

An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance

Authors: Junxian Huang, Feng Qian, Yihua Guo,
Yuanyuan Zhou, Qiang Xu, Z. Morley Mao,
Subhabrata Sen, Oliver Spatscheck

CS577

Brett Levasseur

Outline

- Introduction
- LTE Data And Local Testbed
- LTE Network Characteristics
- Abnormal TCP Behavior
- Bandwidth Estimation
- Network Applications in LTE
- Conclusions
- Questions

- Previously observed LTE network characteristics
 - Higher bandwidths
 - Lower RTT
 - TCP underutilizes links
- This work examines
 - Measurements from real LTE network
 - TCP bandwidth estimation algorithm
 - Power management

LTE Network

- UE – User Equipment
- RAN – Radio Access Network
- CN – Core Network
- Monitor – Author’s data collection point
- PEP – Performance Enhancing Proxy

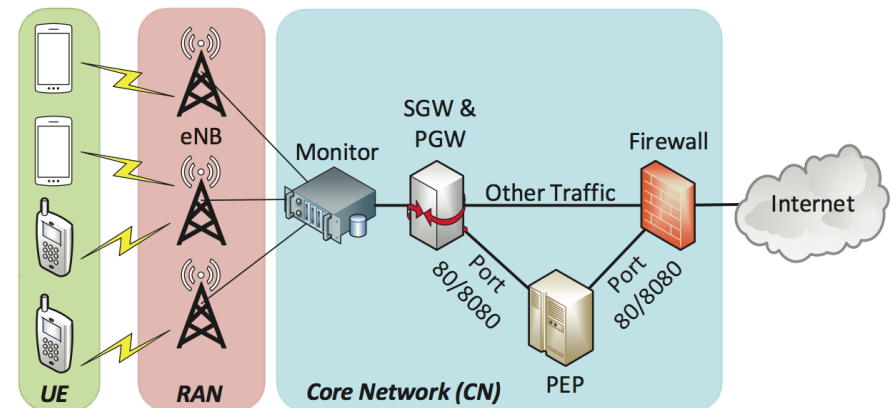
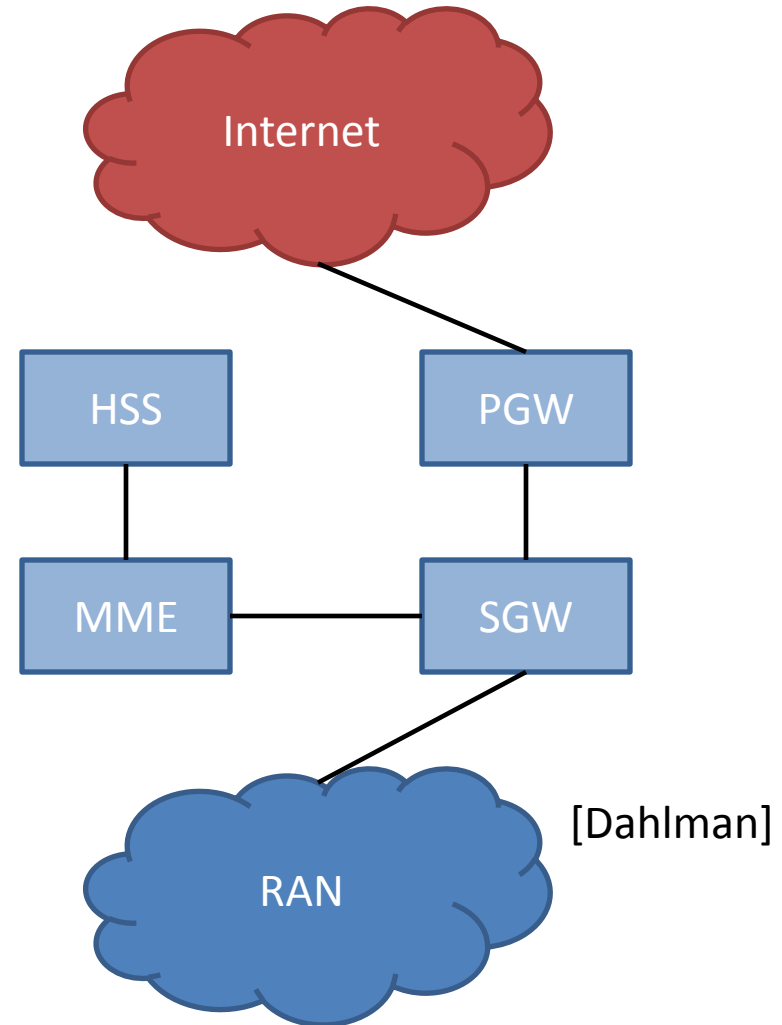


Figure 1: Simplified network topology of the large LTE carrier from which we obtained our measurement data.

LTE Network

- PEP – Not part of normal CN
- Intercepts TCP traffic on ports 80 and 8080
- Splits end to end TCP connection to two
 - UE to PEP
 - PEP to server
- Performs compression and caching



- Covered 22 eNBs in a US city
- Collection from 10/12/2012 - 10/22/2012
- Collected
 - IP and transport headers
 - 64 bit timestamps per packet
 - HTTP headers
 - 3.8 billion packets
 - 2.9 TB of traffic (324 GB of headers)

Local Test Bed

- UE
 - Samsung Galaxy S III
 - Android 4.0.4 / Linux Kernel 3.0.8
- Server
 - 2GB RAM / 2.4 GHz Intel Core 2 CPU
 - Ubuntu 12.04 / Linux Kernel 3.2.0-36-generic
 - TCP CUBIC
- Measured TCP throughput and RTT
- Used two different LTE networks

- Introduction
- LTE Data And Local Testbed
- **LTE Network Characteristics**
- Abnormal TCP Behavior
- Bandwidth Estimation
- Network Applications in LTE
- Conclusions
- Questions

Measurements

- Majority of traffic is TCP
- Majority of the remainder is UDP

TCP Flows (95.3 %)	HTTP (80/8080) 50.1%
	HTTPS (443) 42.1%
TCP Bytes (97.2%)	HTTP (80/8080) 76.6%
	HTTPS (443) 14.8%

TCP Flow Size

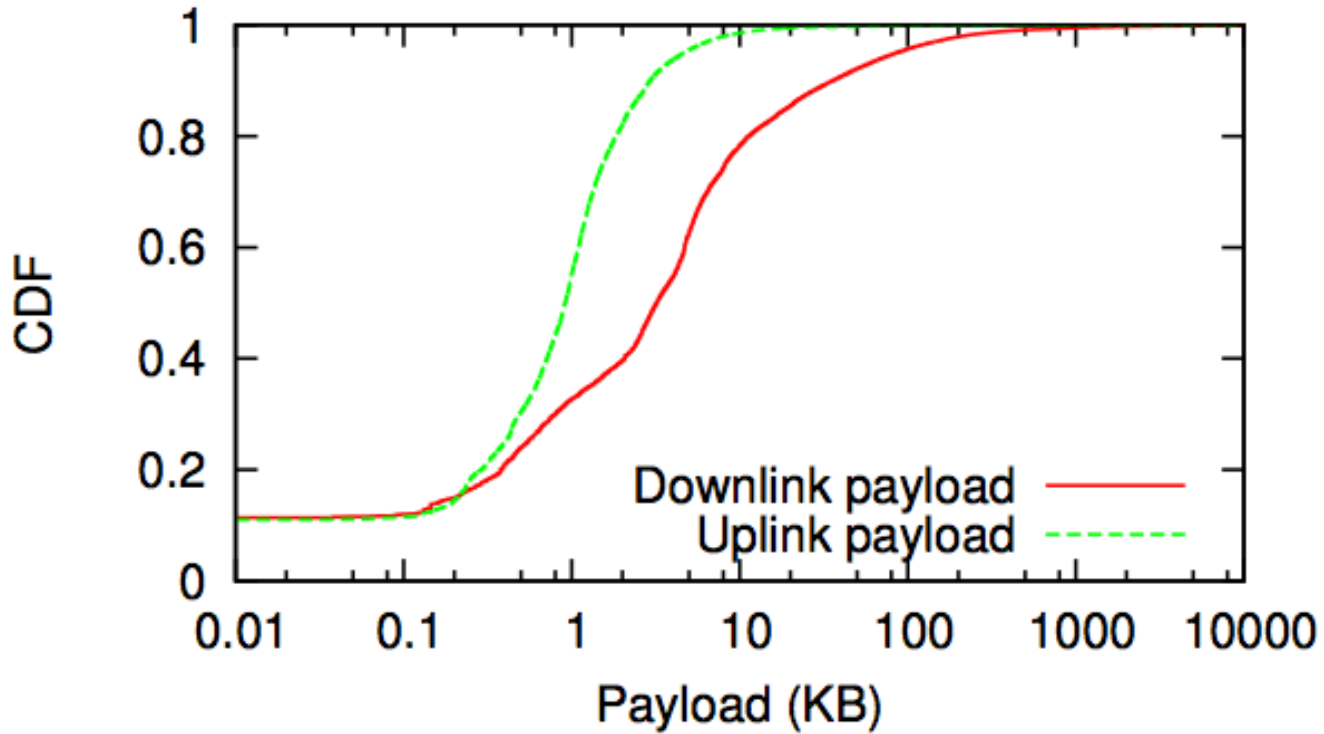


Figure 2: Distribution of TCP flow sizes.

TCP Flow Size

UL Flows	90%, less than 2.9 KB	
	10.9% have no uplink	
	Top 0.1% (by payload) account for 63.9% of total bytes	
	73.6% of the top flows are images	
DL Flows	90%, less than 35.9 KB	
	11.3% have no downlink	
	Top 0.6% (by payload) account for 61.7% of total bytes	
	Top 5% (by payload size)	Payload >= 85.9 KB
		80.3% use HTTP
74.4% video or audio		

TCP Flow Duration





Figure 3: Distribution of flow duration and the duration between the last payload byte to the end of the flow.

TCP Flow Duration

Flows	Duration
48.1%	< 5 seconds
6.8%	>= 3 minutes
2.8%	>= 10 minutes

Flows	Termination
86.2%	TCP FIN
5.4%	TCP RESET
8.5%	TCP SYN (did not connect properly)

Tail Time



10 seconds between last transmission and the UE radio being turned off

Figure 3: Distribution of flow duration and the duration between the last payload byte to the end of the flow.




Figure 4: An example of delayed FIN packet and its impact on radio resource management.

TCP Flow Rate

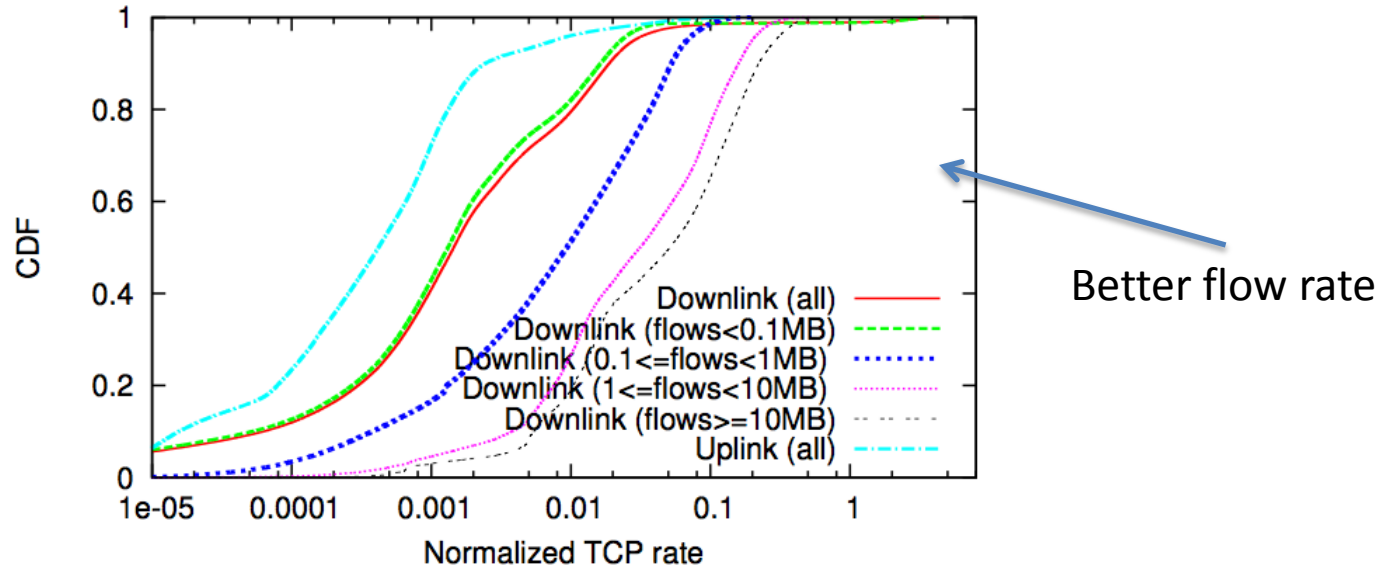


Figure 5: Distributions of normalized TCP flow rates.

- Larger flows send faster than smaller flows
- Flow duration and rate are more negatively correlated than on Internet backbone

TCP Concurrency




Figure 6: Concurrency for TCP flows per user uniformly sampled by time.

- 72.1% of the time there is only one active TCP flow
- Possibly higher for smartphones

RTT

WPI

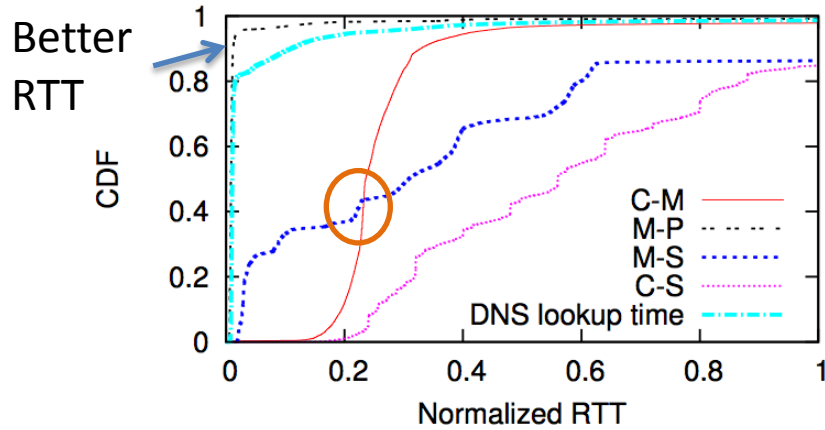


Figure 7: Distributions of normalized handshake RTT and DNS lookup time.

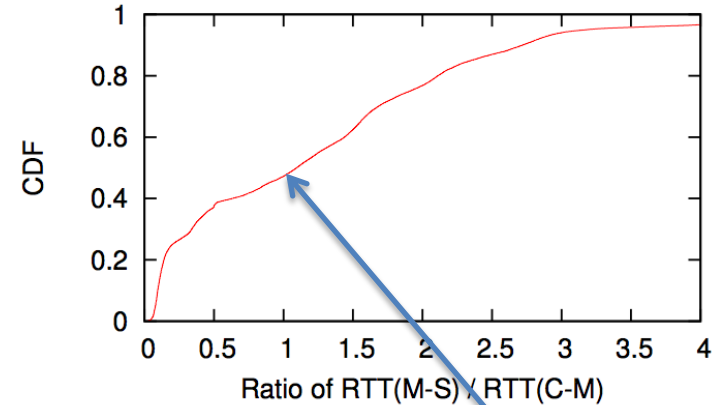


Figure 8: Distribution of the ratio between uplink and downlink RTT (for non-PEP traffic).

C-M = Client to monitor
M-P = Monitor to PEP
M-S = Monitor to server
C-S = Client to server

RTT to M-S > C-M

- RTT Monitor to server > than client to monitor
- Indicates wireless link is not largest delay factor

LTE Promotion Delay

- Time to turn radio on
- $G(TS_b - TS_a) = \text{RTT seen by UE}$
- G – Inverse ticking frequency of UE's clock

Percentile	Promotion Delay
25%	319ms
50%	435ms
75%	558ms

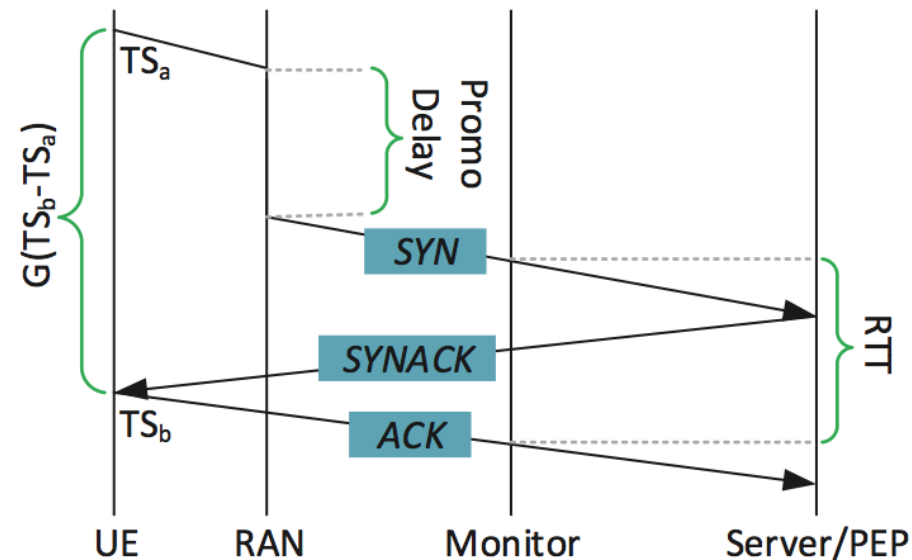


Figure 9: Estimating the promotion delay.

Queuing Delay

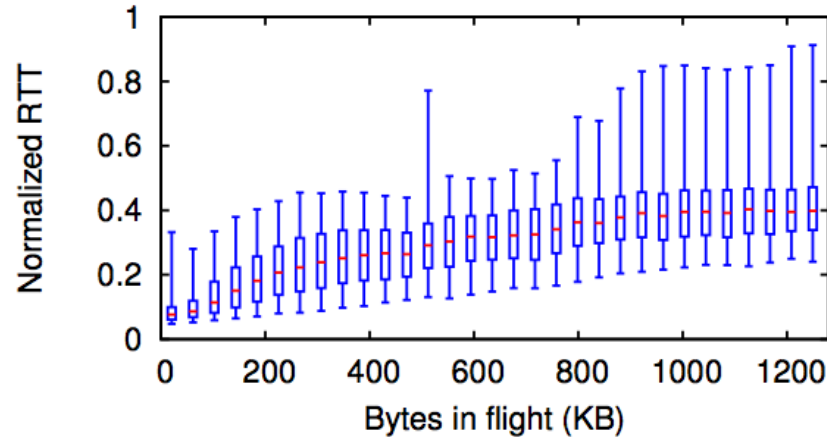


Figure 10: Downlink bytes in flight vs. downstream RTT.




Figure 11: Downlink bytes in flight vs. downstream RTT (controlled lab experiments with LTE Carrier A).




Figure 12: Downlink bytes in flight vs. downstream RTT (controlled lab experiments with LTE Carrier B).

Queuing Delay

- 10% of large flows have > 200 KB in-flight
- Leads to
 - Queue delay
 - Longer RTT
 - Created by long flows but impacts short flows

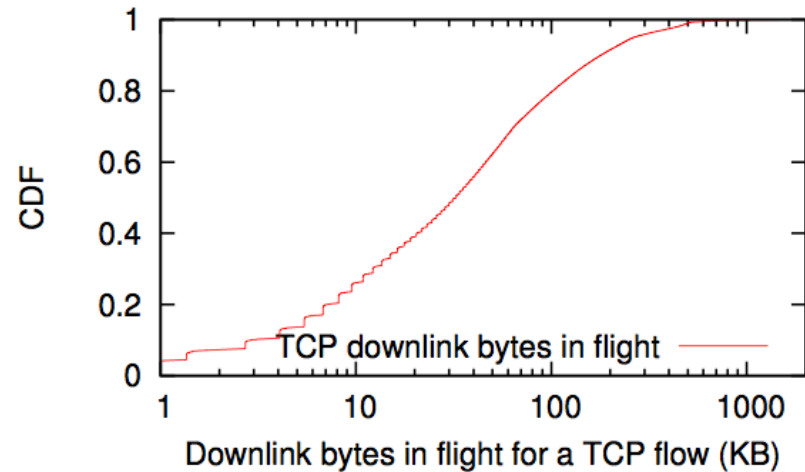


Figure 13: Distribution of downlink bytes in flight for large flows (> 1 MB).

TCP Retransmission Rate

- 38.1% of flows have no retransmission
- 0.06% is the median of flows with retransmission
- Physical/MAC layer retransmission reduced transport layer retransmission
- Study does not look at LTE RLC layer retransmissions

Measurement History

Table 1: Comparing with previous measurement studies

Study Time Location Type	Our Results October 2012 One US Metro Area LTE Only	3GTest [14] Aug to Dec 2009 Across U.S. Four 3G ISPs	4GTest [13] Oct to Dec 2011		SpeedTest [31] February 21 2011 to June 5 2011 (15 weeks)					
			Across U.S.		New York City		Madison WI, US		Manchester UK	
			LTE	WiMAX	Cellular	WiFi	Cellular	WiFi	Cellular	WiFi
5% TCP DL*	569	74 – 222**	2112	431	108	404	99	347	28	267
50% TCP DL	9185	556 – 970	12740	4670	1678	7040	895	5742	1077	4717
95% TCP DL	24229	1921 – 2943	30812	10344	12922	17617	3485	14173	3842	15635
5% TCP UL	38	24 – 52	387	172	52	177	55	168	25	180
50% TCP UL	2286	207 – 331	5640	1160	772	2020	478	1064	396	745
95% TCP UL	8361	434 – 664	19358	1595	5428	10094	1389	5251	1659	5589
5% HS RTT	30	125 – 182	37	89	68	21	99	24	98	34
50% HS RTT	70	160 – 200	70	125	159	54	184	69	221	92
95% HS RTT	467	645 – 809	127	213	786	336	773	343	912	313

* TCP DL: downlink throughput (kbps). TCP UL: uplink throughput (kbps). HS RTT: TCP handshake RTT (ms). 5%, 50%, 95% are percentiles.

** For a range $x - y$, x and y are the result of the worst and the best carriers, respectively, for that particular test.

- LTE outperforms 3G, WiMAX and WiFi
- 4GTest LTE is higher than measurements
 - Possibly due to rate limiting at remote server

- Introduction
- LTE Data And Local Testbed
- LTE Network Characteristics
- **Abnormal TCP Behavior**
- Bandwidth Estimation
- Network Applications in LTE
- Conclusions
- Questions

Duplicate ACKs

- Medians
 - 17 Dup ACKs
 - 2 out of order packets
- Over 29% of flows have > 100 Dup ACKs
- Ratio Dup ACK / out of order
 - 24.7% of flows over 25
 - Some up to 5,000
 - 1 out of order packet can cause many Dup ACKs

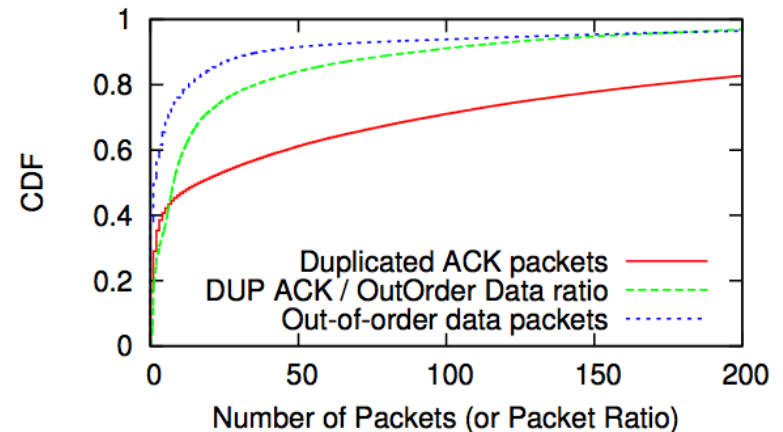


Figure 14: Observed duplicate ACKs and packet reordering in large TCP flows.

Undesired Slow Start

- Undesired Slow Start – Large RTT triggers RTO

- Author's measure undesired slow start with

$$R_{ss} = \frac{\theta_{[100,200]}}{\theta_{[0,100]}}$$

- Where $\theta_{[t_1,t_2]}$ is average downlink throughput from t_1 ms to t_2 ms after last Dup ACK

- $R_{ss} > 1.5$ in slow start

- 20.1% of large flows have ≥ 1 lost packet

- 12.3% of all large flows have ≥ 1 lost packet

Undesired Slow Start




Figure 15: Duplicate ACKs not triggering a slow start.

$$R_{SS} = 1.0$$




Figure 16: Duplicate ACKs triggering a slow start.

Mitigate Undesired Slow Start **WPI**

- Update RTO from duplicate ACKs with SACK
 - Take difference between SACK window of two consecutive duplicate ACKs
 - 82.3% of flows used SACK in dataset
 - < 1% of flows had packet reordering
- Update RTO from duplicate ACKs without SACK
 - Assume duplicate ACKs in response to data packets in order
- Prevent > 95% of observed undesired slow starts

- Introduction
- LTE Data And Local Testbed
- LTE Network Characteristics
- Abnormal TCP Behavior
- **Bandwidth Estimation**
- Network Applications in LTE
- Conclusions
- Questions

TCP Transmission Rate

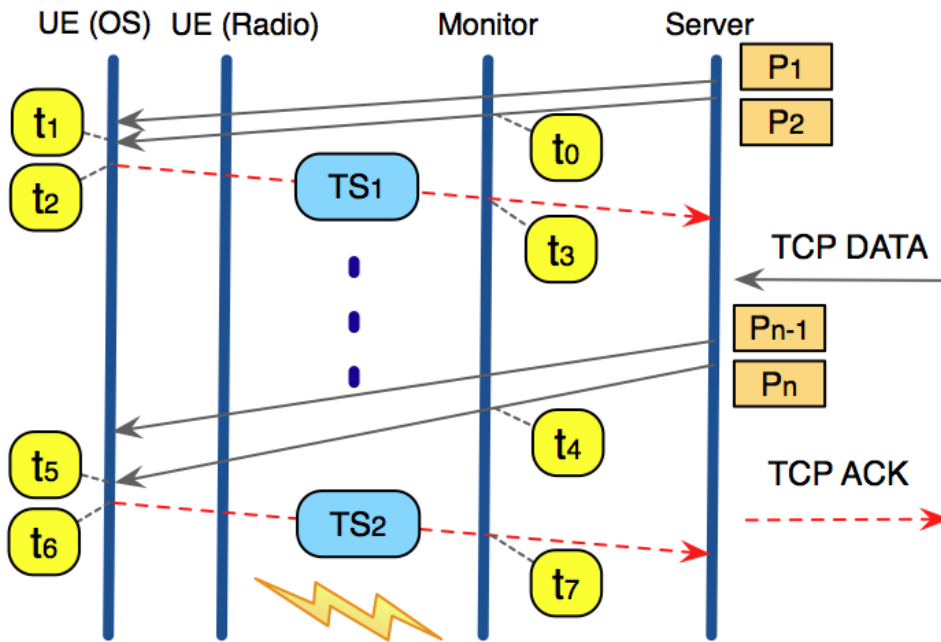


Figure 17: Typical TCP data transfer.

Sending Rate from Monitor

$$R_{snd} = \frac{S(n - 2)}{t_4 - t_0}$$

Receive Rate at UE

$$R_{rcv} = \frac{S(n - 2)}{t_5 - t_1}$$

TCP Timestamps

- Replace t_1 and t_5 with t_2 and t_6
- t_2 and t_6 originate at UE
- Replace t_2 and t_6 with TCP Timestamps
- Infer G

$$R_{rcv} \approx \frac{S(n-2)}{t_6 - t_2}$$

$$R_{rcv} \approx \frac{S(n-2)}{G(TS_2 - TS_1)}$$

$$G \approx \frac{TS_2 - TS_1}{t_7 - t_3}$$

Estimation Accuracy

- For accurate G
 $t_7 - t_3 > \delta_G$
- Error rate of G drops as δ_G grows
- At $\delta_G = 3$ seconds error rate $< 0.1\%$

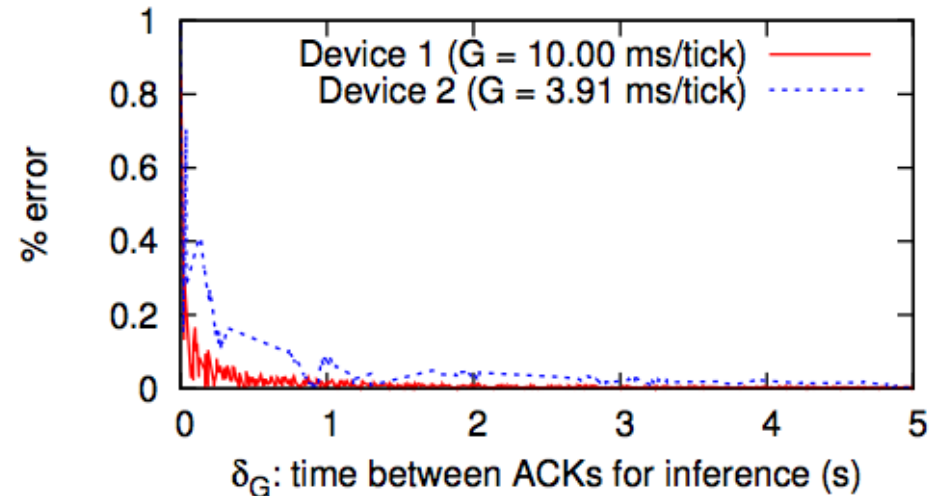


Figure 18: G inference and the selection of δ_G .

Estimation Accuracy

Flows	G
5.9%	NA
57.3%	1ms/tick
36.4%	10ms/tick
0.4%	100ms/tick

- With $\delta_G = 3$ seconds the error rate of inferred $G < 0.1\%$ for the majority of flows
- If G is unknown it is estimated from its formula

Estimation Summary

- G is known or inferred
- Calculate R_{snd}
- If $R_{\text{snd}} \geq C$ AND packets in order AND no duplicates AND last packet is not delayed ACK
 - R_{rcv} calculated
- If C too small – underestimate
- If C too large – not enough samples
- $C = 30\text{Mbps}$

Validation

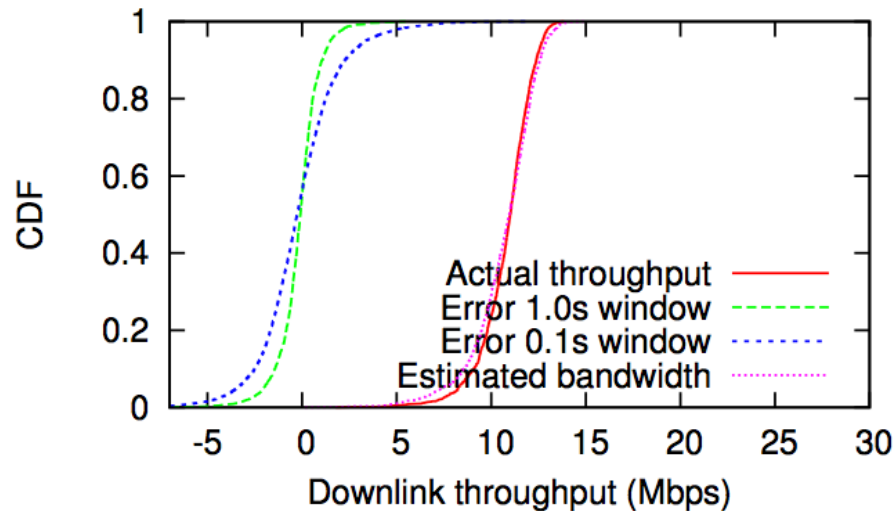


Figure 19: CDF of bandwidth estimation results for LTE network (controlled lab experiments with Carrier A).

- Compare server side estimate and UE trace
- 1 sec sample window average error rate is 7.9%
- 0.1 sec sample window average error rate has higher variation

Validation

- Actual throughput is UE perceived throughput
- Used 1 sec sample window
- Actual throughput varies around 10Mbps
- Error varies by ± 1 Mbps




Figure 20: Time series of bandwidth estimation for LTE network (controlled lab experiments with Carrier A).

Large Flow Utilization

- Median ratio 19.8%
- 71.3% of large flows < 50% utilization
- 6.4% use more bandwidth than estimated
- Average ratio 34.6%
- Low utilization, flows last longer
 - Higher radio usage




Figure 21: BW utilization ratio for large downlink TCP flows.

LTE Bandwidth Variation

- Two large flows
 - Two different users
 - Two different times
- Bandwidth varies over time
 - Condition of the wireless link
 - Movement
 - Resource scheduling

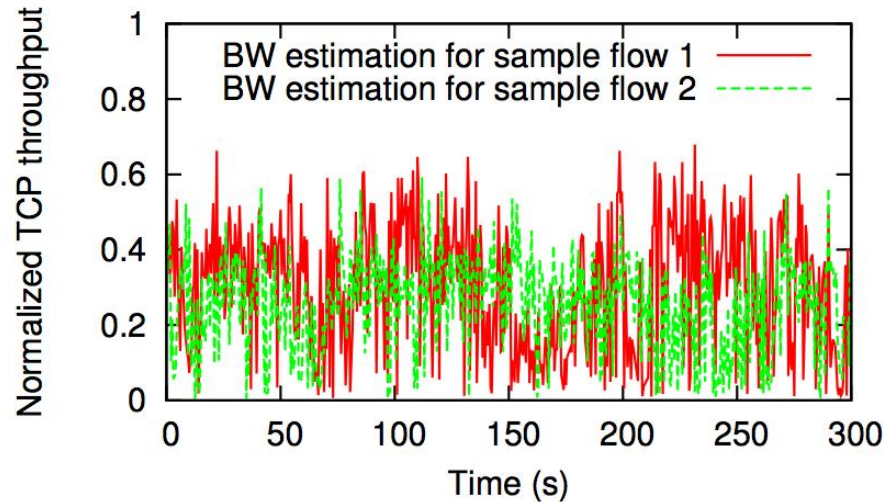


Figure 22: BW estimation timeline for two large TCP flows.

RTT & Throughput

- Experiments with modifiable RTT
 - Used iptables to redirect packets to scheduler
 - Scheduler changes available bandwidth similar to observed LTE
 - Scheduler injects delays to impact RTT
- Under small RTT TCP utilizes 95% of the bandwidth
- $RTT > 400\text{ms}$ utilization drops below 50%

- Introduction
- LTE Data And Local Testbed
- LTE Network Characteristics
- Abnormal TCP Behavior
- Bandwidth Estimation
- **Network Applications in LTE**
- Conclusions
- Questions

HTTP Characterization

HTTP Content	% of traffic
Video	37.8%
Images	19.5%
Text	11.8%
Zip	8.3%
Audio	6.5%
Other	5.6%
Unknown	10.5%

- 12.9% of video content is octet-streams generated mostly by video players

TCP Receive Window

- Shazam app on iOS
- 30s, 1 MB audio
- 0s – 2s 3Mbps
- Recv window full
- 2s – 9s < 300 Kbps
- Download could have been done in 2.5s
- Connection not closed until 30s

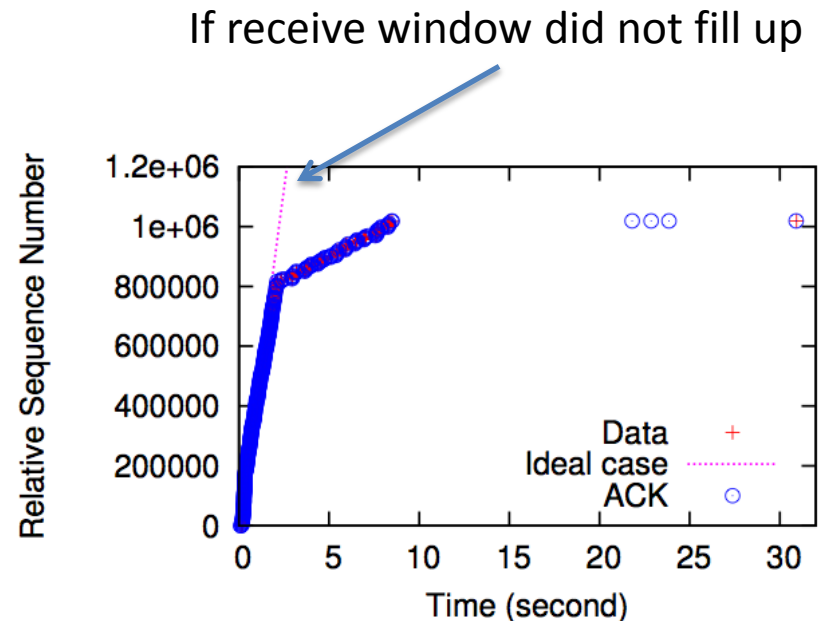


Figure 23: Full receive window slows Shazam player (a popular app) in downloading a 30-second music file.

TCP Receive Window

- Receive Window around 131,712 +- 600 bytes
 - True for > 90% of iOS, Android and Windows Phone flows
- Applications not reading from receive buffer quickly enough
- 52.6% of downlink TCP flows experience full receive window
 - 91.2% of these bottleneck happens in initial 10% of the flow duration

Application Design

- Netflix on iOS
- Multiple HTTP byte-range requests
- Many short HTTP responses
 - $< 1s$
 - 1 – 4MB
- Periodic requests leaves radio idle

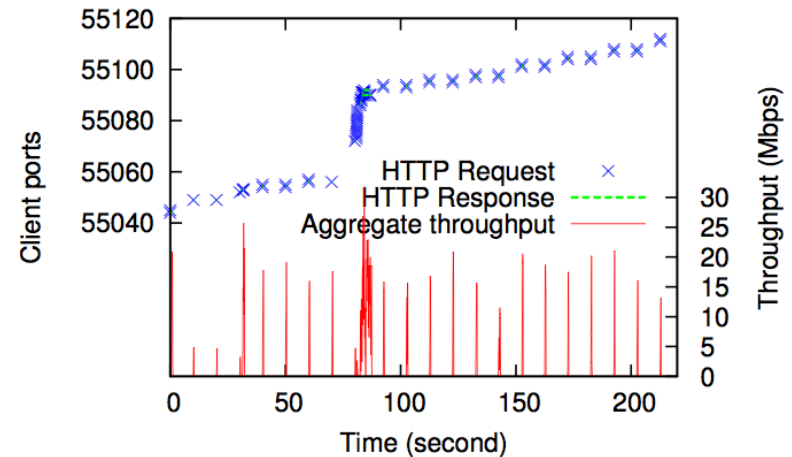


Figure 24: The periodic request behavior of Netflix player limiting its overall throughput.

- Manufactures reduce TCP receive window
 - Decreases buffer bloat
 - Underutilizes network
- Use application buffers
 - Relieves TCP buffers
 - Allows radio interface to close sooner
- Increase amount downloaded per request, decrease number of requests

Conclusions

- Not updating RTT from Dup ACKs causes performance issues with single packet loss
- Bandwidth estimation algorithm
 - 71.3% of large flows have < 50% utilization
 - High variation of network bandwidth
 - cwnd too slow to adjust
- TCP receive window throttles 52.6% of downlink flows
- App design underutilizes bandwidth

Questions

WPI

References

- Huan J. et al, An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance. SIGCOMM 2013.
- Dahlman E. et al, 4G LTE/LTE-Advanced for Mobile Broadband. Academic Press, 2011.