

Profiling Energy Consumption of DASH Video Streaming over 4G LTE Networks

Jingyu Zhang^{1,2*}, Gan Fang¹, Chunyi Peng¹, Minyi Guo², Sheng Wei^{3,4}, Viswanathan Swaminathan³

¹Department of Computer Science and Engineering, The Ohio State University, Columbus, OH, USA

²Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China

³Adobe Research, San Jose, CA, USA

⁴Department of Computer Science and Engineering, University of Nebraska-Lincoln, Lincoln, NE, USA

Email: zhangzhang@sjtu.edu.cn, fang.254@osu.edu, chunyi@cse.ohio-state.edu, guo-my@cs.sjtu.edu.cn, swei@adobe.com, vishy@adobe.com

ABSTRACT

Video streaming is a major source of unprecedented traffic growth in mobile networks, especially in high-speed 4G LTE networks. However, video streaming also poses a critical challenge on energy consumption. The battery-on smartphones and tablets can not sustain active online video streaming for several hours over cellular networks. In this paper, we seek to understand how DASH (Dynamic Adaptive Streaming over HTTP) based video streaming can be energy efficient in the context of LTE networks. We profile the impacts of different streaming strategies and configurations (i.e., segment length and buffer size) on energy consumption. We analyze radio resource control (RRC) to explain and quantify their impacts. Our measurements in real LTE networks show that there still exists a large saving space (more than 30%) for us to improve energy efficiency of mobile devices using appropriate DASH settings. The saving space in some extreme cases can be even larger.

CCS Concepts

•Networks → Mobile networks; *Application layer protocols*; Network measurement;

Keywords

Video Streaming; DASH; LTE; Power Consumption; Energy Saving;

1. INTRODUCTION

Video streaming is primarily going mobile. Thanks to the rapid deployment of high-speed 4G LTE networks, more and more traditional services and businesses on the PCs have

*This work was done when Jingyu Zhang was a visiting student at the Ohio State University.

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.

MOVID'16, May 13 2016, Klagenfurt, Austria

© 2016 ACM. ISBN 978-1-4503-4357-2/16/05...\$15.00

DOI: <http://dx.doi.org/10.1145/2910018.2910656>

moved to mobile devices. There is no surprise that video streaming is the most popular online service which will contribute to 80% of the internet traffic in 2019, up from 64% in 2014 [12]. More and more users watch videos on their handheld smartphones and tablets through cellular networks. It facilitates anytime, anywhere watching experience (at the cafe, airport, bus, subway, train, etc.). Moreover, LTE offers much faster access (more than 10Mbps) than 3G (several hundred Kbps to several Mbps) and supports high-quality video services.

The popularity of mobile video streaming raises one critical issue in the mobile context, which is the energy consumption of mobile devices. The battery technology is far behind the rapid development of mobile processors and mobile networks. Mobile users have to worry about their battery capacity when they watch videos online. For example, the power consumption varies widely from tens of milliwatts (idle) to several watts (video watching) and the battery drains fast. As a result, energy-efficient video streaming is highly desirable to boost mobile user experience.

Video streaming has been actively studied in the recent years. Most research and industry efforts have been devoted to improving QoS (Quality of Service) and user experience on different aspects such as video quality, reliability, download speed [13, 19]. Wei et al. proposed a power efficient video streaming mechanism on mobile devices over cellular networks [21]. They developed an analytical model to identify and quantify the power inefficiency in mobile video streaming, considering the mismatch between HTTP request schedule and the radio resource control schedule. Li et al. designed GreenTube [17], a system that optimizes power consumption for mobile video streaming by judiciously scheduling downloading activities to minimize unnecessary active periods of 3G/4G radio. Huang et al. gave a close examination of power consumption of Android phones running various apps over LTE networks [14].

However, there still has been no comprehensive understanding of real power consumption of mobile video streaming in 4G LTE networks, especially the impact of streaming settings. In this work, we investigate the energy consumption of video streaming strategies. In particular, we study DASH, the most popular protocol for video content providers to delivery high QoS and reliable streaming [22]. We profile the impact of different DASH settings (i.e., video parameters and player configurations). We focus on two

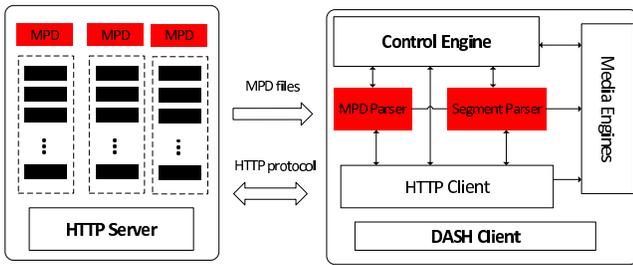


Figure 1: DASH-based video streaming system.

main factors, which allow reconfigurations initiated by the video providers or demanded by mobile clients. We measure the power consumption in real LTE networks in a variety of scenarios and quantify their impact and room for improvement. We also look into 4G LTE network factors to understand the root cause. We examine the RRC procedure, which is the widely used in 4G LTE networks to support video streaming or other mobile data traffic. We uncover that the underlying RRC transition in 4G LTE networks allows sufficiently large room (more than 30%) to improve energy efficiency for video streaming.

We make three main contributions in this work.

- We measure power consumption of video streaming in real 4G LTE networks and profile the impact of DASH streaming parameters (segment length and buffer size).
- We deduce the root causes of power inefficiency through analyzing the real RRC trace in 4G LTE cellular network.
- We quantify the potential energy saving space. We also discuss possible new solutions, which may effectively reduce the energy consumption in DASH-based video streaming.

2. BACKGROUND

HTTP streaming is the most popular video delivery technology for online video play back [15, 23, 16]. Typically, encoded video content in a pre-defined file or delivery format is sliced into many short segments (“chunks”). The URLs and other metadata about these chunks are published by in a manifest file or files associated with the video content. The client downloads the manifest file first to learn about the segments it needs, downloads these segments over HTTP, and plays them back seamlessly in a video session. At any point, the client can switch from a high bitrate segment to a lower bitrate segment (or vice-versa) based on the network bandwidth available or other conditions at the client. This is one of the biggest benefits of HTTP streaming. Also, this leverages the HTTP caching infrastructures at Content Delivery Networks (CDNs) and Internet Service Providers (ISPs) bringing the total cost of operation much lower for large scale video delivery [18].

DASH stands for Dynamic Adaptive Streaming over HTTP, which is an ISO/IEC MPEG standard for HTTP Streaming [9], known as MPEG-DASH. There are several DASH implementations that are widely deployed for large scale video delivery. The media presentation description (MPD) is the DASH-defined manifest file that the DASH client uses to discover the DASH media segments [20].

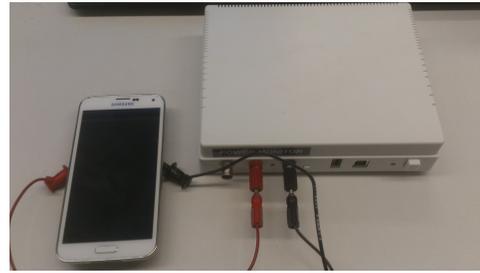


Figure 2: Measurement setup.

Figure 1 depicts a typical DASH video streaming system [22]. A DASH-based server breaks and divides a single video file into segments, which are shown as black boxes in the figure. During the process of video streaming, segments are delivered to clients in order and stored in the player/browser’s buffer. At the very beginning of video streaming, the client retrieves the MPD metadata first and extracts the content and segment information from it.

3. METHODOLOGY

We measure the power consumption of video streaming on the smartphones in real LTE networks to quantify the impact of DASH settings.

DASH Streaming Testbed. We first build a testbed as illustrated in Figure 2. A smart phone in 4G LTE network is connected to a Monsoon Power Monitor and works with the power provided by the power monitor. The DASH player is on the client side when the mobile phone is using online video streaming service in 4G LTE networks. This controlled testbed environment aims to minimize the interference from unknown external factors (e.g., proprietary segmenting) used by public video servers.

We choose DASH-IF, a popular open-source DASH reference player as our video player on the client side [1]. To be compatible with the DASH reference player, we deploy our own video server on the local machine using Jetty [2], which is also open source and supports both HTTP/1.1 and HTTP/2 protocols. The DASH-IF player is running on the client, and it can stream both global and local video sources with adaptive bit rates. To support video streaming, we partition video sources into DASH-supported segments in advance using MP4Box[5].

Measurement Methodology. We use an external Monsoon power monitor [4] to record the real-time power consumption of the test phone (Figure 2). We test with two phone models: Samsung Galaxy S5 and Note 3 over two US LTE carrier networks: T-Mobile and AT&T. We test with three video sources (10 minute, 5 minute, 3 minute) with different video qualities (360p, 720p and 1080p).

We observe the consistent results in such settings, and the results presented in this work are mainly for the S5 phone running a 10-minute video streaming over T-Mobile networks. We run tests at different locations where LTE speed and radio signal strength vary. At each location, we vary the settings of the segment length and the buffer size. Each scenario has 5 runs. We manually disable Adaptive Bitrate Streaming (ABR) in order to eliminate the impact of time-varying video rates and traffic patterns (varying traffic sizes in different runs). To this end, we fix the video source quality (360p, 720p or 1080p) each time.

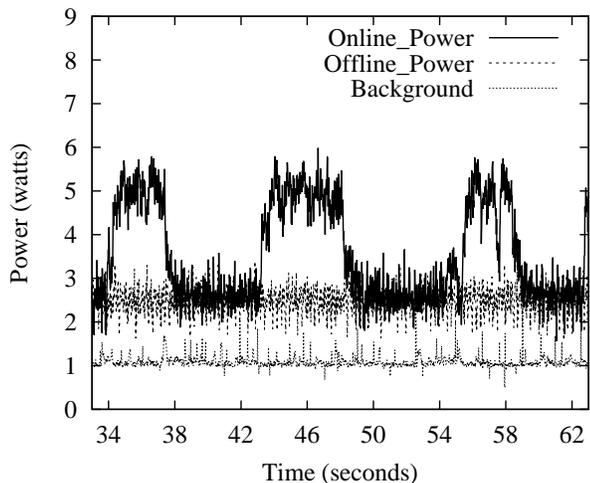


Figure 3: Power logs in case of video streaming (online) and playing (offline) and background (no video activity) mode.

To make traces comparable, we fix the other phone settings for screen brightness and volume, and disable the data usage for background services and other applications in all the tests. In particular, We use the highest brightness and mute voice and play all the videos in the Chrome browser.

In addition to the power consumption logs, we also collect network traffic traces through `tcpdump`. In the meantime, we also collect the cellular network events for further RRC analysis. We develop `MobileInsight` [3], an in-device tool to record RRC message exchanges and retrieve precise instantaneous RRC states. This works similar to `QXDM` [6], a professional cellular network debugging tool developed by Qualcomm.

LTE Power Extraction. In order to extract power consumption solely for video content transmission over LTE networks, we need to tackle the problem that the power consumption is measured as the whole. The power can be used to run basic operations (CPU, memory), play on the display (screen) and transmit video content over LTE network. To separate it, we run both the streaming test (online) and the playing-using-local-file (offline) test and extract network transmission power by comparing them. Figure 3 shows the power logs in the online video streaming, offline video playing and background modes. The gap between the online and offline mode roughly represent the power for video transmission over LTE. On the mobile client side, there are three major activities: (1) periodically or on-demand retrieving segments from the video server, (2) playing the video in a player embedded in a webpage, and (3) performing basic service to maintain and manage mobile network connectivity. The major difference between offline and online modes is that, the online streaming gradually downloads requested video segments during the playback while the offline downloads all the segments in advance and does not incur any transmission during the playback. We also observe that data transmission exhibits a clear active-idle pattern. We compare it with the `tcpdump` traffic trace and learn that data transmission stays active only in a small time window (e.g., 34–38 seconds) and moves to idle afterwards. The cycle matches with the client requests for new segments.

Res	Online(mW)	Offline(mW)	Nw(mW)
360p(2sec)	2208	1877	331 (100%)
720p(2sec)	2434	1913	521 (157.4%)
1080p(2sec)	2632	2013	619 (187.3%)
360p(4sec)	2127	1877	250 (100%)
720p(4sec)	2274	1913	361 (145.0%)
1080p(4sec)	2469	2013	456 (183.1%)

Table 1: Impact of video source quality (resolution) on energy consumption.

4. MEASUREMENT RESULTS

in this section, we now measure the impacts of various DASH-based video streaming factors.

4.1 Video Quality (Resolution)

We first look into the video quality. The video sources with higher resolution usually consume more power than the lower-resolution ones. It affects the power consumption of both video playback and network activities because the higher-resolution ones need to deliver and process/display more content. We test with various resolutions (360p, 720p, and 1080p) using two segment lengths (2s and 4s) and show the average power consumption for data transmission in Table 1. We use one video source called `Timer`, where a timer (of varying durations) is invokes and the video frame changes every second.

The average power of playing the 360p video source using 2-second segments is 2208 mW, where 1877 mW is contributed by video player and other local components, and the remaining 331 mW is consumed by network activities. The transmission power consumption grows as the resolution increases. For those using 2-second segments, the transmission power increases by 190mW (57.4%, 720p) and 371 mW (87.3%, 1080p) than the 360p one. It is similar for those using 4-second segments. The impact of segment is studied next and basically, the absolute power consumption decreases as the length expands. In the meanwhile, non transmission parts (i.e., video playing and other components, the offline power) also consume more power as the resolution grows. It goes up from 1877mW (360p) to 1913mW (720p) and 2013mW (1080p). In total, with the setting of 2-second segment length, the 1080p video consumes 16.1% more power on average than the 360p one. The local activities (offline) consume 7.2% more, and the network activities consume 87.3% more. The 720p video consumes 10.2% more power on average where The local activities consume 1.9% more, and the network activities consume 57.4% more. Figure 4a shows the overall comparison for different resolutions.

Finding 1: Streaming videos with higher resolution will consume more power in the video player, screen, and other player components. In addition, higher resolution also greatly brings up the power consumption of network activities.

4.2 Segment Length

To learn the impacts of different segment lengths, we create 3 different segment length settings (2 seconds, 3 seconds, and 4 seconds) and partition the same video source into different segment sets. We measure the power consumption of streaming the 720p version and the average power consumptions are 2434 mW, 2295 mW, and 2274 mW, respectively.

By subtracting the power consumption for offline activities, which is stable from total power consumption, we ob-

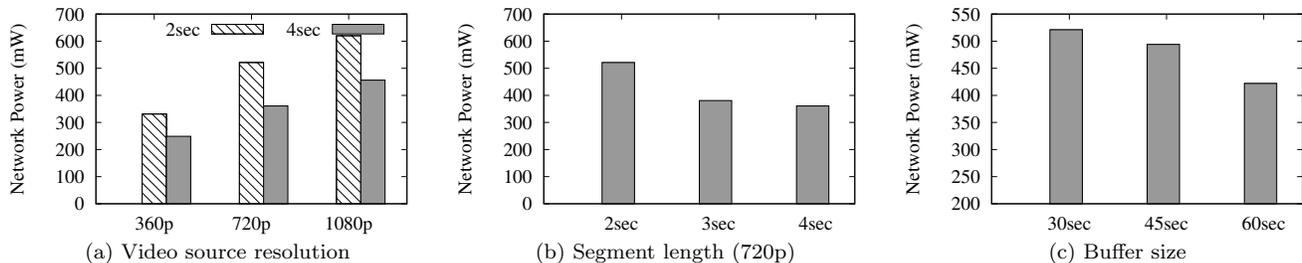


Figure 4: Power consumption under various conditions.

tain that the power consumption of network activities with 2sec segments is 521mW, with 3sec segments is 381 mW, with 4sec segments is 361 mW. For the 720p video source, utilizing 3-second segments can save 26.9% power in network activities, while utilizing 4-second segments can save 30.7%.

From Table 1, we can observe that the average network power consumption for 360p video with different segment lengths. In 360p video streaming, 4-second segments video can gain a saving of 24.8%. In 1080p video streaming, 4-second segments video can reach the saving of 26.5%, which is 164mW. It is obvious that, the traces with 4 second segment length consume lower power than the traces with 2 seconds segments and 3 seconds segments. The total energy gap for a single 600-second video trace between 2 seconds length and 4 seconds length is 48.6 J for 360p, 96 J for 720p, and 97.8 J for 1080p. Figure 4b shows the power comparison with different segment lengths for 720p. We observe that the power consumptions of video player and other components remain the same with different segment length, since the video content remains same. Therefore, the entire energy saving comes from the network activities. The increase in resolution leads to both higher overall and network power consumptions. With the trend of increasing video resolutions in today’s major video streaming services, the room for power savings keeps growing.

Finding 2: The power consumption of video streaming decreases as the video segment length grows. The additional power consumption for smaller video segments is entirely contributed by the network activities.

4.3 Buffer Size

Different buffer sizes also lead to different outcomes of the experiments. To fairly compare the traces with difference only in buffer sizes, we pick the video *Timer* with 720p resolution and 2 second segment length. By adjusting the buffer size parameter in the source code of DASH-IF reference player, we make the player retrieve more segments at one time. Our current approach increases the throughput in the first retrieve in the entire video, while the later deliveries maintain the same under different buffer settings. In the first retrieve, more video segments will be filled into the video buffer if the buffer has larger capacity. We believe the power and energy consumption can be further deducted if more video segments can be filled into buffer in each retrieve throughout the entire session.

In total we have three options for buffer size, including 30 second, 45 second, and 60 second. The 720p video consumes average power of 521 mW with 30 seconds buffer size. However, this number becomes 494 mW if we switch the buffer

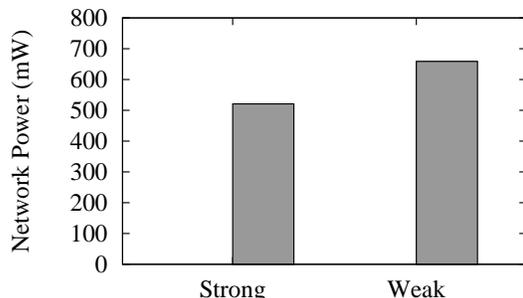


Figure 5: Power consumption under different signal strength.

size to 45 seconds, and 422 mW if we switch the buffer size to 60 seconds. It is obvious that the traces with larger buffer sizes consume less power than the traces with smaller buffer sizes. The DASH-IF Player will continuously accept new segments before its cache is filled. This procedure makes the transmission become more consistent instead of discrete video files transmitting. In this case, some TCP warm-ups, RRC tails, and some other meaningless high power consumption periods are reduced since less deliveries were made, and the entire energy saving emerges. Figure 4c shows the differences for diverse buffer size settings.

Finding 3: Increasing buffer size enables the video player to retrieve more video segments in a single receiving session, which can significantly reduce the power consumption of network activities.

4.4 Signal Strength

In order to see the impact of signal strength on power consumption of video streaming, we monitored and recorded the power and packet traces in video streaming under different network signal conditions.

We pick the video *Timer* with 720p resolution and 2-second segments, and played them in an area with strong LTE signal strength and an area with weak LTE signal coverage. Their radio signal strengths are about (-100, -95) dBm (strong) and (-118, -113) dBm (weak). Measured by the instant speed test, the transmission speed in the strong signal area was around 8 to 20 Mbps, while the speed is only around 3.0 to 3.8 Mbps in the weak signal area. According to our traces, the average power consumption of network activities is 659 mW when the streaming is under weak signal coverage, while it is only 521 mW when the signal is

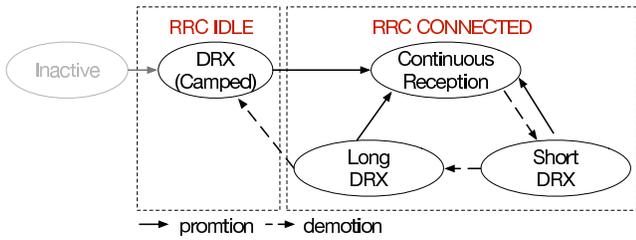


Figure 6: RRC State Transitions.

strong. Figure 5 shows the difference for strong and weak signal strength areas.

Finding 4: Signal strength is also an impact factor. The better the signal strength is, the lower the network power consumption will be.

5. ANALYSIS AND POTENTIAL GAIN

5.1 RRC states

By monitoring the RRC state of the smart phone during the process of video streaming, we see the power waste is closely connected to RRC state Transition.

As a well studied topic, RRC refers to Radio Resource Control protocol [8], is a well-designed mechanism to take charge of connection establishment and release functions. RRC protocol defines two RRC states for a device in the 4G LTE network, including *IDLE* and *CONNECTED*. The device automatically goes into *CONNECTED* when it starts to conduct mobile data transmission, while it will be switched back to *IDLE* once all the transmissions are terminated and no more packets will be delivered [7].

RRC state transitions. When in the *IDLE* state, the User Equipment(UE) is set in DRX (Discontinuous Reception) mode which obviously reduce the UE consuming power into a very low level. When UE is trying to send/receive one packet no matter how much the size is, the RRC state will be triggered into RRC connected state and corresponding UE power state should be promoted into a high level. It adds two sub-states: Short DRX and Long DRX with different periods for discontinuous reception.

Figure 6 shows how RRC states change in 4G LTE networks. Regardless of initial state, the RRC mode will be changed into continuous reception state instantly when data transmission activity begins. The only difference is the promotion time is different. The DRX mode may take more time than Long DRX and then short DRX (DRX > Long DRX > Short DRX). On contrary, RRC experiences demotion procedure from continuous reception to DRX mode when no data is delivered. It is controlled by different inactivity timers. It take a long path and is performed step by step. Here, CX → Short DRX → Long DRX → DRX. This is why one data transmission costs additional power to maintain the RRC state.

RRC impact. A real entire RRC state changing should have this state changing process: *Idle camped* → *Connecting* → *Connected* → *Closing* → *Idle Not Camped* → *Idle Camped*. Based on the preliminary study, our experiments should experience six states except IRAT To LTE Started state and Inactive State during the packet transmission.

Figure 7 has shown the actual RRC states in our online video streaming traces, from this we can see that the RRC

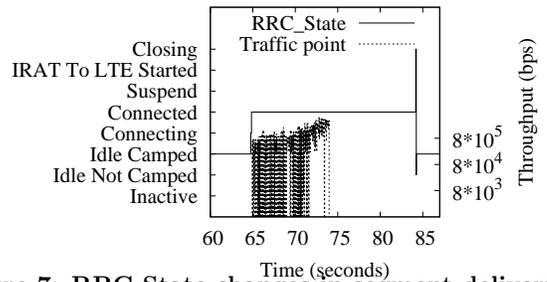


Figure 7: RRC State changes in segment delivery.

state does not return to *IDLE* right after the end of the packet delivery. Instead, it remains in a connected state for approximately 10 seconds before it goes to “Idle Camped” state, which consumes more power than idle state. This period is always defined as RRC tail period, which has been proved to be one of the sources of power wasting.

Every time when the DASH-IF player retrieves video segments to fulfill its buffer, the smart phone has to go into connected state, which brings TCP warm-up and other unnecessary energy wasting periods. After each retrieving, the smart phone has to go through the RRC Tail period, consumes more energy. In other words, fewer retrieving sessions cause fewer energy wasting procedures especially RRC Tail.

5.2 Potential Saving Space

Network power is the part that contains major potential space for energy saving, as the different settings in this section have created variations only in network parameters. From the tables and figures we can learn that the network power decreases with the positive growth in segment length, decreases with the increase in buffer size, decreases with the down grade of video resolution, and increases with the decrease in signal strength and network speed. Compared to the same video *Timer* with 2-second segments, video with 4-second segments can reduce the average network power by up to 30.71%, and save up to 97.8 J of energy in the entire video.

The video played in the player with buffer size of 60 seconds can leverage the average network power by 19% comparing to the player with 30-second buffer. If further modification applied to the DASH-IF reference player can make the buffer retrieve larger set of segments continuously, the reduction can be significantly larger.

In future, delicate designs of video streaming services under LTE may consider the above points in order to achieve better energy efficiency. To reduce the effect of RRC tail, TCP warm-up, and other power-consuming processes, the newly designed video player may merge the segment transmission. When the network condition is good enough, the client can be provided with options of larger segments. The DASH-IF video player can be equipped with larger buffer sizes, which can help the player to retrieve more segments in each single transaction. Based on these new functions, the video player can maintain a better balance between quality of experience and energy efficiency.

6. RELATED WORK

Recent years have witnessed active studies of video streaming in the literature to improve quality of experience (e.g., [15, 19, 16, 18, 13]) and energy efficiency (e.g., [17, 21]). Our work

differs from them in the scope and/or the methodology. Researchers seek to optimize streaming performance/experience over the Internet; They either adjust video rates adaptive to network conditions [15, 19, 13], leverage cooperation or CDN (caching) [16, 18] to improve network transmission quality; They do not target at specific 4G LTE networks. [17] measures the power consumption of various caching mechanisms in video streaming and it mainly considers 3G networks. [21] proposed to optimize power efficiency on mobile devices but it is based on models, not real measurements. Our work is to conduct real experiments to measure the impact of DASH video streaming settings on power consumption over 4G LTE networks. RRC power consumption has been also widely investigated [14, 10, 11]. These studies focus on the energy consumption of the RRC cycles but do not consider the impacts on video streaming.

7. CONCLUSIONS

Throughout this paper we profile and analyze power consumption of different video streaming strategies using DASH in LTE networks. From the measurement and analysis, we learn that the longer segment length and larger buffer size can reduce power consumption and eventually waste less energy. The RRC state should be one important factor that impacts the power. According to our statistics, the potential energy saving on network transmission can be more than 30%, which is significant for mobile devices in LTE networks.

Acknowledgments

We thank the anonymous reviewers for their constructive comments. This work is supported in part by gifts from Adobe Research at Adobe Systems Incorporated.

8. REFERENCES

- [1] DASH industry forum. <http://dashif.org/software/>.
- [2] Jetty - Servlet Engine and HTTP Server - Eclipse. <http://www.eclipse.org/jetty/>.
- [3] Mobileinsight project. http://metro.cs.ucla.edu/mobile_insight.
- [4] Monsoon power monitor. <https://www.msoon.com/LabEquipment/PowerMonitor/>.
- [5] MP4Box GPAC multimedia open source project. <https://gpac.wp.mines-telecom.fr/mp4box/>.
- [6] QxDM Professional Qualcomm eXtensible Diagnostic Monitor. <https://www.qualcomm.com/documents/qxdm-professional-qualcomm-extensible-diagnostic-monitor>.
- [7] 3GPP TR 25.813: Radio interface protocol aspects (V7.1.0), 2006.
- [8] UMTS RRC Protocol specification (version 12.4.0 Release 12), 2015.
- [9] Information technology - dynamic adaptive streaming over HTTP (DASH) - part 1: Media presentation description and segment formats, ISO/IEC MPEG standard, May 2014. ISO/IEC 23009-1:2014.
- [10] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani. Energy consumption in mobile phones: A measurement study and implications for network applications. In *IMC*, 2009.
- [11] C. S. Bontu and E. Illidge. Drx mechanism for power saving in lte. *Communications Magazine, IEEE*, 47(6):48–55, 2009.
- [12] Cisco Visual Networking Index. Global Mobile Data Traffic Forecast Update, 2014–2019, 2015.
- [13] H. E. Egilmez, S. Civanlar, and A. M. Tekalp. An optimization framework for QoS-enabled adaptive video streaming over OpenFlow networks. *IEEE Transactions on Multimedia*, 15(3):710–715, 2013.
- [14] 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 the 10th international conference on mobile systems, applications, and services*, pages 225–238. ACM, 2012.
- [15] T. Y. Huang, R. Johari, N. McKeown, M. Trunnell, and M. Watson. A buffer-based approach to rate adaptation: Evidence from a large video streaming service. *ACM SIGCOMM Computer Communication Review*, 44(4):187–198, 2014.
- [16] L. Keller, A. Le, B. Cici, H. Seferoglu, C. Fragouli, and A. Markopoulou. Microcast: cooperative video streaming on smartphones. In *Proceedings of the 10th international conference on mobile systems, applications, and services*, pages 57–70. ACM, 2012.
- [17] X. Li, M. Dong, Z. Ma, and F. C. A. Fernandes. Greentube: power optimization for mobile videostreaming via dynamic cache management. In *ACM International Conference on Multimedia*, pages 279–288, 2012.
- [18] M. K. Mukerjee, D. Naylor, J. Jiang, D. Han, S. Seshan, and H. Zhang. Practical, real-time centralized control for CDN-based live video delivery. In *Proceedings of the ACM conference on SIGCOMM*, pages 311–324, 2015.
- [19] H. Nam, K. H. Kim, D. Calin, and H. Schulzrinne. Youslow: a performance analysis tool for adaptive bitrate video streaming. In *Proceedings of the ACM conference on SIGCOMM*, pages 111–112, 2014.
- [20] V. Swaminathan. Are we in the middle of a video streaming revolution? *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 9(1s):40, 2013.
- [21] S. Wei, V. Swaminathan, and M. Xiao. Power efficient mobile video streaming using HTTP/2 server push. In *Proceedings of the 17th IEEE International Workshop on Multimedia Signal Processing (MMSP)*, pages 1–6, 2015.
- [22] N. Weil. The state of MPEG-DASH 2015. <http://www.streamingmedia.com/Articles/Editorial/Featured-Articles/The-State-of-MPEG-DASH-2015-102826.aspx>.
- [23] L. Zhan, D. M. Chiu, Y. Hua, and Z. Zhu. A measurement study of mobile video streaming by different types of devices. In *Proceedings of the 7th IEEE international conference on communication systems and networks (COMSNETS)*, pages 1–8, 2015.