Who is the King of the Hill? Traffic Analysis over a 4G Network

Feng Li[1], Xiaoxiao Jiang[1], Jae Won Chung[1], and Mark Claypool[2]

[1]Verizon Labs, 60 Sylvan Rd, Waltham, MA, 02145
[2]Worcester Polytechnic Institute, 100 Institute Rd, Worcester, MA, 01609

***Abstract*—The increase in radio link capacities has brought significant growth in number and variety of mobile applications. To support these applications, network providers and researchers need an up-to-date, thorough understanding of the composition of 4G LTE traffic. In this paper, we conduct a deep study over a tier-1 U.S. wireless carrier to measure and analyze 4G LTE traffic. We collect data on 2.5+ million flows from 5000+ user devices, analyzing cross-layer characteristics for IPv4/6, HTTP/S and newly emerging Quick UDP Internet Connection (QUIC) traffic, with analysis extracted for the most prolific domains. To the best of our knowledge, our QUIC traffic analysis is the first published from a carrier's perspective. Results show data from multimedia sources (audio, video and images) dominate link capacities, with video, in particular, being the "king of the hill". Detailed analysis suggests methods for identifying encrypted video traffic, potentially leading to improved services that can provide network treatments for video.**

I. Introduction

The low latencies and high link capacities of Long Term Evolution (LTE) networks have propelled subscribers to access many Internet services through their mobile devices, which, in turn, have Quality of Service (QoS) requirements beyond that of traditional Web traffic [1], [2].

To improve application throughputs over LTE, many carriers have introduced performance-enhancing proxies (PEPs) as middle-boxes in their core networks [3], [4], [5], [6]. However, growth in firewall and TLS/SSL use makes it increasingly difficult for PEPs to identify content type, and, hence, choose traffic shaping strategies that provide better QoS. This leaves wireless eNodeBs unable to support traffic differentiation through QoS Class Identifiers (QCIs) alone. Blindly increasing data rates for video flows can confuse end-host bandwidth estimators (e.g., DASH clients) that choose high definition videos the wireless links are unable to handle, causing subscribers to experience long buffering times and generally poor quality of experience.

In order to provide enhanced QoS services beyond simply increasing bitrates, carriers need a better understanding of the breakdown of application traffic over current mobile networks. While individual operators may sometimes gather local reports on their networks, analysis of the data takes considerable effort and is typically not shared with the broader community. Given the rapid rate of evolution of wireless networks, devices and the application ecosystem, 4G LTE measurements and traffic classification should be repeated frequently, sampled at many different Internet vantage points – i.e., report everywhere, report often. Confirmation of results that may appear elsewhere is valuable as it suggests stability, while new and/or different results show trends over time and/or location.

This paper presents analysis of traffic from a U.S. tier-1 wireless carrier's network in the south central U.S. (e.g., Oklahoma) in November 2016, with over 5000 devices contributing over 2.5 million flows. Our analysis includes IPv4/6, HTTP/S and emerging Quick UDP Internet Connection (QUIC) traffic. To the best of our knowledge, our QUIC analysis is the first from a carrier's perspective. Our analysis provides the following observations:

1) HTTPS accounts for 38% of traffic volume, more than HTTP (31%). More than 53% of traffic volume is protected by TLS/SSL: 38% from HTTPS, 8% from HTTP2, and 8% from QUIC.

2) YouTube accounts for 22% of traffic volume, the largest of any kind of traffic. Google, Facebook and Apple are the top three traffic contributors by volume.

3) QUIC contributes nearly 64% of UDP traffic volume, while DNS still dominates UDP traffic in terms of the number of flows.

4) IPv6 contributes 53% of traffic volume, with Google's IPv6 traffic alone contributing 23% of all traffic volume.
5) Detecting video traffic may be done via per-packet payload length combined with throughput.

The rest of the paper is organized as follows: Section II summarizes related work; Section III describes our measurement methodology to gather and analyze 4G LTE traffic; Section IV presents cross-layer analysis of network characteristics; Section V presents mobile application analysis; and Section VI summarizes our conclusions and presents possible future work.

II. Related Work

Huang *et al.* characterized mobile network usage and performance with a 10-day local packet capture [7], [1]. While helpful to better understand mobile network usage, their analysis only focused on TCP, especially HTTP traffic, with no application layer traffic analysis, such as video. Our work provides a more recent trace and adds understanding of the share of video traffic over mobile networks.

Xu *et al.* analyzed app usage on smart phones by analyzing the HTTP Agent fields in HTTP headers [2]. While effective for illustrating HTTP use in 2010, their method is ineffective for the massive proportion of encrypted (HTTPS) traffic that has emerged since then. Our more recent observations show more HTTPS traffic than HTTP on mobile networks, making it unlikely that Xu *et al.*'s analysis (and method) provide an accurate representation of app usage over today's mobile networks.

Neither Huang *et al.* [1] nor Xu *et al.* [2] provided results for emerging Quick UDP Internet Connection (QUIC) [8] traffic. The deployment of QUIC presents new challenges for network management and traffic classification since it runs over UDP. To the best of our knowledge, ours is the first QUIC analysis presented from a wireless carrier's perspective.

Research has also illustrated that mobile service providers deployed incorrect traffic differentiation schemes [4], [5], [6], [9]. Flach *et al.* found that globally 7% of connections are identified as policed [4]. Choffnes *et al.* even identified a U.S. carrier that mistakenly rate limited traffic which was not video, and mis-charged subscribers with degraded video streaming from non-partners' sites [10], [5], [6]. The above research motivates the need for more accurate traffic classification, since without which, middle-boxes can only differentiate traffic blindly, often making QoS worse.

III. Methodology

This section describes our methodology to gather and analyze traffic from a tier-1 wireless carrier in the south central U.S.

Figure 1 depicts our measurement architecture. A middle-box labeled "AppDetect" mirrors traffic and builds flow-level information based on 5-tuples (source port, source IP, destination port, destination IP, and transport protocol). Due to the wide deployment of cloud services, reverse DNS using just the server's name is not sufficient for determining application type. For example, Netflix uses Amazon Web Services (AWS) to store video [11], so classifying based on reverse DNS alone may mistakenly identify Netflix video traffic as Amazon web traffic. Thus, we detect application flow types by reverse DNS, HOST names and server name identification (SNI). Once a flow is identified, subsequent packets matching the same flow are used to accrue statistics. When a flow sends a TCP FIN/FIN-ACK or does not transfer a packet for 60 seconds, the flow is marked as "terminated", whereupon the flow record (statistics and application type) are written into a private Cassandra[1] database for off-line analysis.

We use our architecture to gather data during the daytime (local time 9:00am-4:00pm) on two successive days, November 7-8, 2016. The carrier mirrored traffic from one end-user (UE) pool in the south central U.S. with over 5000 users. Our architecture extracted 85 GB of data, containing over 2.5 million flows. Respecting customer privacy, we did not collect subscribers' IDs, geographic locations, device types nor mobile numbers.

[1]http://cassandra.apache.org/

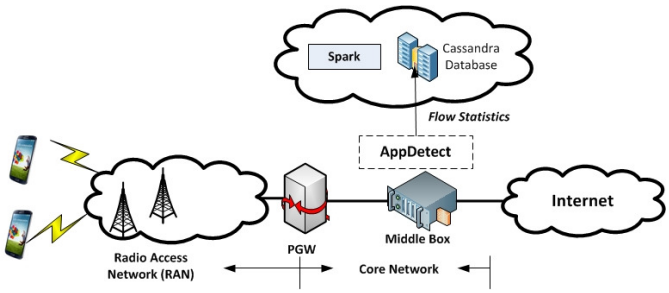

Fig. 1. Traffic Gathering and Analysis Architecture

IV. Network Characteristics

This section analyzes IP protocol, use of encryption, flow statistics and latency.

A. IP Protocol

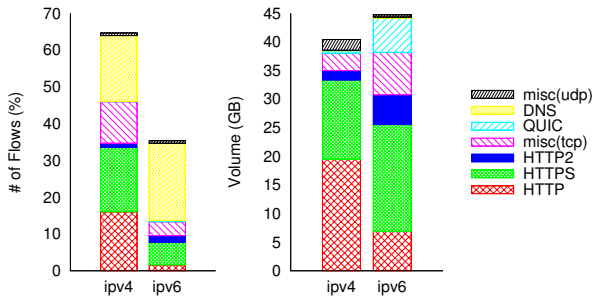

Fig. 2. Application Layer Protocol Distribution

Figure 2 shows the breakdown of traffic for IPv4 and IPv6. The graph on the left is the number of flows (percent) and the graph on the right shows the traffic volume (GBytes). 64% of the flows (1.6 million) are IPv4 compared to 36% for IPv6, but the IPv6 volume is slightly larger than IPv4. The mean size of IPv6 TCP flows is 81 KBytes, 7 times larger than the 11 KBytes IPv4 TCP mean.

B. Use of Encryption

Figure 2 also shows a breakdown of the major protocols in use. The HTTP-family of protocols still dominates mobile networks in terms of number of flows, and accounts for 77% of the overall traffic volume. However, 60% of HTTP traffic is HTTPS or HTTP2, both of which are protected by encryption via TLS/SSL. Since only 30% of traffic is unencrypted HTTP, techniques that classify traffic based solely on the HTTP Agent field [2] are ineffective – new techniques that can differentiate QoS classes for encrypted traffic are needed to provide appropriate QoS support [10].

QUIC [8], a multiplexed stream transport protocol over UDP, accounts for 7% of traffic volume, mostly over IPv6. From a middle-box perspective, QUIC is similar to TCP + TLS + HTTP/2 over UDP. Note, at the time of this study, Google mainly uses QUIC for YouTube services – Sections V and V-C provide more details on QUIC use.

C. Flow Statistics

Of the 85 GBytes total data, 99.9% is either TCP or UDP, with the rest (e.g., ICMP, GTP, SCTP, or IPSpec) contributing an insignificant 10 MBytes total of data. In terms of the number of flows, 990K (39%) are UDP DNS flows, although DNS traffic is insignificant in terms of volume.

We analyze characteristics important for tuning performance in an LTE network (e.g., eNodeB scheduling, per-flow billing, radio access network optimization and video detection over encrypted channels): duration, volume (size), and rate.

Duration is calculated from the time between the first packet and the last packet of a flow. Volume is calculated from the summation of packet lengths within a flow, including IP and transport layer headers. Rate is the throughput, calculated from the volume divided by the duration.

Note, we observe correlations among volume, duration and rate within a flow, similar to other studies [1], [12]. Due to space constraints, we do not present such analysis here.

Duration. Figure 3(a) shows a cumulative distribution function (CDF) of TCP and UDP flow durations (the x-axis, logscale). From the graph, most flows are "dragonflies" (lasting less than 2 seconds) [13] – 55% of TCP flows and 95% of UDP flows are less than 2 seconds. 11% of TCP flows and 0.5% of UDP flows last longer than 1 minute. Both TCP and UDP flow durations exhibit heavy-tailed tendencies, but none of the flows are "tortoises" (lasting more than 15 minutes) [13].

Size. Figure 3(b) depicts CDFs of flow volume (the x-axis, logscale). Flows are broken into TCP and UDP, both up (from the UE) and down (to the UE). Although most flows are small, the distributions of both UDP and TCP flow volumes show heavy-tailed tendencies. 90% of TCP flows are less than 27 KBytes, and 90% of UDP flows are less than 0.5 KBytes. The uplink TCP flows are somewhat smaller than the downlink TCP flows, but the difference is not as great as might be expected. While 90% of TCP uplink flows are smaller than 5 KBytes, the high link capacity of LTE networks coupled with the popularity of social media apps produce more uplink traffic than downlink for many TCP flows. To illustrate this, Figure 3(b) separates the uplink and downlink flow volumes observed for Instagram.[2] Instagram volumes for both uplink and downlink are much higher than the volumes of other TCP flows. Surprisingly, 14% of TCP flows are never successfully established (and not shown), containing fewer than 3 packets – future work could investigate causes for these TCP connection failures.

Meanwhile, the UDP flow volume CDF also shows a heavy-tailed tendency, and some of the UDP flows are more than 10 MBytes. These high volume flows are also long lasting (60+ seconds), contributing over 7 GBytes, about 80% of the total UDP traffic. Starting in 2016, Google deployed UDP-based QUIC [8] in Chrome-based browsers, carrying YouTube and other Google-served traffic in North America. Section V provides more analysis of QUIC and YouTube traffic.

[2]Instagram is a photo sharing app, http://www.instagram.com/

Rate. Because of the asymmetric channel bandwidth and encoding schemes, LTE uplink rates are generally lower than downlink rates. Figure 3(c) depicts CDFs of flow rate (throughput, the x-axis in logscale). 90% of TCP flows have rates less than 285 Kbps down and 60 Kbps up, while 90% of UDP flows have rates less than 815 Kbps down and 379 Kbps up.

Applications running with a request/response protocol (e.g., HTTP) might never send enough data to fill a high capacity LTE pipe [14], making it difficult to improve TCP throughput by simply adding more radio resources. TCP "splitting" via a proxy may improve LTE link utilization [15], but there is little chance of improving throughput for the numerous small flows apparent in Figure 3(b). This suggests focusing throughput treatments on "elephant" flows that can have large congestion windows [16], [15].

D. Network Latency

Round trip time (RTT) is not typically useful for classifying flows but can be meaningful for inferring application performance or for motivating RTT-based congestion control (e.g., BBR) [14]. We estimate the RTT for TCP from the three-way handshake and for UDP from the DNS request-response pairs.

Figure 3(d) depicts a CDF of TCP initial RTTs (the x-axis, logscale). The median is 79 ms, possibly due to UE buffering [17]. The TCP RTT distribution shows a heavy-tailed tendency.

Since propagation delays over 100 ms are unlikely over LTE [7], further investigation shows the large TCP initial RTTs are due to: 1) Retransmission of SYN/SYN-ACK. About 1% of TCP flows have at least one packet retransmission during the three-way handshake. Such retransmissions significantly increase the initial RTT measurements, as shown by the second RTT trendline in Figure 3(d). 2) Backwards compatible 2G/3G devices. 210 TCP flows without SYN/SYN-ACK retransmission also experienced initial RTTs larger than 900 ms, an amount observed on 2G/3G networks [7].

Figure 3(d) also shows the distribution of DNS request times – effectively, the approximate RTT between the middle-box and the DNS server. From the graph, the RTTs between the middle-box and DNS servers are short – 64% of DNS response times are less than 10 ms – confirming the time efficiency of DNS caching for mobile core networks [1].

V. Mobile Applications

This section analyzes applications based on server names (SNIs), extracted primarily from TCP and TLS/SSL handshakes and HOST names in HTTP requests, classifying 88% of the traffic (by volume) – 12% of the traffic (9.1 GBytes of TCP and 0.9 GBytes of UDP) is marked as "unknown". We focus on analysis of traffic from well-known content providers, providing a better understanding of current traffic in mobile networks.

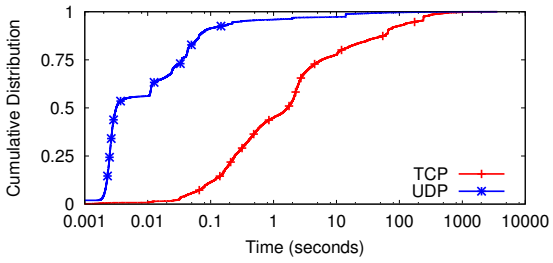

(a) Distribution of Flow Duration

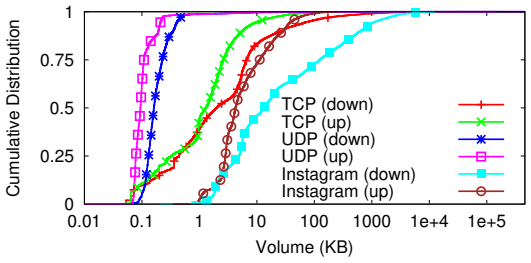

(b) Distribution of Flow Volume

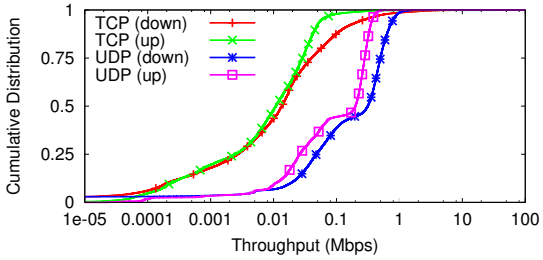

(c) Distribution of Flow Rate

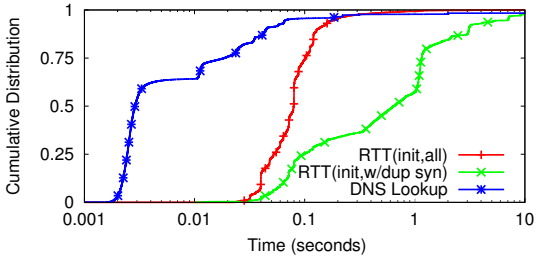

(d) Distribution of Round-trip Times

Fig. 3. Transport Layer Flow Analysis

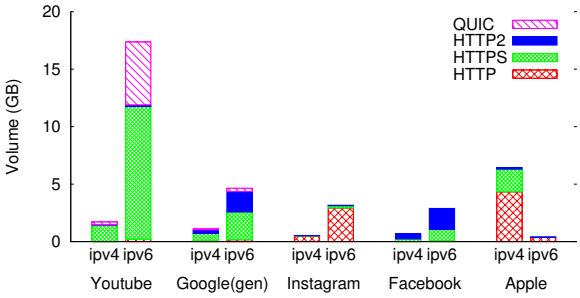

Fig. 4. Traffic from Google, Facebook and Apple

A. Top Content Providers

Figure 4 depicts the traffic volume (y-axis) from the top 5 content providers (x-axis) observed in our traces: YouTube, Google, Instagram, Facebook, and Apple.

Google (including YouTube) is the largest content provider over the carrier's network. Most of Google's traffic is protected by TLS/SSL and is mainly over IPv6. In total, there are only 430 MBytes of unencrypted HTTP traffic from Google and YouTube. Google's 22 GBytes of IPv6 traffic accounts for 49% of total IPv6 traffic, 9x higher than Google's IPv4 traffic (2.7 GBytes). This may mean that Google has moved its default services to IPv6 for the North American market.

A small number of QUIC flows (1250, all IPv6) deliver 5.6 GBytes of data from YouTube. Although only recently announced (Q2 of 2016), QUIC already constitutes 25% of YouTube's overall traffic. This suggests ISP should have a better understanding of QUIC-based applications for planning and treatment.

Instagram is the most popular service from Facebook, constituting about half of all Facebook traffic. From an implementation perspective, Instagram is mainly over HTTP unlike most other Facebook services. Similar to Google, the data may mean Facebook has migrated much of its traffic to IPv6 in North America.

Traffic from Apple consists of a variety of different services such as iTunes, Apple Maps, and iCloud, most of which is over IPv4. Our analysis (not shown) reveals that iTunes constitutes the majority of Apple's traffic.

B. Top Applications

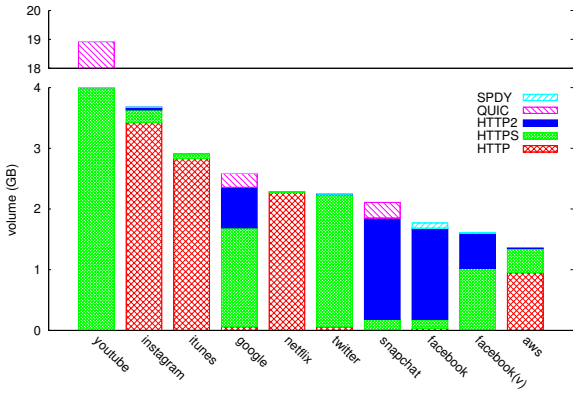

Fig. 5. Traffic from Top 10 Applications

Figure 5 depicts the traffic volume (on the y-axis) from the top 10 mobile apps/servers (on the x-axis) for our traces: YouTube, Instagram, iTunes, Google, Netflix, Twitter, Snapchat, Facebook, Facebook Video, and Amazon Web Services (AWS). Note, YouTube alone has about 19 GBytes of traffic, so the y-axis is "broken" from 4-18 Gbytes. The Google bar only contains Google Web traffic, excluding products such as YouTube, Play, GMail and Google Docs. Facebook has 3

applications in the top 10: Instagram, Facebook "generic", and Facebook Video (labeled as "facebook (v)" in the figure). The label "facebook" consists of 280 MBytes of Facebook Messenger traffic, but because Facebook Messenger is migrating to TLS/SSL, not all Messenger traffic can be differentiated from Facebook "generic" traffic, hence they are grouped together.

The top 10 applications account for about 40 GBytes of traffic, 46% of the total traffic.

Videos and social apps are the two most popular mobile application types. Among the top 10 applications, 3 are video applications (YouTube, Netflix, and Facebook Video) and 4 are social apps (Instagram, Twitter, Snapchat, and Facebook "generic"). Some social apps, such as Snapchat, produce a lot of video content, also.

Snapchat is the only application using QUIC other than Google, but all Snapchat QUIC traffic is served by Google servers. All Snapchat traffic detected is encrypted.

Note Akamai accounts for 1.5 GBytes of data, higher than Amazon AWS (1.4 GBytes). However, since Akamai provides services on behalf of iTunes, iCloud, Instagram, and Facebook, these totals are subtracted from Akamai's volume, leaving Akamai at 12th largest.

C. YouTube Traffic

Since video is projected to make up 75+% of mobile traffic by 2021 [18], identifying video is critical for improving overall QoS over mobile networks [19]. This section analyze flows from YouTube, which accounts 22% of the total observed traffic.

YouTube delivers video services through HTTPS and QUIC [8] in North America. Both HTTPS and QUIC are protected by TLS/SSL, making it difficult for middle-boxes to infer the application type even with deep packet inspection. We differentiate HTTPS-based YouTube flows based on observed SNIs (Table I).

Unfortunately, QUIC does not provide any public header information similar to SNIs. Thus, we classify QUIC into two groups by using reverse DNS: i) *QUIC (youtube)* – QUIC traffic from YouTube servers, and ii) *QUIC (misc)* – QUIC traffic from other Google servers. Our offline observations show QUIC (misc) carries performance statistics on Chrome browsers.

Offline analysis and aggregation produce Table I. Based on frequency, the table suggests flows with SNIs matching *sn-*goog- levideo.com* are likely carrying video content, while flows with SNIs matching **.youtube.com* provide Web pages. Flows with SNIs matching *ytimg*[3] store YouTube thumbnail images. Flows with SNIs matching *manifest.googl- evideo.com* provide manifest files (for video scaling), and likely do not contain video segments.

Figure 6 compares key statistics for YouTube flows – googlevideo (video), ytimg (images), quic youtube (likely video) and quic misc (unlikely video).

Figure 6(a) depicts the distributions of the flow durations (x-axis). From the figure, flow duration by itself is not a good differentiator for video flows since the duration of googlevideo flows (which are videos) or QUIC (YouTube) is even smaller than ytimg flows (which are just images). Ytimg flows only terminate when subscribers close their Web browsers or player apps, making ytimg flows longer than most video flows.

[3]ytimg is an acronym for "YouTubeIMaGe."

TABLE I
SNIs FROM YOUTUBE

Pattern	Count	Freq (%)
r(\d+)—sn-*googlevideo.com	11682	77.64
*.youtube.com or *.youtu.be	1754	11.66
*.ytimg.com	758	5.04
redirector.googlevideo.com	620	4.12
manifest.googlevideo.com	108	0.72
youtube-nocookie.com	42	0.31
without SNIs	77	0.51
Total	15046	100.00

Figure 6(b) depicts the distributions of downlink volume (x-axis) of YouTube flows. Googlevideo and QUIC (YouTube) are larger than other flows – 18% of googlevideo and 16% QUIC (YouTube) flows are larger than 1 MByte. However, 31% of googlevideo flows are smaller than 5 KBytes – too small to carry any actual video content.

Figure 6(c) shows the distributions of downlink throughput (x-axis) of YouTube flows. For reference, the vertical *brown* dashed line at 1.5 Mbps marks the bitrate of a YouTube video encoded at 480p [20]. The throughputs of googlevideo and QUIC (YouTube) are much higher than the others. This suggests that throughput might be a good candidate to use to classify videos, but can only be done after some time (i.e., not upon flow startup). About 25% of googlevideo and 13% QUIC (YouTube) flows yield a throughput more than 1.5 Mbps. One tier-1 carrier rate-limits all video flows to 1.5 Mbps as its video optimization solution [5]. In this case, if their subscribers' video players cannot automatically adapt to lower quality segments, 25% of YouTube video flows would be punished by their traffic policing policy, and the subscribers would experience long buffering times with multiple stalls.

Figure 6(d) shows the average per packet Transport layer payload length (x-axis) for YouTube. From the graph, flows likely carrying video content are dominated by large packets – 60% of googlevideo.com and 25% of QUIC (YouTube) flows have a payload of 1000+ bytes per packet. This suggests per packet payload length, perhaps combined with throughput, might be a good metric for detecting encrypted video flows.

VI. CONCLUSIONS

This paper presents cross-layer traffic analysis from a tier-1 wireless carrier's network serving the south central U.S. (e.g., Oklahoma), providing recent insight into an previously un-analyzed geographic region for comparison/confirmation with previous studies and discovery of new trends. Focus includes IPv4/6, HTTP/S and newly emerging QUIC, broken down into IP, transport and application layers.

Based on our analysis, the high volumes of IPv6 and HTTPS suggest renewed focus on traffic classification that is payload agnostic. QUIC comprises over 60% of UDP traffic

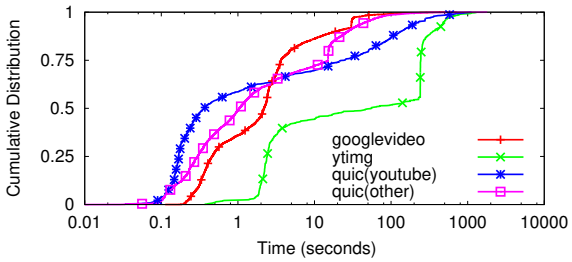

(a) Duration of YouTube Flows

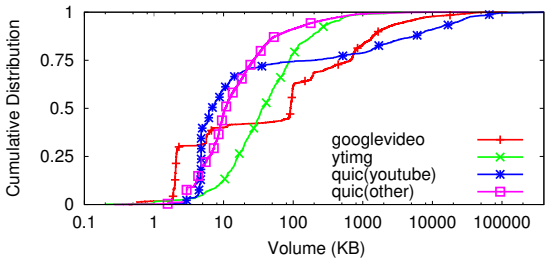

(b) Downlink Volume of YouTube Flows

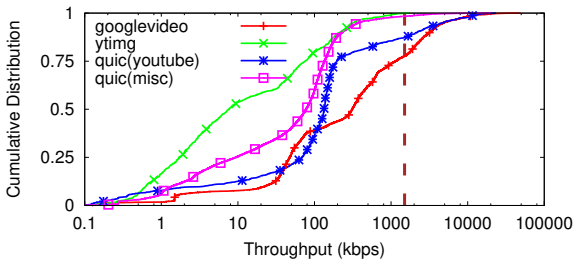

(c) Downlink Throughput of YouTube Flows

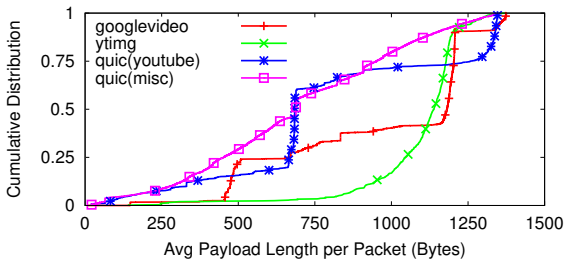

(d) Downlink Average Packet Length of YouTube Flows

Fig. 6. YouTube Flow Analysis

and warrants special attention since it greatly increases UDP volumes. Google, Facebook, and Apple are the top 3 content providers while YouTube alone accounts for over 20% of the total traffic.

In addition to providing a detailed snapshot of a current (late 2016) commercial mobile network, the results from our traffic analysis should be useful for researchers and network providers – the combination of better simulation, emulation, modeling and classification of content can lead to improved wireless network services, including content-based billing and improved quality of service.

Acknowledgments. We would like thank the anonymous ICC reviewers for their valuable feedback, and our colleagues Vijay Nanjundan and Atreya Praveen for their helpful discussions and assistance.

REFERENCES

[1] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, and O. Spatscheck, "An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance," *ACM SIGCOMM Computer Communication Review*, 2013.

[2] Q. Xu, J. Erman, A. Gerber, Z. Mao, J. Pang, and S. Venkataraman, "Identifying Diverse Usage Behaviors of Smartphone Apps," in *Proceedings of ACM SIGCOMM IMC*, Berlin, Germany, November 2011.

[3] N. Dukkipati, M. Mathis, Y. Cheng, and M. Ghobadi, "Proportional Rate Reduction for TCP," in *Proceedings of ACM SIGCOMM IMC*, Berlin, Germany, 2011.

[4] T. Flach, P. Papageorge, A. Terzis, L. Pedrosa, Y. Cheng, T. Karim, E. Katz-Bassett, and R. Govindan, "An Internet-Wide Analysis of Traffic Policing," in *Proceedings of ACM SIGCOMM*, Florianopolis, Brazil, August 2016.

[5] A. M. Kakhki, F. Li, D. Choffnes, A. Mislove, and E. Katz-Bassett, "BingeOn Under the Microscope: Understanding T-Mobile's Zero-Rating Implementation," in *Proceedings of Internet-QoE Workshop*, Florianopolis, Brazil, August 2016.

[6] F. Li, A. M. Kakhki, D. Choffnes, P. Gill, and A. Mislove, "Classifiers Unclassified: An Efficient Approach to Revealing IP Traffic Classification Rules," in *Proceedings of ACM SIGCOMM IMC*, Santa Monica, CA, USA, 2016.

[7] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "A Close Examination of Performance and Power Characteristics of 4G LTE Networks," in *Proceedings of MobiSys*, Low Wood Bay, UK, 2012.

[8] A. Langle et al., "The QUIC Transport Protocol: Design and Internet-Scale Deployment," in *Proceedings of ACM SIGCOMM*, Los Angeles, CA, USA, August 2017.

[9] A. Molavi Kakhki, A. Razaghpanah, A. Li, H. Koo, R. Golani, D. Choffnes, P. Gill, and A. Mislove, "Identifying Traffic Differentiation in Mobile Networks," in *Proceedings of ACM SIGCOMM IMC*, Tokyo, Japan, October 2015, pp. 239–251.

[10] G. Dimopoulos, I. Leontiadis, P. Barlet-Ros, and K. Papagiannaki, "Measuring Video QoE from Encrypted Traffic," in *Proceedings of ACM SIGCOMM IMC*, Santa Monica, CA, USA, November 2016.

[11] V. K. Adhikari, Y. Guo, F. Hao, M. Varvello, V. Hilt, M. Steiner, and Z.-L. Zhang, "Unreeling Netflix: Understanding and Improving Multi-CDN Movie Delivery," in *IEEE INFOCOM*, Las Vegas, NV, March 2012.

[12] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker, "On the Characteristics and Origins of Internet Flow Rates," in *Proceedings of ACM (SIGCOMM)*, Pittsburgh, PA, USA, August 2002, pp. 309–322.

[13] N. Brownlee and kc Claffy, "Understanding Internet Traffic Streams: Dragonflies and Tortoises," *Communications Magazine, IEEE*, vol. 40, no. 10, pp. 110–117, 2002.

[14] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson, "BBR: Congestion-Based Congestion Control," *ACM Queue*, vol. 14, September-October, 2016.

[15] J. Snellman, "Mobile TCP Optimization: Lessons Learned in Production," Telco. Networks, Tech. Rep., August 2015. [Online]. Available: {http://conferences.sigcomm.org/sigcomm/2015/pdf/papers/hotmiddlebox/keynote.pdf}

[16] L. Le, J. Aikat, K. Jeffay, and F. D. Smith, "Differential Congestion Notification: Taming the Elephants," in *Proceedings of IEEE ICNP*, Berlin, Germany, October 2004, pp. 118–128.

[17] Y. Guo, F. Qian, Q. A. Chen, Z. M. Mao, and S. Sen, "Understanding On-device Bufferbloat for Cellular Upload," in *Proceedings of ACM SIGCOMM IMC*, Santa Monica, CA, USA, November 2016.

[18] Cisco System Inc., "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update," Cisco System Inc., Tech. Rep., March 2017.

[19] M. Ghasemi, P. Kanuparthy, A. Mansy, T. Benson, and J. Rexford, "Performance Characterization of a Commercial Video Streaming Service," in *Proceedings of ACM SIGCOMM IMC*, Santa Monica, CA, USA, November 2016.

[20] Google Inc., "Live encoder settings, bitrates, and resolutions: YouTube Help," August 2016. [Online]. Available: {https://goo.gl/Acp7Dt}