Understanding RRC State Dynamics through Client Measurements with Mobilyzer

Sanae Rosen, Haokun Luo,
Qi Alfred Chen, Z. Morley Mao
University of Michigan
{sanae, haokun, alfchen, zmao} @umich.edu

Jie Hui, Aaron Drake, Kevin Lau
T-Mobile USA Inc.*
{Jie.Hui, Aaron.Drake, Kevin.Lau}
@t-mobile.com

ABSTRACT

Understanding how network and application behavior patterns impact client performance on mobile devices is a difficult yet important problem to solve. Often, we are most interested in the performance experienced by end users, but accurately and effectively measuring performance on uncontrolled mobile devices in the wild continues to be a challenging problem. In this paper, we have developed a tool to allow us to more accurately characterize RRC states using client-based measurements. To do so, we have made use of and contributed to an open-source framework called Mobilyzer, which facilitates collecting complex network measurements on client devices without impacting the user. We have deployed our tool on unmodified devices in 23 countries worldwide to directly measure the impact of RRC state machine configurations on individual packets and network protocols. Demonstrating the value of client-based measurements, our large-scale measurement study allowed us to uncover previously unknown performance problems that can increase network latency by several seconds and increase packet losses by an order of magnitude.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: wireless communication; C.4 [**Performance of Systems**]: measurement techniques, performance attributes

Keywords

4G LTE; 3G; smartphones; RRC state machine; cellular network performance

1. INTRODUCTION

As described in previous work [7, 8, 10, 9, 13], Radio Resource Control (RRC) states [4, 5] have a substantial impact on network performance. To balance the need for high performance with the need to conserve energy on resource-constrained mobile devices, these devices transition between states with different performance and energy characteristics. Generally, devices must be in a high-power state to send data, and there is additional latency (a promotion delay) to transition to these states. Devices remain in these high power states for several seconds in case there is more network traffic, before demoting to a lower power state. These timers are set by carriers. Longer timers lead to better network performance, at the cost of a shorter battery life. We summarize the RRC state machine for two network technologies (3G and 4G LTE) in Figure 1.

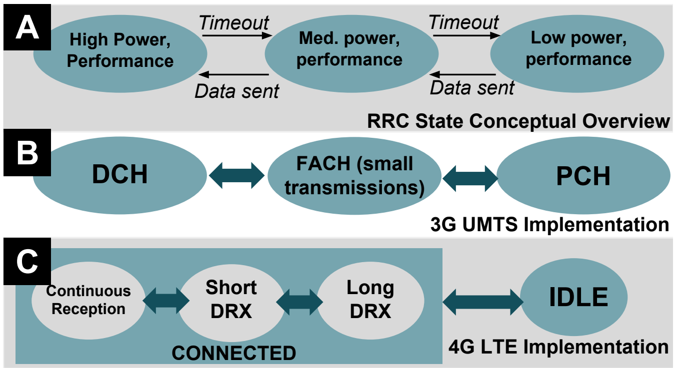

Figure 1: A: Overview of RRC state machine design. B-C: possible 3G and 4G state machines.

In this paper, we focus on describing our client-based measurement approach and our findings from crowdsourcing client-based measurements. In the full version of this paper [12], we investigate fine-grained RRC dynamics and performance in more depth, including through other measurement approaches. Prior work has focused on measuring RRC state configurations in a controlled setting or on specific device types which make RRC transition event logs directly available to applications [8, 10, 9, 13]. Collecting measurements on unmodified client devices, however, allows us to crowdsource the measurement of RRC state timers and performance worldwide. We carry out measurements in 23 countries and on a wide range of device types, technologies and carriers. Through these measurements, we uncovered previously unknown performance problems that can significantly impact users. Our application continues to measure RRC state performance characteristics over time as networks evolve.

Collecting and crowdsourcing measurements on user mobile devices has proven to be an effective technique for understanding mobile networks [6, 2]. However, performing complex network measurements accurately and effectively on unmodified devices,

---

[1] The views presented here are as individuals and do not necessarily reflect any position of T-Mobile.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
S3'14, September 7, 2014, Maui, Hawaii, USA.
Copyright 2014 ACM 978-1-4503-3073-2/14/09 ...$15.00.
http://dx.doi.org/10.1145/2645884.2651821.

in the presence of interfering traffic or noisy network conditions, and without inconveniencing users, remains challenging. To ensure that our RRC measurement techniques remain relevant to the broader community, we have added them to an open-source tool, MobiPerf [1], as well as the underlying Mobilyzer library that MobiPerf uses [11]. Mobilyzer's goals are to facilitate and systematize the development of mobile measurement apps, allowing researchers and other interested parties to easily measure and monitor complex network phenomena while restricting energy and data consumption to ensure these applications meet the needs of users. In addition to describing the results of our RRC measurement tool, we also describe some system challenges in designing the tool, which are likely relevant to crowdsourcing mobile application measurements more generally.

2. CLIENT RRC MEASUREMENTS

To measure RRC states, we start by sending a large UDP packet to ensure the device is in a high power state. After waiting for a period of time, ensuring that no traffic is sent on the device, another UDP packet is sent to an echo server, and the round-trip time is measured. If the demotion timer was shorter than the waiting period, there will be an additional latency delay due to the RRC state demotion. By varying the inter-packet interval and measuring the resulting latency, it is possible to infer RRC state timers. This basic technique was introduced in work by Feng et al [7].

To allow this technique to function in an uncontrolled environment on user Android devices, we detect interfering traffic on the device using the Linux */proc/net* utility, which requires no special privileges to see global packet counts. We discard individual tests where there were more packets than expected. Furthermore, performance on real networks often varies regardless of RRC states. We filter out all devices with less than 5 complete sets of measurements, and then observe changes in the distribution of latencies in our measurements as the time interval is varied. Finally, we also use this technique to measure higher layer protocols such as HTTP requests.

An example of the data collected from one device is shown in Figure 2. For interpacket intervals below 2 seconds, the latency is low for both large and small packets, and from about 4.5 to 9 second intervals, the round-trip time is higher for the larger packet, a characteristic of the intermediate FACH state in 3G. After 10 seconds, the round-trip times for both packets are high. Note that during RRC state demotions, especially for the demotion to FACH, delays are very high. Our large-scale measurement study has determined that a significant number of carriers worldwide exhibit this previously unknown behavior, which we call a *demotion delay*. Note also that the measured round-trip times vary greatly, necessitating a large number of measurements to uncover the underlying pattern.

2.1 Mobilyzer Support

In addition to the challenges in collecting accurate measurements on client devices, there are also systems challenges in performing these measurements efficiently on a large scale. We have added the RRC measurement tool to MobiPerf, an open-source Android app that allows users to measure their network performance and track how performance metrics change over time. Data is collected in aggregate and anonymized on a central server, stripping all user-identifiable data.This tool runs on the Mobilyzer framework, which we have modified as well.

Collecting RRC measurements is not a lightweight task. Comprehensively measuring a large range of interpacket intervals is time-consuming, and interfering traffic on the device causes individual measurements to be discarded and rescheduled. In practice, a full set of tests takes roughly half an hour to complete. Additionally, the number of requests made to test a full range of inter-packet intervals and higher layer protocols can consume hundreds of kilobytes per set of tests. Carefully monitoring and budgeting data consumption over the course of a month is therefore important. We run the RRC task every 36 hours, and reduce the frequency if needed to meet data and power constraints set by the user, dynamically keeping track of data consumption. Measurements are run only when the battery level exceeds a threshold set by the user.

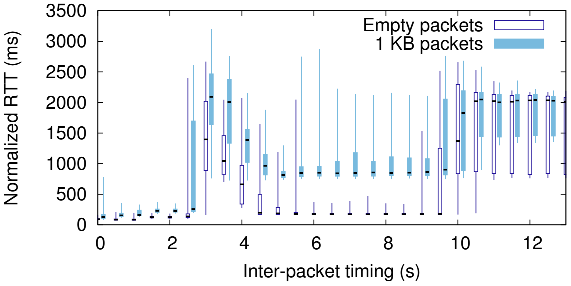

Figure 2: Example of client-based RRC latency measurements while transitioning through RRC states. Median, quartile and 5%/95% values shown.

Additionally, no other measurement tasks should send traffic at the same time as the RRC test. We modified Mobilyzer to support long-term scheduling of background tasks, with a schedule that persists even if the application is forcibly closed. Furthermore, Mobilyzer supports pausing the task (for instance, if a higher priority measurement task is scheduled, or the device switches to WiFi) to allow a test to be completed without losing data. Tasks are scheduled using the global Mobilyzer scheduler, which coordinates across different applications running the library so that different research measurement apps do not interfere.

We also adjust the default schedule (subject to the user's data consumption constraints) from a central coordination server. In this way, once we have collected an initial dataset of RRC performance data, we can reduce the frequency of these measurements, allowing us to continue to monitor changes over time with minimal impact to the user. As data is collected by volunteers, it is essential that the application remain unobtrusive and respect their battery and data consumption limits.

3. RESULTS

We discovered that packets sent during a state demotion can often experience round-trip times several times longer than normal, or in LTE, experience significant packet losses. This appears to be due to a variety of different state demotion implementation problems in different carriers and technologies. Our full dataset covers 44 carriers, but to briefly illustrate our findings, we discuss three example carriers and illustrate the impact that RRC states and RRC state demotion performance problems have on DNS lookups and HTTP requests performed on these networks. These results are shown in Figure 3. To assign each individual test to a RRC state, we first identified the RRC states using the method discussed above, as well as the transition period between each state. For these three carriers, we confirmed these timers using QxDM [3], a tool that allows us to directly read RRC state transitions.

For carrier C1, there are three distinct RRC states in 3G: DCH, a high-power state; FACH, an intermediate state where only a

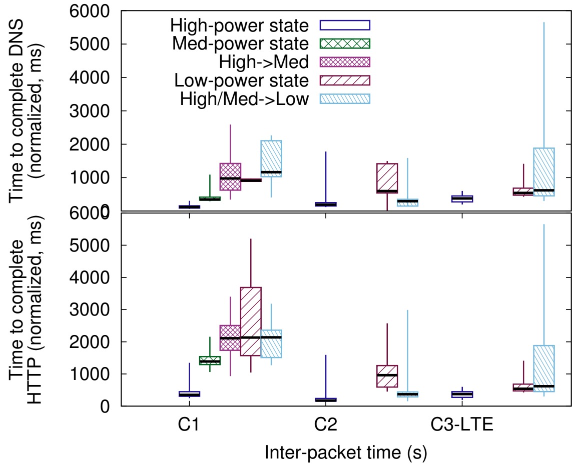

Figure 3: Example of how RRC state transitions affect DNS and HTTP performance for three sample carriers.

small amount of data can be sent; and PCH, a low-power state. DNS requests are generally small enough they can be completed in FACH without needing to promote to DCH, and so the average delay in FACH is significantly less than in PCH. HTTP requests, being larger, often require a transition to DCH to complete and so do not benefit as much from FACH. However, in both cases, packets sent during the transition from DCH to FACH (shown as a solid purple box) can result in very high delays. This is especially pronounced for the smaller DNS lookup requests.

In contrast, for carrier C2, there is no FACH state. Furthermore, requests sent during state demotions do not have unexpected delays. As we discuss in the full paper, this is due to implementation differences between the two carriers.

Finally, for carrier C3, we examine LTE. LTE's lower state promotion overhead means that overhead of sending a packet in IDLE is minimal. However, delays for packets sent during the demotion process experience high delays and loss rates an order of magnitude higher than normal. This demotion delay problem occurs for fewer carriers for LTE than 3G, however. Furthermore, the demotion period is shorter and so packets are less likely to experience these delays.

3.1 Discussion and Future Work

In this paper, we have discussed the challenges and benefits of using client measurements to better understand the impact of RRC states on performance. In addition to the novel results we have found (discussed in more depth in our full-length paper [12]) this approach allows us to continue to maintain a database of cellular network performance worldwide that reflects the evolving cellular network. We have also contributed to the Mobilyzer project [11] aimed at developing a large, comprehensive dataset of network measurements worldwide from the client perspective.

Our previously undiscovered findings on RRC demotion delays motivates the importance of client-based measurements. A crowdsourcing approach allows for large-scale, global deployments, unconstrained by the geographic location of the researchers involved. The resulting measurements can give a great deal of insight into how complex network phenomena affect users in the wild, ensuring that results are representative of real user experiences, and allows a wide range of devices, carriers, network types and network performance characteristics to be investigated. Furthermore, we have demonstrated that a highly complex measurement task that is sensitive to interference and environmental conditions can be automated and run on unmodified, actively-used mobile devices. In addition to improving our understanding of the real-world impact of RRC states, we argue for the value of future work in crowdsourcing measurements of cellular networks worldwide.

4. ACKNOWLEDGEMENTS

We would like to thank Ashkan Nikravesh, Hongyi Yao and Dave Choffnes for their assistance with Mobilyzer and their work on developing the Mobilyzer framework. We would also like to thank Kranthi Sontineni, Randy Meyerson, Warren McNeel, and the Product Experience and QoE Lab team at T-Mobile for their assistance. This research was supported in part by the National Science Foundation under grants CNS-1059372, CNS-1039657, CNS-1345226, and CNS-0964545, as well as by an NSERC Canada PGS M scholarship.

5. REFERENCES

[1] MobiPerf. `http://mobiperf.com`.

[2] NU JamLogger: A Study of User Activity and System Performance on Mobile Architectures. `http://www.ece.northwestern.edu/microarchitecture/jamlogger/`, 2009.

[3] QxDM Professional Proven Diagnostic Tool for Evaluating Handset and Network Performance. `http://www.qualcomm.com/media/documents/files/qxdm-professional-qualcomm-extensible-diagnostic-monitor.pdf`, 2012.

[4] 3GPP TS 35.331: Radio Resource Control (RRC) - UMTS, 2013.

[5] 3GPP TS 36.331: Radio Resource Control (RRC) - LTE, 2013.

[6] C. C. Tossell, P. Kortum, A. Rahmati, C. Shepard, and L. Zhong. Characterizing web use on smartphones. In *ACM CHI*, 2012.

[7] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. Characterizing Radio Resource Allocation for 3G Networks. In *Proc. ACM IMC*, 2010.

[8] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. A Close Examination of Perfomance and Power Characteristics of 4G LTE Networks. In *Proc. ACM MobiSys*, 2012.

[9] J. Wigard, T. Kolding, L. Dalsgaard, and C. Coletti. On the User Performance of LTE UE Power Savings Schemes with Discontinuous Reception in LTE. In *Proc. International Conference on Communications*, 2009.

[10] L. Zhou, H. Xu, H. Tian, Y. Gao, L. Du, and L. Chen. Performance Analysis of Power Saving Mechanism with Adjustable DRX Cycles in 3GPP LTE. In *IEEE VTC*, 2008.

[11] A. Nikravesh, H. Yao, S. Rosen, Y. Jia, D. Choffnes, and Z. M. Mao. Mobilyzer: Mobile network measurements made easy. `http://mobilyzer-project.mobi/about.html`.

[12] S. Rosen, H. Luo, Q. A. Chen, Z. M. Mao, J. Hui, A. Drake, and K. Lau. Discovering Fine-grained RRC State Dynamics and Performance Impacts in Cellular Networks. In *Proc. ACM MobiCom*, 2014.

[13] N. Vallina-Rodriguez, A. Auçinas, M. Almeida, Y. Grunenberger, K. Papagiannaki, and J. Crowcroft. Rilanalyzer: A comprehensive 3g monitor on your phone. In *Proc. ACM IMC*, 2013.