2G/3G Network Measurements in Rural Areas of India

Amitsingh Chandele, IIT Delhi
Zahir Koradia, IIT Bombay
Vinay Ribeiro, IIT Delhi
Aaditeshwar Seth, IIT Delhi
Sipat Triukose, NICTA Australia
Sebastien Ardon, NICTA Australia
Anirban Mahanti, NICTA Australia

1. INTRODUCTION

Recent years have seen rapid cellular expansion in urban and rural India [1], providing an avenue to bridge the digital divide. However there is little understanding of the performance of cellular data connectivity in different geographies. We take a first step towards this. We are planning to characterize the performance of cellular data networks available across different locations, rural and urban, in India through a large scale experimental setup consisting of more than 50 measurement points across the country. We hope our findings will reveal capacity provisioning and network design characteristics that telecom operators follow in deploying 2G/3G connectivity in different areas.

Here we present results from several lab experiments and two field experiments: one rural (15 days in Dholpur, Rajasthan) and one urban (7 days in New Delhi). In the field experiments, we conducted basic uplink and downlink throughput and latency tests once every two hours. The results of these experiments have provided us early insights into the performance of Airtel GPRS, Airtel HSDPA and MTNL HSDPA in India. Some of the interesting observations are:

- Airtel and MTNL connections were unavailable during night hours at the rural site.
- Server side optimization techniques such as TCP Segmentation Offloading (TSO), Generic Receive side Offloading (GRO), and delayed acknowledgements, all commonly turned on by default, increase the burstiness of TCP resulting in packet loss.
- Large (> 2 minutes) stall durations were noticed in TCP flows, likely due to buffering in middleboxes.

We next present these results in greater detail.

2. RESULTS

2.1 Availability

During our field test at the rural site we noticed a strange diurnal pattern: both Airtel and MTNL connections got disconnected in the evenings and reconnected in the mornings. This can be seen in Figure 2.1 of TCP throughputs obtained in the rural area on Airtel GPRS connection over one week. Base stations in rural regions often operate on alternate sources of power (diesel/solar), and it is likely that the base stations at Dholpur may have been turned off at night to save energy. Figure 2.1 also shows instances where the devices were disconnected even during the day, indicating that the availability of the connections is low in rural areas.

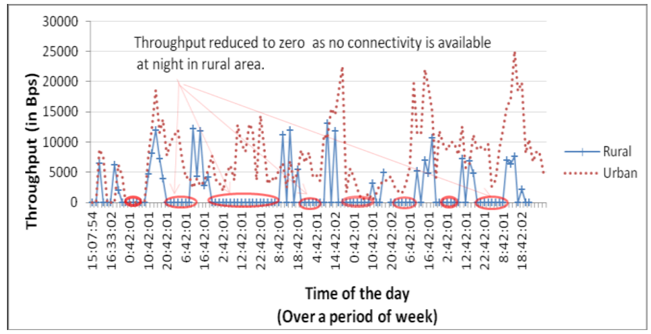

Figure 2.1: Availability of Airtel GPRS connection.

2.2 TCP specific parameters

We also found through analysis of TCP packet traces that several server optimization techniques made the flows bursty, causing packet losses. We present some of these observations below.

TCP Segmentation Offloading (TSO)

When TSO is enabled, a TCP sender transfers a segment much larger than the MSS to the network interface card (NIC), and leaves it to the NIC to segment it further into MSS sized segments. This is intended to reduce the CPU utilization at the server, by transferring the segmentation job to the NIC controller.

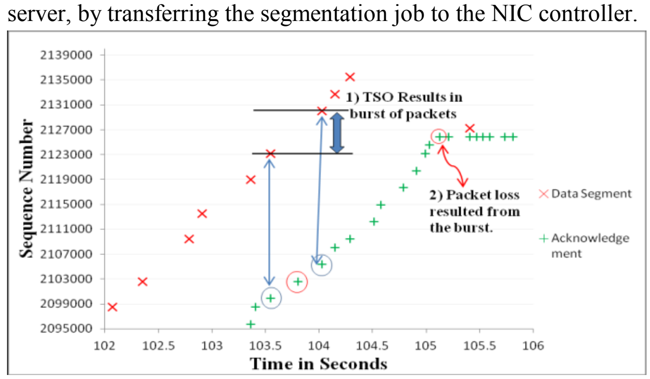

Figure 2.2.1: TSO causing burstiness in TCP over Airtel HSDPA connection in downlink direction.

Enabling TSO however means that the sender will wait for the congestion window to open enough before a sizable segment can be sent to the NIC, and all smaller segments are then dispatched by the NIC in a burst. Figure 2.2.1 shows how TSO causes a burst of packets to be sent, which in-turn results in packet loss. Our lab experiments showed that disabling TSO caused a 20% reduction in the number of retransmissions for low bandwidth Airtel GPRS connections.

Generic Receive side Offload (GRO)

GRO is a server side optimization technique which allows a NIC to receive multiple packets from the same flow, assemble them

DEV '13, January 11-12, 2013 Bangalore India

into a single large segment, and then send the segment to the higher layers. This again is intended to reduce CPU utilization. However, when GRO is used to combine multiple acknowledgements of a downlink TCP flow, the server's congestion window increment is very large and causes a burst of packets to be sent. Figure 2.2.2 shows how GRO causes burstiness.

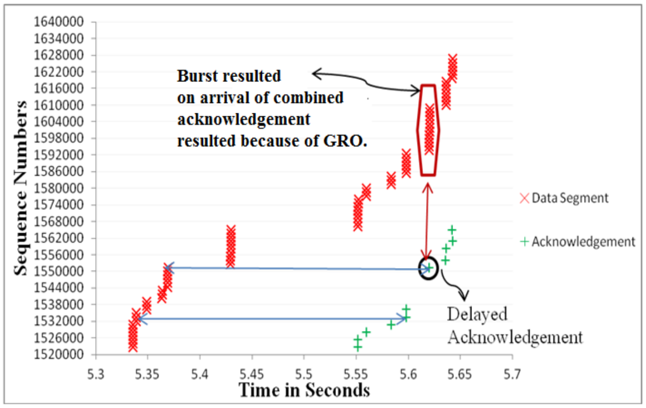

Figure 2.2.2: Combined acknowledgement causing burst in flow over Airtel HSDPA connection in downlink direction.

2.3 Long flow breaks

Our analysis of packet traces also showed long periods of inactivity, where the sender sent several packets that never arrived at the receiver. We call these periods of inactivity, *long flow breaks*. Figure 2.3 shows one such trace with several flow breaks.

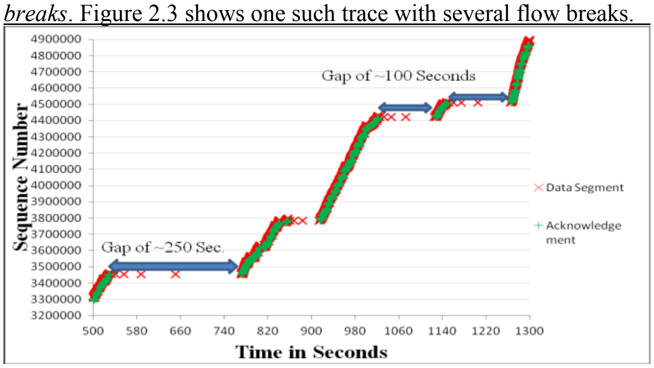

Figure 2.3: Long flow breaks in Airtel GPRS connection in downlink direction.

Across all our urban field measurements an Airtel GPRS connection, suffering from long flow breaks, spent an average of 24.8 seconds in flow breaks out of 118 seconds. Similarly, Reliance EVDO connections spent an average of 23.44 seconds in flow breaks and MTNL HSDPA connections spent 17.37 seconds in flow breaks. One possible explanation for these gaps could be attributed to the impact of middle boxes and firewalls in the network. A combination of firewalls buffering out of order packets [2] coupled with the forward RTO extension in TCP, could be creating this problem. We are further investigating the cause of these stalls.

3. RELATED WORK

Optimization of TCP for cellular networks has been well studied over the past 15 years, including appropriate values of MSS, initial window size, and TCP SACK and timestamp options [3]. Manipulations of sending window [4] and congestion control and loss recovery algorithms [5] have also been explored. However, we are the first to observe impact of server side optimizations on cellular data connections.

The behavior of TCP in environments where the per-flow throughput is less than one packet per RTT, referred to as sub-packet regime, which is often the case in campus networks with limited resources in developing regions has also been studied [6]. TCP flows in our target networks do not operate in sub-packet regimes, but a large number of flows from the same client using a GPRS connection could lead to such a degenerate condition. We intend to study the interaction of multiple simultaneous flows with each other in the future, especially for access to Facebook and Youtube like websites which are known to open multiple connections.

Recent measurement studies in cellular networks have measured latencies [7] [8], and throughputs and capacities [9] of cellular deployments in the developed world. These studies have found (a) improved latencies in HSDPA networks compared to WCDMA networks, (b) dependence of latency on service provider configurations, (c) throughputs roughly correlated to advertised bandwidth, and (d) wide variations in cell capacities across service providers as well as geographical locations. However, no such study exists for India. To the best of our knowledge we are the first to attempt understanding cellular networks in India.

4. ACKNOWLEDGEMENT

This work was funded as part of a Ford Foundation grant to design and deploy an Internet measurement framework for cellular data connectivity in India, a grant by the Department of Information Technology India to design and deploy a content distribution network for rural areas in India, and a grant by the Department of Science and Technology for Indo-Australia collaboration for innovations in rural Internet connectivity.

5. REFERENCES

[1] Latest survey by Juxtconsult. http://www.juxtconsult.com

[2] Zhaoguang Wang, Zhiyun Qian, Qiang Xu, Z. Morley Mao, Ming Zhang, "An Untold Story of Middleboxes in Cellular Networks", SIGCOMM'11, August 15–19, 2011.

[3] P. Benko, "A large-scale, passive analysis of end-to-end TCP performance over GPRS", IEEE INFOCOM 2004.

[4] R. Chakravorty, J. Cartwright, I. Pratt, "Practical Experience with TCP over GPRS", IEEE GLOBECOM 2002.

[5] T. Goff, J. lvloronski, D.S. Phatak, V. Gupta, "FreezeTCP: A true end-to-end TCP enhancement mechanism for mobile environments". IEEE INFOCOM 2000.

[6] Jay Chen, Janardhan Iyengar, Lakshminarayanan Subramanian, Bryan Ford, "TCP Behavior in Sub-Packet Regimes", SIGMETRICS'11, June 7–11, 2011.

[7] M. Jurvansuu, J. Prokkola, M. Hanski, P. Perala, "HSDPA Performance in Live Networks", ICC 2007.

[8] A. Elmokashfi, A. Kvalbein, J. Xiang, K. Evensen, "Characterizing Delays in Norwegian 3G Networks", PAM12.

[9] W. L. Tan, F. Lam, W. C. Lau, "An Empirical Study on the Capacity and Performance of 3G Networks", IEEE Transactions on Mobile Computing, June 2008.