool measures the UDP jitter and datagram loss.
- (8) Test result upload: When all measurements finish, the testing result is uploaded to our database server in the text file format for post-processing.

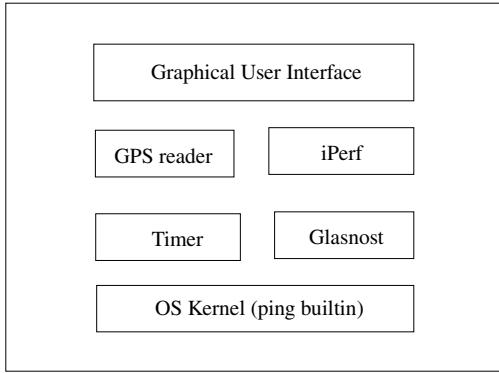


Figure 2. Software architecture of measurement software.

A. iPerf

iPerf is a software tool to measure the network throughput and performance [1]. It is developed based on the

client-server architecture where the client side software requests a TCP connection to the server side software. After the connection is established, the tool measures the throughput between the two ends. For the TCP upload bandwidth measurement, the client software makes a TCP connection to the server software and sends data streams to the server for 10 seconds. After that, the server software sends data streams back to the client for download bandwidth measurement. Note that the original iPerf tool uses two separate TCP connections for the upload and download measurements, respectively. But since many network operators block incoming TCP requests to client device through the use of NAT/Firewall, we reuse the TCP connection used for the upload measurement and use it for the download measurement as well. As for other measurement parameters, the tool uses 64 Kbytes window size (-w 64k) and executes four threads concurrently (-P 4), which can increase the data volumes to be exchanged between the client and server in many cases. For the UDP jitter and datagram loss rate measurement, the iPerf client sends data to the server for either one or five seconds. Note that there's no data streams from the server to the client in the UDP measurement. As other parameters, we use 220K buffer length (-l 220) and 88K bits/sec bandwidth (-b 88k) to simulate VoIP data.

The iPerf tool is composed of several threads such as the main, reporter, listener, client, and server threads. When it runs on a server machine, the main, reporter, and listener threads execute continuously to handle multiple requests from several clients. For a client request, the listener thread launches a server thread that handles the request. Additionally, the listener thread launches a client thread to measure the bandwidth from the server to the client later. The reporter thread presents all measurement results with time information.

B. Glasnost

Glasnost is an open software system to detect whether an ISP (Internet Service Provider) is differentiating between flows of applications [2]. It compares the throughput of a pair of flows to determine if traffic differentiation is present. For instance, Glasnost determines if an ISP rate limits BitTorrent protocols that are used for peer-to-peer file sharing. This type of traffic is commonly suspected of being subject to rate limiting because it is often used to transfer large files which consume network bandwidth. One of the flows contains a BitTorrent payload and the other flow is a TCP connection containing unidentified payload bytes. If an ISP is rate limiting BitTorrent protocols, the throughput experience by this control flow will be higher than that of the BitTorrent flow. This is how Glasnost determines if an ISP is rate limiting their traffic. However, network connections are often subject to noise interference, due to cross traffic on the network. To overcome this problem, flows are run multiple times back to back and along the same network path. In addition, flows of the same type are compared and if the identical flows differ significantly from one another the connection is determined to be noisy and is not used for analysis.

Rate limiting could occur when an ISP intentionally slows down a user's Internet connection speed. ISPs would employ rate limiting for the purpose of preserving available bandwidth. For example, if multiple users are downloading BitTorrent files and consuming a large amount of bandwidth, the overall network connection for all the users can be affected. By limiting the amount of data that individual users can download or upload ISPs conserve bandwidth and make the general network connection faster for the majority of users. When ISPs, however, advertise that they provide their users with unlimited Internet connection and then employ rate limiting, they are not truly providing unlimited Internet connection. The purpose of Glasnost is to promote network transparency by allowing Internet users to see if their ISP is rate limiting their connection. While we used Glasnost to examine only BitTorrent, the presence of rate limiting may be indicative of similar behavior for other traffic types.

C. Tool Integration

1) Android Version: Our Android integration application is based on the Network Diagnostic Tool (NDT) Android client that is available at the Measurement Lab [3, 4]. The original program allows the user to connect to a selectable NDT server, and then runs a 10 second test, displaying intermediate test results on a scrolled text screen area. The test results could be sent to the user via an email message. Our tool employs a two-tabbed screen, the front tab labeled Standard which displays a scrolled text area summary of the program status, and the other labeled Statistics which displays the detailed scrolled output of each test while it is executing. There are four buttons on the bottom of the Standard screen: Standard starts the standard test, Custom allows running of pre-selected custom tests through the menu button, Quit exits the current test run, and Upload Results allows uploading of the result files from previous measurements that have not been uploaded due to the weak or no wireless connection. The application allows a tester to select measurement parameters and testing server.

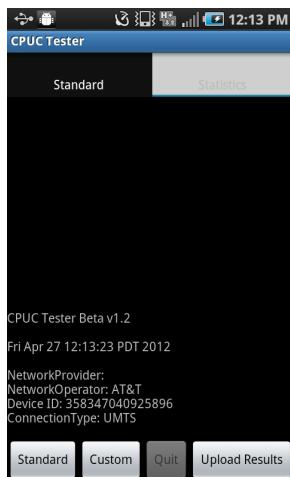


Figure 3. Screenshot of Android-based smartphone application.

2) Laptop Version: Our laptop application is developed to measure the performance of a wireless data card used in a laptop. It is written in Java and designed to run under the Windows 7 operating system. Since no on-board GPS exists in a typical laptop, an USB GPS receiver is required for the application. The laptop also serves as a station for receiving program updates and uploading the Android application to the phone via the USB port. The laptop program also employs a two-tabbed screen. The front or Standard Test tab displays the testing date and time, location ID, provider, progress bar, and scrolling text area for a summary of the test. It also has START, QUIT, UPLOAD, and EXIT buttons. The Custom Settings tab provides measurement parameter setting, GPS configuration, and detailed test results.

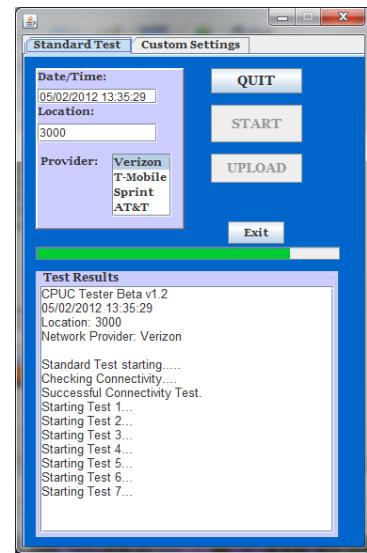


Figure 4. Screenshot of Windows-based laptop application.

III. FIELD TEST AND TEST RESULTS

Using the software mentioned in Section II, we conducted our first field test at 1,200 different locations throughout the state of California in May 2012 [5]. The locations were randomly selected in urban (34%), rural (55%), and tribal (11%) areas. For the testing, eight testers drove over 35,000 miles and equipped with four smartphones and four data cards for four major wireless network providers such as AT&T, Sprint, T-Mobile, and Verizon. The testers ran the testing software in a stationary environment inside an automobile.

A. Field Testing Issues

During the field tests, several technical and non-technical issues came up. The following presents some issues.

GPS reading: Tests were intended to be conducted for all four carriers at the exact same location. Due to minor differences in GPS readings as well as practical problems of testers sometimes being unable to conduct all tests as the

exact same location, GPS coordinates for all tests performed at each location were averaged to a single set of latitude/longitude coordinates. Sometimes, GPS data was not collected because there was no valid GPS signal. In that case, the initial estimated GPS information originally assigned to that location was used.

Abnormal test result values: In some locations, especially weak wireless signal areas, we found several abnormal test results. For instance, when a test starts, we run the ping program to our West coast server to check the wireless connectivity. If a location failed the connectivity test three times, we concluded that the location has "no effective service". Although we measured the TCP upload and download for 10 seconds respectively, our measurement tool could spend more than 120 seconds because of too slow data communication between a client device and a server. If our tool failed to get the TCP testing result within 120 seconds, our testing result generated "timeout". Sometimes, our program had an iPerf internal software error message such as "write1 failed", "write2 failed", and "connect failed", which could have many reasons for the error message, including very weak wireless signal. In our test results, we marked it as a "connect error".

Testing time and difficulties: For each field test, eight testers visited a total of 1,200 locations throughout the state, a challenging schedule for both the testers and the support team. There were four weeks of testing which meant many long days of driving and testing, many of which were in very remote areas. Additionally, we used four different data cards in a single laptop, which is not usual case to the general user. The process of changing the datacards on the laptop, waiting for the device to be recognized, and then running the connections managers tended to be tedious and error prone.

B. Test Results

Since it was the first testing of a series drive tests, it is premature to make concrete conclusions about the wireless performance of the four major operators. However, we observed the following [6]:

- (1) The large gap in throughput between non-rural versus rural and tribal locations demonstrates that the rollout of Long Term Evolution Networks (LTE) has been focused primarily in higher density, urban areas.
- (2) AT&T and Verizon appear to provide the best coverage and highest average throughputs, with Verizon outperforming A&T at the higher speeds, but AT&T delivering a much wider geographic distribution of mid-range speeds across the 1,200 locations we tested.
- (3) The number of locations meeting the combined 6 megabits per second downstream and 1.5 megabits per second upstream threshold on mobile is still limited.

For more detailed information on the test result and analysis, read the Spring 2012 mobile broadband field testing report [6]. By definition of the Federal Communications Commission (FCC), mobile service qualifies as broadband if the downstream throughput is 768 kbps (kilobits per second)

or higher and upstream throughput is 200 kbps or higher. Note that we collected measurements even when they fell below this threshold for the study. This included locations in California where no effective service was available for any mobile operator. Meanwhile, we would like to emphasize the limitations of mobile testing in our study. Planning and operating networks take into account the probability of a user's location and forecasting of aggregate demand for a cell site, both which may vary, depending on factors such as the time of day, location and topography of the test location, network loading and congestion, and device hardware and software limitations. In addition, test results show end user experience at a specific time and specific location and do not necessarily represent end user experience for an entire census tract or census tract.

IV. RELATED WORKS

Measuring bandwidth of communication networks is important to consumers, policy makers, and network carriers [7]. However, measurement results depend on many factors such as the particular tool used for measurement, testing time, network environment at the measurement location, indoor/outdoor, etc. There are two fundamental approaches to be used to measure network performance: a hardware approach and a software approach [8]. Hardware approaches involve a hardware device that is physically connected with a consumer's Internet connection. The device can capture the data flow between the consumer's home and remote servers on the Internet. This approach can provide accurate testing result compared to software approach. However, the hardware-based measurement requires hardware cost and limited samples. Furthermore, this approach is not easily utilized by the general public. Because of these reasons, software-based measurements are more common. For example, a customer can open a web browser on a desktop computer and start a measurement using a web-based test tool. However, the software-based approaches do not always measure the accurate bandwidth.

Measuring Broadband America project conducts a nationwide performance study of residential wireline broadband service, focusing on DSL, cable, and fiber-to-the-home technologies, in the United States [8]. The study measured the services of several wireline broadband providers using a hardware device. Approximately 9,000 volunteers participated to the study, and TCP upload/download, UDP latency, UDP packet loss, RTT data were tested in the study.

Measurement Lab (M-Lab) is an open, distributed server platform for researchers to deploy Internet measurement tools so that the Internet users can test their connection using the tools which include Network Diagnostic Test (NDT), Glasnost, Pathload2, ShaperProbe, etc [4]. NDT is one of the major tools on the M-Lab which instruments the operating system of a server to capture detailed connection statistics during the measurement. In an NDT test, the client exchanges as much data as possible with a server to measure the upload and download speed between the two machines.

Huang and his colleagues measured and analyzed 3G network performance using a tool called 3GTest which is

now evolved to MobiPerf [9]. The tool collected TCP throughput, RTT, DNS lookup, HTTP request to landmark servers, etc. and used by more than 30,000 users from all over the world on several devices such as iPhone, Android G2, and Windows Mobile phones.

V. CONCLUSIONS

In the paper, we presented our software to conduct carrier-independent performance measurement of mobile broadband data services in California. We have developed the testing software for the Android-based phone and Windows-based laptop, and conducted our first field tests at 1,200 different locations throughout the State of California. Testing was done in the following sequence: simultaneous Android smartphone tests, establish a valid GPS signal using a GPS receiver on the laptop, and finally running a test on the laptop for each datacard. At a subset of locations, testers ran a Glasnost test on the laptop for each provider to measure the presence of carrier rate limiting or traffic shaping.

In the future, we will continue field testing every six months until 2014. Since we will be sampling every six months, we should be able to have precise trends both in coverage and performance for California mobile wireless broadband. We expect this trend information to be of great value to policy makers. Furthermore, we will make our tools available to the public so that they can download them and measure the performance and quality of their wireless broadband networks. We believe that our open-source nature of the project can serve as a model for other states to replicate. People will also be able to view the measurement result from the California broadband availability database.

ACKNOWLEDGMENT

This project is funded by a State Broadband Initiative Grant issued by the National Telecommunications and Information Administration to the California Public Utilities Commission pursuant to the American Reinvestment and Recovery Act.

REFERENCES

- [1] NLANR/DAST, iPerf, <http://sourceforge.net/projects/iperf/>
- [2] M. Dischinger, M. Marcon, S. Guha, K. P. Gummadi, R. Mahajan, S. Saroiu, "Glasnost: Enabling End Users to Detect Traffic Differentiation," Proc. USENIX Symposium on Networked Systems Design and Implementation (NSDI), San Jose, USA, April, 2010.
- [3] R. Carlson, "Developing the Web100 Based Network Diagnostic Tool (NDT)," Proc. Passive and Active Measurement (PAM), 2003.
- [4] Measurement Lab (M-Lab) <http://www.measurementlab.net/>
- [5] Mobile Broadband Drive-Test Blog <http://calbroadbanddrivetest.blogspot.com/>
- [6] Spring 2012 Mobile Broadband Field Testing Initial Staff Report, <http://broadbandcouncil.ca.gov/WorkArea/DownloadAsset.aspx?id=110>
- [7] S. Bauer, D. Clark, and W. Lehr, "Understanding broadband speed measurements," Proc. TPRC 2010: Telecommunications Policy Research Conference, October, 2010
- [8] Federal Communications Commission, Measuring Broadband America – A report on consumer wireline broadband performance in the U.S., <http://www.fcc.gov/measuring-broadband-america>
- [9] J. Huang, Q. Xu, B. Tiwana, Z. M. Mao, M. Zhang, and P. Bahl, "Anatomizing application performance differences on smartphones," Proc. 8th international conference on Mobile systems, applications, and services (MobiSys '10), San Francisco, USA, June 2010