

BingeOn Under the Microscope: Understanding T-Mobile’s Zero-Rating Implementation

Arash Molavi Kakhki
Northeastern University
arash@ccs.neu.edu

Fangfan Li
Northeastern University
li.fa@husky.neu.edu

David Choffnes
Northeastern University
choffnes@ccs.neu.edu

Alan Mislove
Northeastern University
amislove@ccs.neu.edu

Ethan Katz-Bassett
University of Southern California
ethan.kb@usc.edu

1. INTRODUCTION

The popularity of mobile devices for ubiquitous Internet access has led to exploding demand for relatively scarce cellular bandwidth. As a result, cellular operators increasingly manage their customers’ demand using various techniques such as traffic shaping [10], transcoding [11], and *zero-rating* [9]. With zero-rating, Internet Service Providers (ISPs) do not charge users for traffic sent to/from certain services, often because those services agree to use limited bandwidth resources.

Perhaps the most well-publicized case in the U.S. is T-Mobile’s “*BingeOn*” service, which zero-rates video streams from a large number of partner sites but limits the bandwidth they can use (i.e., rate limits). *BingeOn* works as follows [2, 7]: First, a user must participate in *BingeOn*, which is enabled by default. If the user explicitly opts-out of *BingeOn*, traffic is not zero-rated or rate-limited by T-Mobile. If *BingeOn* is enabled, video streaming may be affected based on whether the content provider is part of *BingeOn* (and meets its technical requirements), opted out, or neither. Traffic for providers that opt-in is zero-rated and rate-limited, traffic for providers that explicitly opt-out is neither zero-rated nor rate-limited, and traffic for providers that neither opt-in nor out is rate-limited but not zero-rated. The provider opt-out option was not present in the original *BingeOn* program, but was present in March [2].

BingeOn has been highly controversial, due to concerns over network neutrality, user confusion, and tech-

nical downsides for users [12]. The resulting debate has led the EFF to call for T-Mobile to abandon *BingeOn*, and generated several responses from T-Mobile’s CEO John Legere [1]. Importantly, there is little rigorous empirical data to inform the implications of zero-rating on network neutrality principles or customer-perceived Quality-of-Experience (QoE).

In this paper, we address this issue by conducting a study of T-Mobile’s zero-rating policy and implementation to understand its implications for users and content providers in terms of data quota, performance, and QoE. We focus on T-Mobile and *BingeOn* due to their recent prominence, but we believe that lessons learned from this exercise will readily apply to other carriers using similar technologies to implement their policies.

We leverage our prior work on differentiation detection [11] to develop a suite of controlled tests. We deploy these tests to multiple *BingeOn* (and non-*BingeOn*) T-Mobile devices in the U.S., and correlate our tests with measurements of the billing records from these accounts. As a result, we are able to make significant headway in understanding *BingeOn*. Overall, we make five key contributions.

First, we characterize how *BingeOn* differentially impacts participating and non-participating providers. We determine that *BingeOn* is implemented solely by rate-limiting specific flows to 1.5 Mbps (consistent with the EFF’s study [3]). Despite the fact that T-Mobile claims that *BingeOn* provides “optimized streaming”, there is no transcoding or optimization taking place.

Second, we show how these differences translate to QoE, and how this QoE may impact users. We find that with *BingeOn* enabled, non-partner video flows see the same rate-limit (even though users are charged for these degraded flows). As a result, video services that do not support multiple quality levels could cause users to have poor video QoE.

Third, we reverse-engineer the classifier used for enabling zero-rating, and show that *BingeOn* as imple-

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 the author(s) 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.

Internet-QoE, August 22-26 2016, Florianopolis, Brazil

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4425-8/16/08...\$15.00

DOI: <http://dx.doi.org/10.1145/2940136.2940140>

mented can have collateral damage by throttling flows that are not video content. We find that *BingeOn* is implemented using simplistic string matching, which can lead to false positives and negatives.

Fourth, we show that this implementation admits vulnerabilities such as free-riding. We demonstrate that the above implementation opens the door for abuse, as users can proxy their traffic with fake `Host` headers and receive zero-rating for arbitrary flows. We have responsibly disclosed this vulnerability to T-Mobile.

Fifth, we discuss how the above vulnerability is difficult to address using the current DPI-based approach, or alternative DPI approaches. This suggests the need for an alternative policy that can be completely and correctly implemented in existing middleboxes.

Together, our results show that the current *BingeOn* implementation can negatively impact both users and T-Mobile. If a user enrolls in *BingeOn*, all video traffic is currently rate-limited (even non-partners), and this can cause poor video quality for certain videos. Moreover, the implementation of *BingeOn* allows users to “steal” arbitrary amounts of data from T-Mobile.

Note that our results reflect the state of *BingeOn* in early March, 2016. The implementation details can (and have) changed over time, but our high-level findings hold at the time of writing (May, 2016).

2. METHODOLOGY


There are two parts to our methodology: characterizing the behavior of *BingeOn* and reverse-engineering its implementation. Specifically, to characterize *BingeOn*, we need controlled experiments that reveal how *BingeOn* is impacting network performance and QoE. After confirming that *BingeOn* impacts our test traffic, we seek to reverse-engineer its implementation to determine how T-Mobile determines that traffic should be classified as falling under *BingeOn* (and whether this approach is susceptible to subversion).

2.1 Characterization


We begin by describing our approach to characterizing the behavior of *BingeOn*. We aim to address three questions with our methodology:

1. *Does T-Mobile’s BingeOn give differential treatment to video traffic?* According to T-Mobile’s *BingeOn* documentation, “*BingeOn* optimizes video quality for smartphone screens. It provides a great DVD-quality experience (typically 480p or better) for all detectable video, which can minimize buffering and maximize quality while using a fraction of the data” [7].

To determine whether T-Mobile indeed gives differential treatment to video traffic, we adopt the methodology used in our prior study of traffic differentiation [11]. To summarize, we record traffic traces from arbitrary mobile apps, then replay those recorded byte streams between a replay client and replay server under our control. As such, the only difference at the application layer



(a) *BingeOn* Disabled



(b) *BingeOn* Enabled

Figure 1: CDF of throughput for Netflix traffic replayed with *BingeOn* disabled and enabled. When *BingeOn* is off, no rate limit is imposed on Exposed or VPN-tunneled flows. With *BingeOn* enabled, VPN-tunneled traffic evades T-Mobile’s detection and rate limits, while the Exposed replay is classified as Netflix and limited to 1.5Mbps.

between our recorded traces and replays is that our replay server usually does not have the same IP address as the server in the original recorded traces. As we demonstrated previously [11], not having the same IP address as the server does not affect DPI-based classification for all tested networks, including T-Mobile.

To determine if there is differentiation, we compare the throughput distribution for three types of replays: (1) the original packet payloads, (2) the original payloads encrypted via a VPN tunnel, and (3) the original payloads replaced with the same number of random bytes. If a DPI device is being used for classification, it should classify the exposed traffic (case 1) based on the application that generated it, and it should not be able to do so for encrypted or random bytes. To test whether this has an impact on performance, we conduct multiple rounds of tests in series and use statistical tests (a modified KS-test [11]) to identify differences in performance (i.e., throughput, latency) that indicate differentiation.

We conduct these tests separately with *BingeOn* enabled and disabled. We tested the apps listed in Table 3, which comprise popular streaming apps, some of which are participants in *BingeOn* [7].

2. *Is BingeOn traffic zero-rated?* Using the same methodology as described above, we determine whether the replayed traffic counts against the data quota for our testbed data plan. We do so using the self-service codes provided by T-Mobile for users to determine their data usage. Specifically, before we run each test, we determine the amount of data used, and we check again after the test. We observed the data usage counters

to be updated in real-time, and we also confirmed this with T-Mobile’s customer service.

3. *What is the impact on QoE?* To determine the impact of *BingeOn* on QoE, we use a popular video streaming service and gather commonly used QoE metrics such as video quality, quality changes, and rebuffering events. We conduct tests with *BingeOn* enabled and disabled, across several video bitrates.

2.2 Reverse Engineering

We investigated the *BingeOn* implementation with two central questions:

1. *How does BingeOn classify traffic?* Our prior differentiation work [11] suggested that DPI devices classify applications using regular expression matches on certain fields of HTTP requests and responses, and SNI fields in TLS handshakes. We thus use a modified version of the record/replay approach described in the previous section, where we modify various fields and determine their impact on *BingeOn* classification. We use two sources of ground truth to validate whether *BingeOn* is being triggered on our test traffic: whether the traffic is zero-rated and/or it receives differentiated service.

2. *Is BingeOn susceptible to subversion?* Based on our observations from reverse engineering the classifier, we develop techniques to “trick” the *BingeOn* classifier into thinking that an application not participating in *BingeOn* is subject to *BingeOn*. For this, we develop a custom traffic-rewriting proxy and test the result by checking whether traffic is zero-rated and differentiated.

2.3 Testbed and Dataset

We conduct our experiments using Android phones with T-Mobile SIM cards. Each card is equipped with a data plan consisting of 6 GB of high-speed (4G) data. At any time, one of the SIM cards has *BingeOn* disabled, and the other one has it enabled. When running experiments, both SIM cards are in the same location, and their devices are connected to the same cell tower and have good reception. For tests using tethering, we used USB to tether a Mac laptop to the Android phones.

All of the results presented in this paper are based on tests conducted in Boston, MA. In total we ran more than 2000 traffic replays and more than 400 video streams against real apps using T-Mobile network, which add up to more than 20GB of data. We conducted a small number of tests in Los Angeles, CA and found no difference in our results.

3. CHARACTERIZING BINGEON

In this section we quantify *BingeOn*’s impact on performance and video-streaming quality of experience.

3.1 Performance

Does *BingeOn* reduce video-streaming performance? To investigate this question, we begin us-

ing replays of Netflix traffic. Figure 1 shows individual examples of the impact of *BingeOn* on throughput. Each figure shows the cumulative distribution (CDF) of throughput samples for replays of Netflix traffic with *BingeOn* disabled (top) and enabled (bottom). The top figure shows that there are no significant visual differences between *VPN* and *Exposed* replays when *BingeOn* is disabled, but there is a clear difference when *BingeOn* is enabled in the bottom figure. We find that the average throughput for *Exposed* with *BingeOn* enabled is 1.5Mbps, about 80% lower than *Random* replays. Note that our results are consistent with EFF’s results [3].

We conducted 60 Netflix replays in the Boston area to build a large sample of data to characterize *BingeOn*. Table 1 shows the average and standard deviation of throughput, RTT, and retransmissions when replaying traffic with *BingeOn* enabled and disabled. We summarize our results below.

- *The BingeOn rate limit is ≈ 1.5 Mbps (Table 1).* This rate is sufficient for 480p videos on YouTube and “low quality” Netflix streaming, but is often below requirements for higher video quality (e.g., 720p for YouTube requires 1.5–4 Mbps [5] and Netflix SD quality video requires 3 Mbps [6]). Google’s VP9 codec supports 720p at bitrates lower than 1.5 Mbps, but the only smartphones that support it are select Android ones. Thus, T-Mobile’s claim of supporting “480p or better” [7] video quality is technically true, but in practice support for “better” depends on the device and video service.
- *The BingeOn infrastructure does not “optimize” video.* T-Mobile’s CEO claimed that “*BingeOn* includes a proprietary technology to not only detect the video stream, but select the appropriate bit rate to optimize to the mobile device” [1]. Our differentiation detection methodology trivially reveals whether an ISP modifies content, e.g., for optimization. We found no modification to our replay content, nor any evidence that *BingeOn* behavior changed in response to the device we used (smartphone, or laptop).
- *BingeOn is implemented using policing.* There is low jitter and high retransmission rates when *BingeOn* is enabled (Table 1). This indicates a token bucket with a small (or no) queue which results in packets being dropped when there are no tokens available.¹
- *BingeOn’s rate limit is cumulative for all flows from a single SIM.* If a customer streams simultaneous *BingeOn*-eligible videos (via tethering), the average throughput per stream will be lower than 1.5Mbps.

Is *BingeOn* traffic zero-rated while other video traffic is not? We address this question by investigating data plan consumption before and after each test we run to tell if the traffic was zero-rated (no charge against the data plan). T-Mobile customer service con-

¹Our replays do not adapt bitrates; rather, video traffic is replayed at the same bitrate it was recorded, which may be higher than supported by *BingeOn*.

	<i>BingeOn</i> disabled		<i>BingeOn</i> enabled	
	Exposed	Random	Exposed	Random
Throughput (Mbps)	7.707 (0.404)	7.231 (1.238)	1.540 (0.003)	7.205 (2.264)
RTT (ms)	206 (33)	201 (31)	126 (2)	141 (48)
Retransmission rate (%)	0.099 (0.093)	0.121 (0.1217)	18.158 (0.732)	0.272 (0.135)

Table 1: Netflix’s average (std) throughput, RTT, and number of retransmissions over T-Mobile’s network with and without *BingeOn*. When *BingeOn* is enabled flows are consistently throttled to 1.5Mbps. When *BingeOn* is disabled, the average rate is higher and each replay varies based on available bandwidth and signal strength. The low RTT and high retransmission rates when *BingeOn* is enabled suggests that *BingeOn* uses policing for rate-limiting.

Quality	Time to Start (sec)		video loaded in 1 min (%)		#rebuffers		Buffer/Play Time (%)	
	<i>BingeOn</i> On	<i>BingeOn</i> Off	<i>BingeOn</i> On	<i>BingeOn</i> Off	<i>BingeOn</i> On	<i>BingeOn</i> Off	<i>BingeOn</i> On	<i>BingeOn</i> Off
auto*	1.54	1.32	4	3	0.0	0.0	2.48	2.15
hd1080	4.64	2.6	1	3	7.22	0.11	43.63	4.2
hd720	2.76	1.64	2	5	2.5	0.0	7.44	2.65
large	1.76	1.32	3	8	0.0	0.0	2.82	2.19
medium	1.58	1.08	4	12	0.0	0.0	2.56	1.76
small	1.26	1.0	7	15	0.0	0.0	2.07	1.63

Table 2: QoE metrics for a YouTube video in different qualities, averaged over 10 runs. *Auto selects hd1080 with *BingeOn* disabled and medium (360p) with *BingeOn* enabled.

firmed that that “*BingeOn* services are whitelisted in their data usage counters,” and that the counters were updated in real time.

We found that *BingeOn* traffic is typically zero-rated while traffic from services not participating in *BingeOn* was charged. When we tested with replays of video traffic from providers not participating in *BingeOn*, the data used was charged against our data plan. Thus, at the time of our experiments, YouTube traffic was throttled to 1.5Mbps and we were charged for the data, whereas Netflix was similarly throttled and there was no data charged.² Such differential treatment of video services calls into question the policy’s legality in the face of the FCC’s Open Internet Order.

We found that *BingeOn* behavior is not entirely consistent over time. We encountered a small number of cases where a *BingeOn*-participating service’s traffic was not zero-rated. These cases were transient, suggesting they are due to reasons such as buggy or overloaded infrastructure that supports *BingeOn*.

We conducted the same experiments for other *BingeOn* participants (e.g., HBOGo, ShowTime, and Hulu) and non-participants (e.g., Vimeo and Veoh), and found identical results. We also recorded and replayed other types of traffic (e.g., image download), and observed no rate limits imposed on them.

3.2 Video Quality of Experience (QoE)

We now investigate how the policies revealed in the previous section affect video-streaming QoE. We developed a tool that uses the YouTube iFrame Player API [4] to open an hour-long YouTube video, play it for 60 seconds and log the following QoE metrics: time to start the video, video quality, quality changes, re-

²Note that list of services participating in *BingeOn* changes over time (e.g., YouTube joined *BingeOn* just before submission of this paper), but we confirmed this behavior is true for every tested video service not participating in *BingeOn*.

buffering events, and fraction of video loaded. We use this tool to play YouTube videos with different bitrates, and with *BingeOn* enabled and disabled. We picked YouTube because it offers a large variety of video qualities, it supports HTTPS which prevents in-network caching, and it supports a rich API for gathering QoE metrics. We believe our results apply to other video services because the previous section showed *BingeOn* applies the same policy to all video traffic.

Table 2 summarizes our findings.³ As quality increases, QoE metrics degrade, leading to hd720 and higher qualities becoming unwatchable with *BingeOn* enabled. In contrast, the hd1080 quality can stream with good QoE when *BingeOn* is disabled.

BingeOn provides sufficient bandwidth for 480p and lower, but it increases download times, which may increase battery consumption by keeping the radio in an active state for longer. Interestingly, when *BingeOn* is enabled, YouTube selects medium (360p) quality, lower than the 480p specified by T-Mobile.

The lower bitrate occurs independent of the device screen size. For example, YouTube attempts to stream HD for tablets, but T-Mobile’s rate-limiting forces YouTube to adapt to lower qualities that result in visibly low resolution on large screens.

4. DISSECTING BINGEON

We now investigate how T-Mobile detects *BingeOn*-eligible traffic (§ 4.1) and how it can be exploited to zero-rate arbitrary traffic (§ 4.2).

4.1 Reverse Engineering Classification

T-Mobile claims that *BingeOn* includes a “proprietary technology” to detect video streams [1]. To understand which features of the traffic are used to trigger

³We omitted 144p, hd1440, and hd2160 from the table as they follow the same trend.

Application	Detection criteria		How to evade detection?	
	<i>BingeOn</i> /Music Freedom*	Video	<i>BingeOn</i>	Video
Netflix	Specific GET arguments and the term “Netflix”	Same as <i>BingeOn</i>	Randomize GET argument and “Netflix” in response	Same as <i>BingeOn</i>
HBOGo	Host header ends with “hbogo.com”	Content-Type header (video/mp2t)	Randomize Host header	Randomize Content-Type header
ShowTime	Host header ends with “showvodhls.edgesuite.net**”	Content-Type header (video/mp2t)	Randomize Host header	Randomize Content-Type header
Hulu	Host header ends with “hulu.com”	Content-Type header (video/mp2t)	Randomize Host header	Randomize Content-Type header
Amazon Video	Host header ends with “amazonvod.loris.llnwd.net***”	Content-Type header (video/mp2t)	Randomize Host header	Randomize Content-Type header
Veoh	Not part of <i>BingeOn</i>	Content-Type header (video/mp4)	N/A	Randomize Content-Type header
Vimeo	Not part of <i>BingeOn</i>	Unknown***	N/A	N/A
YouTube (HTTP)	Host header ends with “googlevideo.com”	Content-Type header application/octet-stream	Randomize Host header	Randomize Content-Type header
YouTube (HTTPS)	Server name in the SNI ends with “googlevideo.com”	Same as <i>BingeOn</i>	Randomize “googlevideo” in the SNI	Same as <i>BingeOn</i>
Spotify*	“Spotify” in Host and User-Agent headers	N/A	Randomize the term “Spotify”	N/A
Pandora*	Host header ends with “p-cdn.com”	N/A	Randomize Host header	N/A

Table 3: Summary of how each app is detected by T-Mobile, and what changes can evade detection. If a flow matches a *BingeOn* app, the traffic will be zero-rated and throttled. Otherwise, it will check for video signatures and throttle such traffic (but not zero-rate) if it matches. **Music Freedom*, is a program similar to *BingeOn*, which zero-rates music streaming apps [8]. *edgesuite.net* (run by Akamai) and *LLNWD.net* are CDNs, serving ShowTime and Amazon videos. It is possible that these providers use other servers/CDNs too, hence different host names exist that result in zero-rating. ***We did not observe rate-limiting, which we use to reverse engineer Video detection criteria. We suspect that *BingeOn* classifiers are not properly configured for Vimeo.**

detection, we use our testbed to replay a recorded trace multiple times, each time modifying different portions of traffic content (e.g., change the Host HTTP header) to observe its effect on classification. Our goal is to exhaustively understand how popular video services are detected, but we do not cover all video services.

To detect the impact of traffic contents on detection as *BingeOn*-eligible traffic, we enable *BingeOn* for a SIM card and log the data consumption and average throughput of each replay. Based on our analysis in the previous section, we detect the following policies:

- ***BingeOn*-eligible:** This applies to our replay if it is zero-rated (data consumption does not increase) and its average throughput is 1.5 Mbps.
- ***BingeOn*-ineligible video:** The replay is *not* zero-rated, but has an average throughput of 1.5 Mbps.
- **Not video:** Traffic is not zero-rated or throttled.

Table 3 summarizes our findings for several popular video streaming services. A key take-away is that *BingeOn* uses a DPI device that matches regular expressions to detect video and *BingeOn*-eligible traffic. As we discuss below, this opens T-Mobile to free-riding, and means that their policy can easily be erroneously applied to traffic. Our key findings are as follows:

- ***BingeOn* uses regular expressions to match on Host, Content-Type, and SNI fields.** The left two columns in Table 3 show that *BingeOn* uses

simple regular-expression string matching based on pre-defined signatures. Classification first prioritizes the signature for *BingeOn* apps; if there is no app-specific match then *BingeOn* uses Content-Type signatures to detect non-*BingeOn* video streams.

- ***BingeOn*’s string-matching is brittle.** Once T-Mobile successfully matches a *BingeOn*-specific string in its classifier, it ignores all other fields that might support *or contradict* the classification. As we will show in §4.2, this will allow exploitation of *BingeOn* to zero-rate arbitrary traffic.
- **IP addresses and port numbers have no impact on classification.** Our replay experiments use different IPs from the recorded services, and we also experimented with changing port numbers. We find that classification does not depend on the port, e.g., a request with Host: hbogo.com and a video Content-Type to port 55555 (as opposed to the standard 80) on our replay server is still detected by T-Mobile as HBOGo.
- ***BingeOn* uses app-specific strings to throttle without zero-rating.** The “YouTube (HTTPS)” row in Table 3 shows that *BingeOn*-ineligible video streaming apps⁴ are specifically targeted for throttling, in addition of general rules for all videos.

⁴Before YouTube was added to *BingeOn*.




Figure 2: Subverting *BingeOn*: a) a local proxy copies the Host header into an optional parameter and overrides it with a *BingeOn*-enabled domain b) T-Mobile incorrectly classifies the traffic and zero-rates it c) an external proxy reverts the changes and forwards the request to destination.

4.2 Subverting BingeOn

The previous section indicates that *BingeOn* traffic is classified strictly using simple text matching, which suggests that *we can exploit BingeOn to free-ride on T-Mobile* by modifying arbitrary HTTP traffic to masquerade as *BingeOn*-enabled (and thus zero-rated) activity. *We built a proxy that does this, and confirmed that it allows free-riding.* Our current implementation uses two proxies (Fig. 2): a local proxy on the user’s device, and a proxy sitting outside T-Mobile’s network (e.g., in a cloud data center).⁵ The proxy works as follows. First, the local proxy stores the `Host` header in an `X-Host` header, then rewrites the `Host` header with a *BingeOn*-enabled host (e.g., `hbogo.com`) and forwards the request to a proxy located outside of T-Mobile’s network. This causes T-Mobile’s classifier to detect the traffic as *BingeOn*-enabled and zero-rate it. Next, the proxy outside of T-Mobile reverts the local proxy’s changes and forwards the request to the final destination. Note that at the time of writing there is no need to modify any of the traffic in the reverse direction.

Free-riding on T-Mobile violates their terms of service. We notified T-Mobile of the vulnerability, and they are currently mitigating it via abuse monitoring.

5. DISCUSSION AND CONCLUSION

This paper provided a detailed look at zero-rating in a popular US carrier, both to characterize its impact on traffic as well as understand how the policy is implemented. We discovered a variety of behaviors that have important implications for users and regulators. To summarize, by default *BingeOn* throttles all video traffic but charges for video from services not participating in *BingeOn*, there are no video- or screen-specific optimizations being used, and this policy can have a negative impact on video-streaming QoE metrics.

Further, we identified how T-Mobile classified traffic for *BingeOn*, and found that its regular-expression-based approach is brittle, potentially inaccurate, and easily subverted to free-ride on T-Mobile. We have notified T-Mobile of these issues, but we believe that fixing such vulnerabilities is difficult to do using the current DPI-based approach. Namely, if T-Mobile changes their regular expressions to match the ones our proxy uses, we could easily use new ways of encoding our original header information. Further, if they use a static

list of *BingeOn*-enabled apps’ server IP addresses, they will not be able to easily account for dynamically redirected IPs commonly used by video-streaming CDNs. Last, more detailed traffic analysis will require more DPI resources, perhaps reducing the cost-effectiveness of any solution. We believe that the above cat-and-mouse game, if played out over time, will become disproportionately complex and expensive for T-Mobile. This issue arises because the simplistic DPI-based approach to traffic classification *does not match T-Mobile’s BingeOn policy.* We believe the only way to prevent subversion is to specify a policy that can be *completely and correctly implemented* in existing middleboxes.

6. REFERENCES

- [1] <https://twitter.com/JohnLegere/status/685201130427531264>.
- [2] Content provider technical requirements for binge on. <http://t-mobile.com/content/dam/tmo/en-g/pdf/BingeOn-Video-Technical-Criteria-March-2016.pdf>.
- [3] EFF Confirms: T-Mobile’s Binge On Optimization is Just Throttling, Applies Indiscriminately to All Video. <https://eff.org/deeplinks/2016/01/eff-confirms-t-mobiles-bingeon-optimization-just-throttling-applies>.
- [4] YouTube IFrame player API. https://developers.google.com/youtube/iframe_api_reference.
- [5] Live encoder settings, bitrates and resolutions. <https://support.google.com/youtube/answer/2853702>, Mar. 2016.
- [6] Netflix internet connection speed recommendations. <https://help.netflix.com/en/node/306>, Mar. 2016.
- [7] T-Mobile Binge On. <http://www.t-mobile.com/offer/binge-on-streaming-video.html>, Mar. 2016.
- [8] T-Mobile Music Freedom. <http://www.t-mobile.com/offer/free-music-streaming.html>, Mar. 2016.
- [9] Zero rating: What it is and why you should care. <https://www.eff.org/deeplinks/2016/02/zero-rating-what-it-is-why-you-should-care>, Feb. 2016.
- [10] 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 *Proc. of SIGCOMM*, 2016.
- [11] A. Kakhki, A. Razaghpanah, A. Li, H. Koo, R. Golani, D. Choffnes, P. Gill, and A. Mislove. Identifying traffic differentiation in mobile networks. In *Proc. of IMC*, 2015.
- [12] B. Schewick. T-Mobile’s Binge On violates key net neutrality principles. *Stanford Law School*, 2016.

⁵A colluding HTTP server could replace the second proxy.