

Impact of Acknowledgments on Application Performance in 4G LTE Networks

Brett Levasseur
MIT Lincoln Laboratory
244 Wood St
Lexington, MA
brettl@ll.mit.edu

Mark Claypool
Worcester Polytechnic Institute
100 Institute Rd
Worcester, MA
claypool@cs.wpi.edu

Robert Kinicki
Worcester Polytechnic Institute
100 Institute Rd
Worcester, MA
rek@cs.wpi.edu

ABSTRACT

4G LTE, the new cellular network standard, provides the capacity to support Quality of Service (QoS) for wireless multimedia applications. Recent developments have modified the LTE QoS setup with MAC layer schedulers and changed its current architecture. However, research has yet to examine the effects of the various LTE retransmission choices on QoS. This paper examines the impact of using acknowledgments and timer adjustments to recover lost data over LTE on VoIP and FTP applications. NS-3 simulations show that LTE retransmissions improve FTP throughput, while the benefit to VoIP applications varies with wireless loss rate.

1. INTRODUCTION

Mobile technology has increased both in capability and Internet traffic penetration. Cisco reports over 500 million mobile devices and connections were added in 2013 and 77% of this growth came from smartphones [3]. Global mobile data traffic grew 81% and mobile video traffic was 53% of all mobile data traffic in 2013. In addition, 4G connections generated 14.5 times more traffic on average than non-4G connections. While accounting for only 2.9% of all mobile connections, 4G generated 30% of the mobile traffic and is predicted to handle half of all mobile traffic by 2018.

A common challenge for networks is managing loss and minimizing response times. This is problematic for wireless networks when the communicating nodes are mobile and receiving signals vary in signal quality due to interference from obstructions and radio wave emitting sources. Wireless networks address this variability in signal quality with modulation and encoding redundancy and retransmissions.

4G Long Term Evolution (LTE) provides retransmissions at two different network layers with several parameters that control its retransmission strategies. These parameter settings impact the effect of LTE retransmission schemes on the quality of service experienced by wireless end users' applications. Selecting the best parameter settings is challenging

This work is sponsored by the Department of the Air Force under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and not necessarily endorsed by the United States Government.

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.




Figure 1: Evolved Packet Core

because applications such as email or file transfer need error-free transmissions while other applications such as Voice over IP (VoIP) require minimal network delay and can accept some loss. Additionally, LTE timers interact with these parameters to determine if wireless data is lost and when to trigger retransmissions. However, to the best of our knowledge there has been no systematic exploration of the effects of LTE retransmissions and timer settings on application performance.

This work examines the effect of 4G LTE configurations for retransmission of last hop data on delay sensitive applications (e.g., VoIP) and throughput sensitive applications (e.g., file transfer). We enhanced the NS-3 simulator to support the use of negative acknowledgments (NACKs) in LTE acknowledged mode (AM), part of the LTE specification that was not previously implemented in this simulator. Detailed simulations with varied wireless loss rates demonstrate the sensitivity of application performance to LTE timer settings and guide recommendations for when LTE should use acknowledged mode versus unacknowledged mode for delay sensitive and throughput sensitive applications.

This paper is organized as follows: Section 2 provides 4G LTE background; Section 3 describes our NS-3 LTE experiments; Section 4 analyzes the results; and Section 5 summarizes our conclusions and mentions possible future work.

2. BACKGROUND

An LTE network has the two components depicted in **Figure 1**. The Evolved Packet Core (EPC) connects a wireless access point or eNodeB to an IP network and the Evolved Universal Mobile Telecommunications System Terrestrial Radio Access (E-UTRA) connects the eNodeB to phones, tablets and computers (User Equipment or UEs). The E-UTRA interface includes two LTE link layers between the physical and IP layer – the Media Access Control

(MAC) and Radio Link Control (RLC) layers. Sitting above the physical layer, the MAC layer handles scheduling, notification of transmission opportunities and retransmissions. Located above the MAC layer, the RLC handles out of order arrival, error correction and, optionally, retransmission of data not recovered by the MAC layer.

2.1 Hybrid Automatic Repeat Request (HARQ)

The LTE transport block is the data packet sent across the wireless link [1]. The MAC handles retransmission of transport blocks with Hybrid Automatic Repeat Request (HARQ). When a transport block arrives, the MAC uses forward error correction (FEC) and parity bits to check for errors. If it detects no errors, the MAC receiver sends an ACK to the MAC sender and passes the data up to the higher layers. If the MAC receiver detects an error, it sends a NACK and keeps the transport block to possibly combine with multiple retransmitted copies of the damaged data to recreate the original data through a soft combination process. For downlink transmissions (from the eNodeB to the UE), HARQ sends up to three retransmissions, although some research suggests that a lower maximum would be better under poor wireless conditions [7]. If any MAC packets arrive without error or if the MAC soft combine produces a valid packet, HARQ sends the packet up to the RLC layer. If, after three retransmissions, HARQ cannot reproduce the transport block through soft combining, the data is treated as lost.

2.2 Radio Link Control (RLC) Acknowledged/Unacknowledged Modes

The RLC layer provides segmentation and reassembly between the original data frames and those encoded in the transport blocks [1]. The RLC retransmission scheme employs two transmission modes – acknowledged mode (AM) and unacknowledged mode (UM). Both RLC modes wait a fixed time interval to correct out of order data (due to MAC layer HARQ retransmissions). UM RLC only waits for reordering before sending current data to higher layers, whereas AM RLC sends ACKs and NACKs for data not recovered by the MAC layer.

To support these transmission modes, the RLC uses timers and state variables [1]. Both AM and UM modes use the *t-Reordering* timer to control the RLC wait interval on out of order MAC data before either: 1) considering the data lost and handing any received data off to the next network layer (UM), or 2) updating which RLC layer packets to ACK or NACK (AM). This timer gives the MAC's HARQ process a chance to recover the lost data. If *t-Reordering* is set too low, data the MAC layer could still recover may be discarded as lost (UM) or NACKed prematurely for retransmission (AM). If *t-Reordering* is set too high, received data may be held unnecessarily long before requesting a retransmission (AM) or delivering it to upper layers (UM).

AM uses an additional *t-StatusProhibit* timer to control status messages [1]. STATUS is an RLC control message that preempts user data messages. An AM STATUS message causes ACKs and NACKs to be exchanged between the eNodeB and the UE. STATUS messages are either polled by the sender or triggered on the receiver during certain events. AM mode sets *t-StatusProhibit* after sending a STATUS message to prevent a node from sending another STATUS message while the timer is running. If *t-StatusProhibit*




Figure 2: 4G LTE simulation network

is set too low, duplicate STATUS messages may clog the transmission medium. If *t-StatusProhibit* is set too high, the RLC sender may continue transmitting new data when old data needs to be retransmitted. Research from [9] indicated that for normal Web browsing, other selective repeat mechanisms perform better than AM, but AM performs better for FTP. This research did not however consider the MAC layer's HARQ process.

3. APPROACH

This research uses the NS-3 simulator¹ to examine the impact of configurable RLC parameter adjustments on performance of applications running over 4G LTE networks. Specifically, this work considers two mobile applications that differ in QoS criteria and characteristics: Voice over IP (VoIP) and FTP. VoIP applications send small packets at a relatively low rates and need low packet delays and loss rates. In contrast, FTP applications send large packets at a relatively high throughput and need zero loss but can tolerate some delay. A more in-depth analysis of these experiments and additional experiments with MPEG video are available in [8].

Since the LTE module² for NS-3 version 3.16 did not support NACKs in AM, we joined an existing discussion on this issue at the NS-3 user forum³, and, over the course of several months, added NACK support to the simulator.

All our simulations use the network topology in **Figure 2** with a server node sending data over 4G LTE to a wireless UE. **Table 1** itemizes the Internet and LTE settings for our NS-3 simulations.

NS-3 has a built-in trace fading model which takes as input a file with a matrix of time, a section of the radio spectrum, and the Signal to Noise Ratio (SINR). The model applies the trace SINR at the specified time and section of the radio spectrum by adjusting the base SINR of any data transmitted at a given time and location. We use a two-state model to induce wireless traffic loss into the LTE simulations – either the fading trace makes no adjustments or the SINR is set low enough that any data sent at that time and position is lost. Using a Gilbert-Elliot [5] model, we define two LTE states: 1) good where no packets are lost, and 2) bad where all packets are lost. Our implementations employ a set of equations from [6] to find settings for the transition state variables in the Gilbert-Elliot model based on a desired input of lost transmission opportunities (where any data sent

¹<http://www.nsnam.org/>

²<http://networks.cttc.es/mobile-networks/software-tools/lena/>

³<https://groups.google.com/forum/#!topic/ns-3-users/CEfmMX3IRBw>

Table 1: NS-3 Settings

Internet	
Bandwidth E-UTRA to server	100 Gbps
Delay E-UTRA to server	20 ms
MTU	1,500 bytes
LTE	
Distance eNodeB to UE	5 Km
EARFCN UL	100
EARFCN DL	18,100
# PRBs	100
Tower Tx Power	20 dBm
UE Tx Power	10 dBm
CQI (no loss)	8
RLC AM Buffer Size	Unlimited
RLC UM Buffer Size	9,999,999 Bytes
RLC AM maxRetxThreshold	5
pollPDU	1
pollBytes	50 bytes
tPollRetransmit	100 ms
tReordering (default)	40 ms
tStatusProhibit (default)	20 ms

at that time is lost) and input gap lengths. With the chosen transition state variables, the Gilbert-Elliot model generates the fading trace file for NS-3 to simulate bursty LTE traffic behavior.

While varying RLC settings, our investigation focuses on VoIP and FTP performance. During a series of NS-3 LTE simulations, we adjust the wireless loss rate, the use of RLC AM and UM and the settings of the timers *t-Reordering* and *t-StatusProhibit*. The suggested settings from the LTE specification [1] guide our timer settings.

The simulated UDP VoIP application sends CBR traffic at 64 Kbps to align with the G.711 encoding standard [11]. As VoIP conversations are sensitive to delay and packet loss, the simulator continually computes the mean opinion score (MOS) [2], a scale from 1 (bad) to 5 (good). While the overall MOS average over time may be high, cellular users react negatively when a VoIP conversation experience a low MOS at some point. Hence, this study uses the worst MOS score as its VoIP QoS metric with the simulator determining a single worst MOS value over a talkspurt (the average amount of time a person speaks continuously before pausing, 4.14 seconds based on [10]) for every possible time frame within the simulation.

Using TCP New Reno (the default congestion control algorithm in NS-3), the simulated FTP application transmits as much data as possible. Hence, we use throughput as FTP's QoS performance metric.

4. RESULTS

This section analyzes the results of three sets of 4G LTE simulation experiments. The first experiments employ uniform wireless loss to compare the performance of RLC using acknowledged mode (AM) against RLC with unacknowledged mode (UM) for VoIP and FTP. Using a bursty loss setup, the next set of tests vary the *t-Reordering* and *t-StatusProhibit* timers to evaluate the impact of these settings on the behavior of wireless VoIP and FTP applications. After selecting one fixed setting of timers for VoIP and another fixed setting for FTP, the final simulations vary

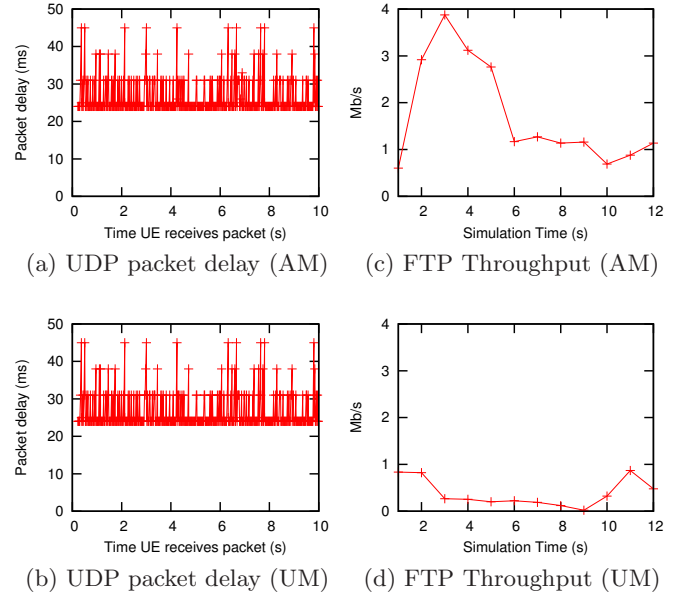


Figure 3: VoIP packet delay and FTP throughput

wireless loss rates to assess the affects of using RLC AM and UM for the two applications.

4.1 Basic RLC AM and UM

The first set of experiments study the impact of RLC using AM versus UM on VoIP and FTP. The wireless loss for these tests is set with a uniform loss probability of 25%. The graphs of UDP delay for each VoIP packet received in **Figures 3a** and **3b** show little difference between acknowledged mode (AM) an unacknowledged mode (UM). As VoIP puts only a small capacity demand on LTE, with an 25% uniform probability loss rate many loss events occur during intervals when the VoIP application is not transmitting. Additionally, the low VoIP CBR means RLC AM retransmissions have little impact on UDP packet delay when compared to UM delay results.

Figures 3c and **3d** clearly indicate that LTE using AM yields higher FTP throughputs than when LTE uses UM. By recovering lost encapsulated TCP packets via AM retransmissions, the RLC layer reduces the number of TCP packet losses. When using New Reno, the simulated TCP server reduces its sending rate after three duplicate acknowledgments. That is, the receiver sends back a total of four acknowledgments for the same TCP segment due to a TCP segment that has been missing for four receptions of other out of order segments. Thus, AM mode indirectly reduces the TCP duplicate ACK traffic while lowering the probability of the TCP server entering the fast retransmit state which lowers TCP's congestion window before retransmitting the TCP packet.

4.2 Adjusting RLC Timers

While the above results imply little difference between AM and UM in LTE for VoIP and that FTP prefers AM Mode, this is not the whole story. Adjusting timers *t-Reordering* and *t-StatusProhibit* affects the time RLC takes before considering data lost and requesting retransmissions. The next set of tests use the Gilbert-Elliot model to create bursty loss

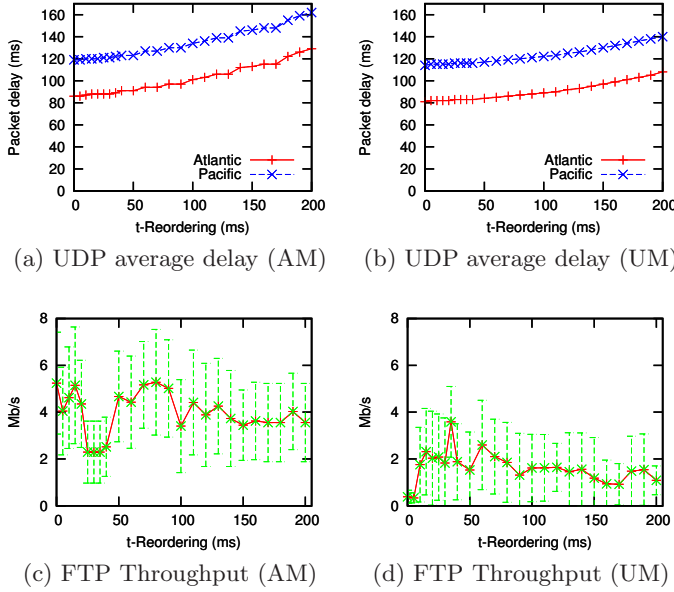


Figure 4: Adjusting t -Reordering for VoIP and FTP

with the overall loss rate set to 10%, the acceptable upper bound LTE uses when adjusting modulation and encoding schemes for transmission [4]. The VoIP test simulates a single VoIP application server communicating over a 4G link to the UE. A typical VoIP scenario would have another UE and two communication flows, but this simplified setup allows for a focus on LTE downlink performance. The simulated UDP end-to-end packet delay includes both the delay on the core network to reach the 4G network and the time to traverse the LTE network itself. To account for core network delays, we add the average delays reported for two of Verizon’s core networks (namely, 77 ms for its trans-Atlantic line and 110 ms for its trans-Pacific link) [12] to the LTE delays the simulator reports.

Figures 4a and 4b show the result of adjusting t -Reordering while t -StatusProhibit remains fixed at its default of 20 ms. Regardless of RLC mode, as t -Reordering increases, the average UDP VoIP packet delay increases. However, AM retransmissions cause extra wireless delay which yields higher UDP packet delay in Figure 4a than the UM delays seen in Figure 4b. Over both AM and UM tests for the Atlantic and Pacific core networks, the worst MOS scores varies little from 4.5 which corresponds to good user call quality. Hence, the strategy of setting t -Reordering to its lowest value seems attractive for providing optimal VoIP QoS. However, setting the timer too low stifles MAC layer recoveries. To avoid unnecessary lost MAC packets in UM and extra retransmissions in AM, t -Reordering must permit the MAC layer recovery process to complete (i.e., approximately 28 ms [4]). The closest recommended timer setting larger than this interval is 30 ms [1].

Figures 4c and 4d graph throughput over a series of FTP simulations that vary t -Reordering for both AM and UM. AM FTP throughput is higher than UM FTP throughput for all t -Reordering settings except for the 25-45 ms range when AM throughput unexplainably drops.

There is no one setting that produces the best results in

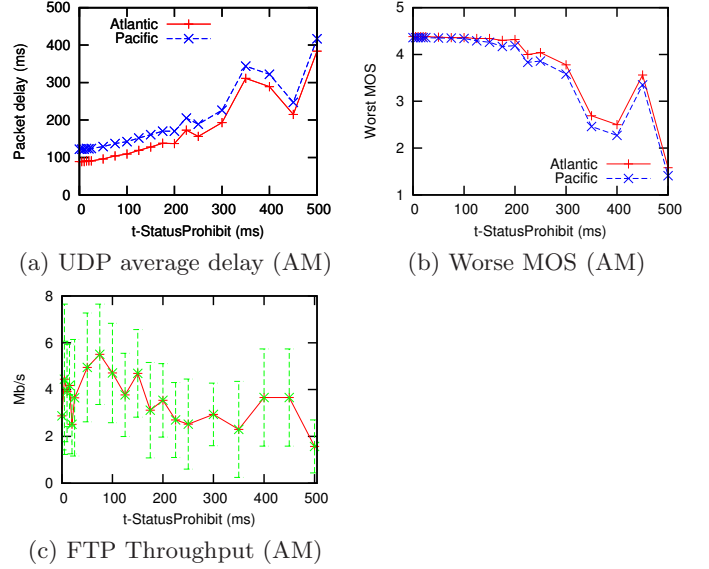


Figure 5: Adjusting t -StatusProhibit for VoIP and FTP

both AM and UM. Most of the best results for AM come with the timer set from 50 to 90 ms. Some of the best UM timer settings range from 15 to 60 ms. As our third set of experiments compare AM to UM when the timers are fixed, to split the higher AM and lower UM, we chose 50 ms for t -Reordering for both AM and UM.

The next series of experiments fix t -Reordering to 40 ms (the NS-3 default) and vary t -StatusProhibit. As t -StatusProhibit controls STATUS messages containing ACKs and NACKs, it only applies to AM. Figures 5a and 5b indicate that t -StatusProhibit has a greater impact on VoIP QoS than the t -Reordering timer. Generally, lowering t -StatusProhibit reduces UDP packet delay and increases the worst case MOS. However, there are higher settings where the delay and MOS are better. The 450 ms setting yields better QoS than at 400 ms and 500 ms settings, likely due to interaction between the two timers. When RLC enables t -StatusProhibit, the node cannot send STATUS messages but does update the set of packets to retransmit after t -Reordering expires. If t -StatusProhibit starts and then t -Reordering expires, any new packets that need to be NACKed have to wait until t -StatusProhibit expires. For example, if t -StatusProhibit is 400 ms and t -Reordering expires slightly later, almost 400 ms must pass before the NACK STATUS message is sent. However, if t -StatusProhibit is set to 450 ms, t -Reordering may expire when t -StatusProhibit is not running and a STATUS message can be sent earlier.

While lower t -StatusProhibit yields better VoIP performance for AM in terms of delay and MOS, the timer must not be set too low because lower settings increase STATUS message frequency. Since STATUS messages preempt user data, they reduce user uplink throughputs. While analysis of uplink traffic is outside the scope of this investigation, we limit impact by choosing 50 ms for t -StatusProhibit.

Similar to the negative impact on VoIP, setting t -StatusProhibit too high has a negative impact on TCP throughput. Figure 5c tracks FTP throughput while varying t -

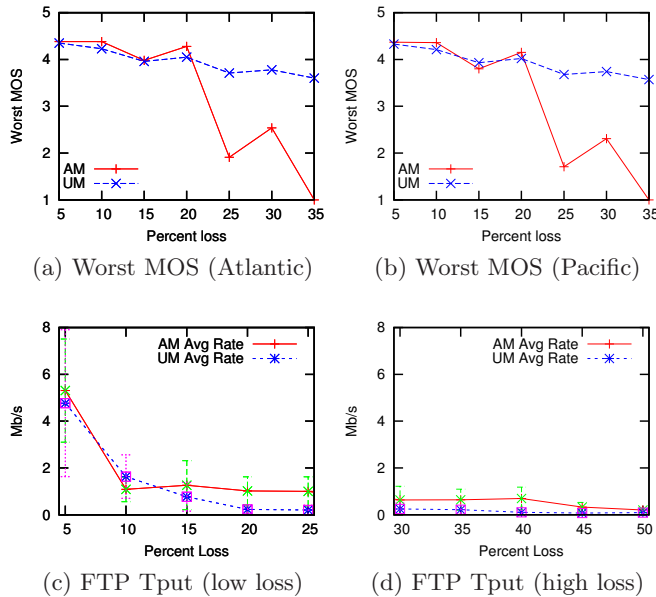


Figure 6: Fixed t -Reordering and t -StatusProhibit with different loss rates

StatusProhibit. The figure indicates the best throughput when t -StatusProhibit is set to 75 ms.

4.3 Fixed Timers and Varied Wireless Loss

This section presents VoIP and FTP simulations that use the bursty loss model with fixed timer settings based on the previous section's results.

The VoIP experiments fix t -Reordering and t -StatusProhibit to 30 ms and 50 ms respectively while varying the wireless loss rates from 5% to 35% in five percent increments. **Figures 6a** and **6b** provide worse MOS scores accounting for the Atlantic and Pacific core network delays, separately. With fixed timers, RLC AM slightly outperforms UM up to 20% loss. However, the differences are negligible as call quality at or near MOS 4 is considered good. At higher loss rates, the negative impact of many AM retransmissions on VoIP delays significantly outweighs more lost UDP packets when VoIP runs over UM.

The FTP tests run with t -Reordering and t -StatusProhibit set to 50 ms and 75 ms, respectively, with loss rates varied from 5 to 50% in five percent increments. As two distinct input sets are required for the equations that generate the fading trace files, we present the FTP throughput results as two separate figures – **Figures 6c** and **6d**, with 5 - 25% loss rates and 30 - 50% loss rates, respectively.

At an average loss rate of 5%, AM outperforms UM since MAC layer retransmissions can recover much of the lost data while reducing RLC AM retransmissions. A loss rate of 10% represents a crossover point where AM sends more retransmission to make up for lost data. However, above a 10% average loss rate, AM yields higher throughput than UM until very high rates at 45 and 50% where neither AM nor UM deals with the losses well.

5. CONCLUSION

This study examines the impact of the t -Reordering and

t -StatusProhibit timers along with the choice of RLC Acknowledged Mode (AM) versus Unacknowledged Mode (UM) on VoIP and FTP applications running over 4G LTE. In the NS-3 VoIP experiments with t -Reordering and t -StatusProhibit set to 30 ms and 50 ms, respectively, extra AM retransmissions improve call quality for up to a 20% packet loss rate on the wireless link. The FTP experiments with t -Reordering and t -StatusProhibit set to 50 ms and 75 ms, respectively, demonstrate that RLC with AM provides higher TCP throughput than RLC with UM.

Delay sensitive applications like VoIP experience better QoS when run over RLC UM while throughput sensitive applications like FTP perform better with the extra retransmissions of AM. Of the RLC settings examined, the timer t -Reordering is best set at a level sufficiently high to permit the MAC layer to effectively recover LTE transport blocks. The timer t -StatusProhibit is best set low to not adversely delay RLC ACKs and NACKs but not so low that the network spends much of its transmission opportunities sending higher priority AM STATUS messages. Further investigation of 4G LTE link layer architectural choices are important to improving modern application QoS over LTE networks.

Future work includes expanding on the applications and RLC layer settings used in our testing, as well as adding more features to the NS-3 LTE simulator.

6. REFERENCES

- [1] 3GPP. 3GPP Specification series, April 2013. <http://www.3gpp.org/ftp/Specs/html-info/36-series.htm>.
- [2] R. A. Bangun, E. Dutkiewicz, and G. Anido. Analysis of multi-player network games traffic. In *IEEE Multimedia Signal Processing*, pages 3–8, 1999.
- [3] Cisco. Cisco visual networking index: Global mobile data traffic forecast update, 2013-2018. Technical report, Cisco, 2014. <http://goo.gl/2ZfwuI>.
- [4] E. Dahlman, S. Parkvall, and J. Sköld. *4G LTE/LTE-Advanced for Mobile Broadband*. Academic Press, Oxford, 2011.
- [5] E. N. Gilbert. Capacity of a burst-noise channel. *Bell System Technical Journal*, 39(5):1253–1265, 1960.
- [6] N. Gordo. Evaluating video quality degradation from data packet losses in an LTE environment. Master's thesis, Lulea University of Technology, 2009.
- [7] T. M. Kawser, N. I. B. Hamid, N. Hasan, M. S. Alam, and M. M. Rahman. Limiting HARQ retransmissions in downlink for poor radio link in LTE. *International Journal of Inf. and Elec. Eng.*, 2012.
- [8] B. Levasseur. Impact of acknowledgments on application performance in 4G LTE networks. Master's thesis, Worcester Polytechnic Institute, 2014.
- [9] M. Makidis. Implementing and evaluating the RLC/AM protocol of the 3GPP specification. Master's thesis, Athens University of Econ. and Bus., 2007.
- [10] A. Norwine and O. Murphy. Characteristic time intervals in telephonic conversation. *Bell System Technical Journal*, 17(2):281–291, 1938.
- [11] I. T. Union. Pulse Code Modulation (PCM) of voice frequencies. Technical Report G.711, ITU-T, Nov. 1988. <http://goo.gl/PGaHfN>.
- [12] Verizon. IP latency statistics, Apr. 2014. <http://www.verizonenterprise.com/about/network/latency/>.