A Campus-Level View of Netflix and Twitch: Characterization and Performance Implications

Michel Laterman Martin Arlitt Carey Williamson
University of Calgary

***Abstract*—Video streaming is a major contributor to modern Internet traffic. In this paper, we use data collected from a campus edge network to characterize two popular video streaming services: Netflix and Twitch. While these video streaming services provide inherently different types of content, namely video-on-demand and live-streaming, they nonetheless exhibit many similarities in traffic patterns, protocol usage, content popularity, and growth. We identify seven similarities and differences, and discuss how these could be leveraged to improve streaming video content delivery on the Internet in the future.**

I. Introduction

Video streaming sites have experienced tremendous growth within the past few years, and this growth is expected to continue into the foreseeable future [5]. In fact, media streaming, specifically for video, is the largest category (by byte volume) for incoming Internet content at our university.

In this study, we focus on characterizing Netflix [13] and Twitch [18] usage by our campus community of over 30,000 students, faculty, and staff. Both of these streaming sites tend to generate long-duration high-bandwidth sessions, and serve enough content to rival some of the larger broadcast television networks [8]. For example, Netflix already serves more traffic than two of the four major US television networks, while Twitch is projected to be among the top 25 networks [8].

At our university, the three most accessed video services (in terms of bytes received) are YouTube, Netflix, and Twitch. YouTube is a video streaming service for user-generated content, specializing in short videos that are several minutes in duration. Netflix's catalog specializes in TV shows and movies, including well-known syndicated TV series as well as online-only Web content produced by Netflix itself. Netflix charges a monthly subscription fee for unlimited access to its content. Twitch is a site that focuses on the live-streaming of video games being played by professional video gamers.

For our study, we collected information about Netflix and Twitch traffic for a five-month period, spanning from December 2014 through April 2015. This time span includes an entire academic semester (January to April), as well as the month prior. Our dataset provides a snapshot of aggregate usage of these two video streaming services by our campus community.

At the time of our study, both Netflix and Twitch used unencrypted HTTP, facilitating our traffic analysis of URLs to identify video content information. Since mid-2015, however, Netflix has transitioned to a new Web interface and Transport Level Security (TLS) using Secure HTTP (HTTPS). As such, our study provides a "last look" at Netflix prior to its transition to end-to-end encryption. In our work, we ignore YouTube traffic, since it is already well-characterized in the literature [4], [6], and already uses HTTPS. We note, however, that Netflix's traffic volume on our network is already commensurate with that of YouTube.

The research questions behind our study are the following:

- How much network traffic is from Netflix and Twitch?
- How are these two video services similar and/or different?
- What are the performance implications of video streaming traffic on the campus network?

This study provides three main contributions. First, we characterize the content access patterns for Netflix and Twitch. Our dataset provides a final look at Netflix traffic before encryption, and (to the best of our knowledge) is the first Twitch study using a network-level dataset. Second, we characterize the connections and responses used to deliver video traffic. Third, we identify several characteristics that appear similar for both Netflix and Twitch. They are summarized in Table I, and explained in detail in subsequent sections.

Our measurement results are of value to network operators, protocol designers, and content providers. The results can be used by network operators to plan for future resource allocation, and help protocol designers with improving video streaming in the future. Service providers want to improve quality of service for popular applications, while reducing operational costs. Our campus-level study provides a glimpse of possible future demands for streaming on enterprise, ISP, and mobile networks, and constructive advice on how to handle such future traffic growth.

The rest of this paper is organized as follows. Section II discusses prior related work. Section III presents our data collection and characterization methodology. Section IV describes our overall traffic for the two services. Section V studies content-access patterns. Section VI examines video streaming protocol usage, focusing on connection usage and response characteristics. Section VII discusses performance implications of this work. Section VIII concludes the paper.

II. Related Work

There are numerous previous studies on Web and video traffic on the Internet. These studies span YouTube, Hulu, Netflix, Vimeo, and many other video service providers.

Borghol *et al.* [3] conducted a large-scale study of several different video content sites on the Internet. The primary focus in their study was characterizing the popularity of individual videos, and modeling the rise and fall in popularity over time.

TABLE I
MAIN SIMILARITIES AND DIFFERENCES OBSERVED BETWEEN NETFLIX AND TWITCH

Characteristic	Similarities	Differences	Section
Traffic Volume	Both services are high volume and continue to grow.	Netflix traffic volume is 5-10x higher than Twitch.	IV
Access Patterns	Both services show strong diurnal traffic patterns.	Netflix has two daily peaks, while Twitch has only one.	IV-B
Platforms	Content is accessed from diverse platforms and browsers.	Twitch access is primarily from Windows desktops.	IV, VI-B
Mobile Devices	Mobile devices are used to access both services.	Netflix has 40% mobile devices, while Twitch has 10%.	IV, VI-B
Content Popularity	Access is heavily skewed toward popular content.	Twitch has greater volatility in its content popularity.	V
Connection Usage	Both services use multiple connections to transport content.	Twitch only uses multiple connections at start of session.	VI-A
Responses	Both use DASH as a basis for content delivery.	Twitch has faster response times than Netflix.	VI-B

Xu *et al.* [19] characterized home networks in 2014. Their study provided evidence of the popularity of two video services, namely YouTube and Netflix. They also found that there are strong diurnal patterns on home networks.

Adhikari *et al.* [1], [2] studied how Netflix connects to clients across the United States, using data traces from 2011. Many technical details, including hostnames, CDNs, and usage of Silverlight, have changed since the study was published.

Martin *et al.* [11] conducted a study of Netflix in 2013. Certain details about Netflix's infrastructure have changed since then. For example, Netflix was using third-party CDNs to deliver video traffic; we did not observe the same CDNs. They also found that Netflix's implementation of DASH (Dynamic Adaptive Streaming over HTTP) defaults to TCP congestion control under heavy network traffic.

Summers *et al.* [17] used server logs to characterize Netflix traffic, in an effort to understand and improve server-side performance. A main emphasis was on pre-fetching video segments, and determining a good prefetch size based on chunk size, streaming bit rate, and network characteristics. As part of their work, they also studied the startup behavior of Netflix streams, which use multiple connections to determine suitable quality levels for adaptation [17].

Zhang and Liu [20] studied the characteristics of Twitch traffic. They used the Twitch API to crawl Twitch in the fall of 2014. The authors noted strong diurnal patterns with viewership, and found that most viewers watch from a desktop as opposed to a console device such as an XBox or Playstation. When examining the streamers themselves, they observed that about 1% of the streamers accounted for 70% of the views.

There have been several studies of social aspects of Twitch. Hamilton *et al.* [7] presented a general overview using streams of many different sizes, while Nascimento *et al.* [12] focused on streamers involved in electronic sports (eSports) for the game StarCraft 2. Their observations were based on data collected using the Twitch API, and interactions observed in Twitch chat. They found that viewers exhibited different behaviors, such as channel surfing and early exit. Kaytoue *et al.* [9] found that many streams (41%) originate on the west coast of North America, 19% on the east coast, and the rest mostly from Europe or south-east Asia. They also observed fluctuations in game popularity. These fluctuations occurred when a new game was released, with new games often receiving a surge of popularity.

While there are many previous studies on video traffic analysis, we believe that we are the first to provide detailed network-level comparisons between Netflix and Twitch.

III. METHODOLOGY

A. Data Collection

Our data was collected from a mirrored stream of all traffic that passes through the university's edge router. We can observe all traffic that has one endpoint in the campus network and the other in the Internet. Because our monitoring infrastructure is set up for long-term data collection, we do not record any packet-level or payload information; only connection-level traffic summaries are produced, using scripts that process packets on the fly.

We used the Bro network security monitor [14] to observe traffic on our network from December 2014 through April 2015. The Bro connection logs are used to study the network-level characteristics, and Bro's HTTP logs are used to quantify application-level characteristics. The connection logs list general information about each observed (TCP or UDP) connection (e.g., start time, endpoints, bytes/packets transferred by each endpoint, duration, and termination state). The HTTP logs contain information about each HTTP request-response pair, with information such as start time, endpoints, request/response body length, domain, path, referer, etc. We extended Bro's default behavior to collect extra information about HTTP request-response transactions, such as start and end times for requests and responses, pipelining, caching-related headers, and response type.

There are several limitations to our data collection. First, we do not record any cookies or other user-identifying information, and thus are unable to track sessions or users. Second, we do not record meta-data about media file names, types, formats, or resolutions, and thus cannot analyze video bit rates. Nonetheless, the Bro logs provide a valuable summary of Netflix and Twitch traffic.

B. Netflix Information

Netflix is a globally popular video-on-demand streaming site with over 80 million subscribers [13], [17].

At the time of our study, the structure of a typical Netflix session was as follows. Upon visiting `http://www.netflix.com` for the first time, Netflix responded with an HTTP redirect (301) to `https://www.netflix.com`, from which a subsequent redirection to a country-specific Netflix server may be required to handle content geo-restrictions. Next, Netflix processed login authentication over HTTPS. After logging in through HTTPS, Netflix reverted back to unencrypted HTTP for communication.

After logging in, the client was redirected to the Web Interface home (`WiHome`) to select the user's profile. On the Web interface home page, there was a menu with Netflix-suggested content for the user, including recently added Netflix content, content the user had not finished watching, and content recommendations based on prior viewing patterns.

Upon selection of an item from the menu, the browser sent an HTTP request of the form `www.netflix.com/WiPlayer?movieid=<id>...` that resulted in a JavaScript player being loaded. Content was then transported with a different set of HTTP requests over one or more TCP connections.

The `movieid` in the URL was an essential item for our analysis. It uniquely identified content on the Netflix server, whether it is a movie, a TV show, or a specific episode within a TV series. We used this identifier to track content popularity and byte volume in our traffic analysis. (In June 2015, the Netflix Web interface changed, along with the semantics of the `movieid`, making it context-dependent. Furthermore, the `movieid` attribute is no longer visible under HTTPS.)

From our data collection vantage point, we observed five CIDR subnets being used for Netflix content delivery: `108.175.32.0/20`, `198.45.48.0/20`, `198.38.96.0/19`, `23.246.0.0/18`, and `192.173.64.0/18`. Netflix owns additional IP address ranges, but no traffic was observed on these at our site. Other domains involved when visiting Netflix include CDNs operated by Netflix and by third parties to load thumbnail images (e.g., movie/series covers, still frames).

C. Twitch Information

Twitch is a subscription-based live-streaming site for video game play [18]. Users can watch professional game players in action, accompanied by audio commentary or analysis of the game play. Popular streamers can partner with Twitch to monetize their efforts, by partaking in tournaments, offering multi-player invitations, embedding advertisements into streams, and promoting particular games.

At the time of our study, the Twitch homepage showcased one of the featured live streams in the middle of the page, with a brief description of the stream to the right. Directly beneath the stream was a short icon list of featured streams, and further down the page was a list of featured games.

Once logged in, a user requested a specific stream. The request path when accessing such a page included the username of the streamer. An example URL was `www.twitch.tv/ddrjake`. The Web page had information about the stream itself, such as the title, the game being played, the streamer, and the streamer's avatar picture. Some pages had multiple media streams (e.g., game screen, Web camera on streamer, audio channel). User interaction with the Twitch site was handled using Flash. (In July 2015, Twitch transitioned to an HTML5-based video player with underlying Flash content.)

We observed Twitch video content originating from two different domains owned by Twitch: `twitch.tv` and `ttvnw.net`. From December 1, 2014 through March 16, 2015, video traffic was delivered primarily by `twitch.tv`, but from March 16 until the end of our collection period in April 2015, `ttvnw.net` was used. Other domains owned by Twitch, such as `jtvnw.net` and `justin.tv`, were used by Twitch to deliver other elements, such as static documents. Additionally, `jtvnw.net` had a CDN domain for serving images for Twitch. Almost all video content from Twitch (from `twitch.tv` or `ttvnw.net`) came from servers running Apple's HTTP Live-Streaming (HLS) service. HLS is an implementation of the DASH protocol.

D. Traffic Analysis

Using the Bro logs, we characterize the similarities and differences between Netflix and Twitch. Table I shows the main properties that we focus on, namely data volume, traffic patterns, platforms and mobile devices, content access patterns, connection characteristics, and response characteristics.

The rationale for selecting these characteristics is as follows. Data volume is used to show the overall levels of traffic for the services. Usage patterns allow us to see when video streaming services are used. Examining browser and mobile device usage is of interest to understand user preferences. Content access patterns offer insight into what users are viewing. Connection characteristics show how the video-streaming protocol influences the network traffic. Response characteristics are application-level properties that show the differences between on-demand and live-streaming content.

IV. High-Level Traffic Characterization

This section addresses our first research question, regarding the prevalence of Netflix and Twitch traffic on our campus network. We focus on traffic volume, diurnal patterns, as well as platforms and browsers used. Similarities and differences between the two video services are also highlighted.

A. Traffic Volume

Over the five-month period, we observed that 91% of the inbound campus traffic was TCP. Together, HTTP and HTTPS accounted for 88% of the inbound TCP traffic (1.40 PB).

Within the HTTP and HTTPS traffic, YouTube served 239.26 TB, Netflix served 217.15 TB, and Twitch served 19.49 TB. YouTube's traffic is encrypted, so we are unable to characterize the content-level details, and thus ignore it for this study. Netflix and Twitch were the largest (by volume) unencrypted video services accessed from the university network. We observed 305 million HTTP request-response transactions to Netflix on 14.3 million TCP connections. Twitch traffic involved 54 million HTTP request-response transactions on 1.6 million TCP connections. The video traffic generated by these two services accounts for much of the inbound data traffic during peak usage periods.

B. Diurnal Patterns

Figure 1 shows a typical week of traffic for Netflix and Twitch. Note that the scales for the two plots differ since Netflix's traffic levels are much higher than Twitch's.

Netflix and Twitch both exhibit the typical diurnal patterns associated with human-generated traffic, corresponding to when the majority of people are on campus. The busy period starts in the late morning, with usage peaking mid-day, and continuing into the late evening. The "light" period starts late in the night and lasts until the early morning. Traffic levels are lower on weekends and during university holiday breaks.

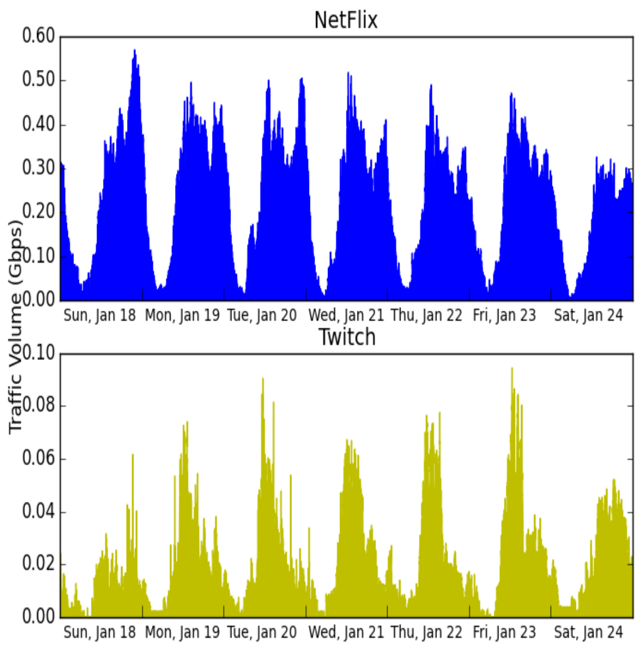

Fig. 1. Netflix and Twitch Weekly Traffic (January 18-24, 2015)

Figure 1 shows that Netflix often has two distinct peaks per day, with one in the early afternoon and one in the evening. The latter is not particularly surprising, since movie-viewing is often an evening activity in the student residences. Twitch, on the other hand, only peaks in the afternoon, and does not transmit a lot of traffic at night. This may be due to Twitch's live-stream nature. That is, content on Twitch is only available when there are streamers active. This nightly drop-off suggests that popular streamers viewed from campus are somewhere in North America; this would be consistent with the global streamer locations reported by Kaytoue *et al.* [9].

C. Platforms and Browsers

We analyzed the user agents reported in HTTP request headers to identify the platforms and browsers used. Nearly 40% of the total requests for Netflix were made by mobile devices. The total volume of mobile video content from Netflix was 54.01 TB, while desktop video content was 162.6 TB. For Twitch, mobile devices made up only 10% of requests.

Table II shows a breakdown of desktop and mobile agents. Netflix requests made with an empty user-agent string are

TABLE II
USER-AGENT PLATFORM SUMMARY FOR NETFLIX AND TWITCH

Netflix			Twitch		
Type	Platform	Requests	Type	Platform	Requests
Desktop (59.7%)	Macintosh	35.3%	Desktop (84.8%)	Windows	76.2%
	Windows	24.4%		Macintosh	8.6%
Mobile (39.2%)	Android	26.8%	Mobile (10.6%)	Android	5.0%
	iPad	7.1%		iPhone	3.8%
	iPhone	4.6%		iPad	1.8%
	ios-app	0.5%			
	ChromeOS	0.2%			
Other/Unknown		1.1%	Other/Unknown		4.6%

TABLE III
USER-AGENT BROWSER SUMMARY FOR NETFLIX AND TWITCH

Netflix			Twitch		
OS	Browser	Reqs	OS	Browser	Reqs
Macintosh (35.5%)	Safari	17.6%	Windows (76.4%)	Chrome	66.7%
	Chrome	15.1%		Firefox	7.7%
	Firefox	2.0%		IE	1.5%
	Other	0.8%		Other	0.5%
Windows (24.6%)	Chrome	18.3%	Macintosh (8.8%)	Chrome	5.1%
	Firefox	3.5%		Safari	2.1%
	IE	2.8%		Firefox	1.3%
	Other	$< 0.1\%$		Other	0.3%
iOS (12.2%)	iPad	7.2%	iOS (5.6%)	iPhone	3.8%
	iPhone	4.6%		iPad	1.8%
	ios-app	0.4%			
Other (27.7%)	Android	26.9%	Other (9.2%)	Android	5.0%
	Linux	0.3%		Linux	3.8%
	ChromeOS	0.2%		Other	0.4%
	Other	0.3%			

counted as Android requests, since this was the case observed in our testing. The use of mobile devices for Netflix shows that when data is "free" for the user (since they do not have to pay cellular network fees when using campus WiFi), they do not mind using a smaller screen on a mobile phone or tablet. The differences between desktop and mobile traffic for Netflix are highlighted with response characterization in Section VI-B. Mobile requests to Twitch include the user-agent string, and use the same URI as desktop requests. Given the low volume of mobile traffic for Twitch, we do not differentiate between its mobile and desktop traffic.

The results in Table II show that Twitch is accessed primarily by users on Windows desktop platforms. This observation does not hold for Netflix, which is accessed from a wider variety of devices, including mobile. This pattern likely reflects the breadth and maturity of the commercial market for the Netflix service, compared to Twitch, which is targeted for the gaming community.

Table III provides a further breakdown of the browsers used, according to the user agent strings reported in requests. The results here reinforce the observations made above. While there are diverse browser platforms used for both services, Twitch is Windows-dominated, while Netflix is not.

V. CONTENT CHARACTERIZATION

This section analyzes video content popularity, in an attempt to identify additional similarities and differences between

Netflix and Twitch. These results help answer our second research question.

When characterizing content from Netflix and Twitch, we observe similar behaviors, with highly non-uniform access patterns. That is, a small subset of the content accounts for a large proportion of the traffic from these services. Over the five-month period, we observed 16,501 unique movieids on Netflix, and 6,677 unique Twitch streamers.

On Netflix, 50% of the data traffic volume arose from only 25 titles (2,801 movieids). Twitch is not as skewed, with 50% of the traffic volume accounted for by 42 streams (one of which was renamed during the observation period). Figure 2 shows the cumulative bytes for content accessed from both services. The dashed lines in the graphs represent the top 25 titles on Netflix and the top 25 streams on Twitch. Separate analysis (not shown here) indicates Zipf-like content popularity, with a long-tailed power-law distribution [10].

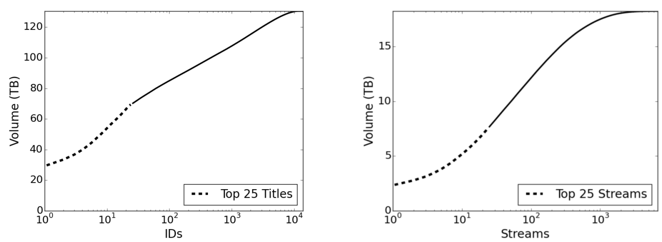

(a) Netflix Content Popularity (b) Twitch Stream Popularity

Fig. 2. Content Popularity: (a) Netflix; (b) Twitch

Table IV lists the top 20 titles from Netflix (by data volume) over the five-month period. The items in the table are ordered by their cumulative overall rank (on the left), with monthly ranks indicated in the columns on the right. Entries that have a dash instead of a number indicate no traffic for that title in that month. For Netflix, this is because the content had not yet been added to the catalog.

Table IV leads to several interesting observations about content popularity on Netflix. On initial inspection, we find that TV shows are much more popular than movies on Netflix. The reason is that a TV series offers a lot more content than a typical two-hour movie.

There are two general patterns in Netflix content: short-term popularity and long-term popularity. Examples of short-term popularity include *House of Cards*, *Suits*, and *Daredevil*. Examples of long-term popularity include *Friends*, *Grey's Anatomy*, and *Gossip Girl*.

New content on Netflix often exhibits short-term popularity. For example, when a season for a popular show is added to the catalog, viewers consume the new content quickly, resulting in a surge of popularity for a month or two. We can see this behavior in Table IV with *House of Cards*; the series surged when the third season was added in February, and the surge lasted at least two months before waning. Our data suggests that viewers on Netflix tend to "binge watch" newly added content (short-term), then return to watching favorite shows (long-term).

Examples of long-term popularity on Netflix include *Friends* and *Grey's Anatomy*. These series have a lot of content a viewer can watch, and are very popular month to month. *That 70's Show* is another example of content exhibiting long-term popularity. Table IV shows a jump from rank 49 to 4 between January and February – this surge-like behavior is easily explained based on Netflix catalogs. Viewers watching the series in December and January were doing so using special mechanisms to bypass Netflix's geo-restriction policies, so that they could view content from another country's catalog. Once this content was added to our country's catalog, many more users on campus were able to access it directly.

One reason for long-term popularity is the sheer volume of content available. For example, *Friends* has 10 seasons of content, making it difficult for a viewer to consume it fully in a short period of time. Users also find long-term content interesting; such titles on Netflix are often rated 7 or higher in the Internet Movies Database (`www.imdb.com`).

Table V provides a corresponding look at the top 20 streams in Twitch. When we look at the streams from Twitch, we see that their monthly ranks change more frequently than Netflix. This difference in stability is also observed in day-to-day rankings (not shown here); the top Netflix monthly titles tend to be among the top titles for any given day. This property does not hold for Twitch, which has a smaller user community, and is heavily driven by live events.

Selected streams from Twitch also show short-term and long-term popularity. On Twitch, the short-term streams are driven by events, such as eSports (professional tournaments). For example, streams like `esltv_lol`, `esl_lol`, and `esl_csgo` draw many viewers during eSports competitions.

Several streams on Twitch exhibit long-term popularity. For example, `riotgames`, `beyondthesummit`, and `imaqtpie` were popular throughout the observation period. The lower rankings for `riotgames` and `imaqtpie` in December likely reflect end-of-semester effects (e.g., studying, final exams, Christmas vacation). The popularity for `beyondthesummit` dipped in March when another channel featured a major tournament for that game.

VI. Streaming Protocol Usage

Netflix and Twitch both use DASH [16]. DASH works by breaking a larger file (or video stream) into a sequence of many smaller files that can be easily transmitted over the Internet. DASH servers can provide the video files in different quality levels. Clients interacting with the server dynamically choose the best quality possible, based on network conditions when requesting the next file in the sequence [15]. With live-streaming content, if a file cannot be transmitted in time, it is skipped and the next one is requested. DASH is the basis of Apple's HTTP live-streaming (HLS), which is what Twitch uses for their live-streaming solution.

A. Connection Characteristics

Both Netflix and Twitch use multiple connections per video to transport content. In our dataset, Netflix had a total of 14.3

TABLE IV
NETFLIX TOP VIDEO CONTENT RANK BY MONTH (DECEMBER 2014 TO APRIL 2015)

Rank	Netflix Title	Description	Dec	Jan	Feb	Mar	Apr
1	*Friends*	TV sitcom set in NYC, Warner/NBC, 1994-2004	-	1	1	1	1
2	*Grey's Anatomy*	TV medical drama set in Seattle, ABC, 2005-	1	2	2	3	2
3	*House of Cards*	Web TV political drama, Netflix, 2013-	20	16	3	2	9
4	*Gilmore Girls*	TV comedy-drama, Warner, 2000-2007	2	4	9	10	5
5	*Gossip Girl*	American teen TV drama, The CW, 2007-2012	3	3	7	7	7
6	*That 70's Show*	TV period sitcom in Wisconsin, Fox, 1998-2006	42	49	4	4	6
7	*Suits*	TV series, American legal drama, USA, 2011-	6	5	10	5	10
8	*The Mindy Project*	TV romantic comedy in NYC, Fox/Hulu, 2012-	8	7	16	9	4
9	*Supernatural*	TV series, fantasy horror, WB/CW, 2005-	5	10	6	12	11
10	*House M.D.*	TV medical drama set in NJ, Fox, 2004-2012	7	9	5	13	14
11	*How I Met Your Mother*	TV sitcom, flashback theme, CBS, 2005-2014	4	12	13	11	13
12	*The 100*	TV science fiction drama, CW, 2014-	12	14	8	8	28
13	*White Collar*	TV criminal drama, USA, 2009-2014	13	6	12	16	18
14	*90210*	TV teen drama, California, The CW, 2008-2013	17	41	15	18	8
15	*The Vampire Diaries*	TV supernatural drama, The CW, 2009-	16	11	11	14	39
16	*The Office*	TV sitcom, Pennsylvania, NBC, 2005-2013	11	20	19	15	15
17	*Archer*	TV adult animated spy comedy, FX, 2010-	18	8	14	17	36
18	*Daredevil*	Web TV superhero action drama, Netflix, 2015-	-	-	-	-	3
19	*Family Guy*	TV adult animated sitcom, Fox, 1999-	9	22	21	19	17
20	*Dexter*	TV crime drama mystery, Showtime, 2006-2013	14	19	24	23	12

million connections transporting content, while Twitch had 1.6 million connections transporting video and non-video content.

Connections for both Netflix and Twitch are persistent and use pipelined requests. Netflix uses concurrent connections to servers to retrieve content. Requests made from the different connections in the browser may be interleaved (though this behavior was not seen on clients using the new Web Interface in June 2015). We observed concurrent connections and interleaved requests on Twitch only at the start of a new connection; after the initial requests, Twitch uses a single connection. Parallel connections may be used by these services to improve throughput to the client.

TABLE VI
NETFLIX AND TWITCH CONNECTION CHARACTERISTICS

	Netflix		Twitch	
Month	Avg. Size	Avg. Dur	Avg. Size	Avg. Dur
December	27.72 MB	166 sec.	22.67 MB	118 sec.
January	28.44 MB	169 sec.	22.22 MB	114 sec.
February	26.47 MB	169 sec.	20.46 MB	100 sec.
March	24.66 MB	166 sec.	21.63 MB	116 sec.
April	24.30 MB	165 sec.	23.38 MB	132 sec.

Table VI shows monthly TCP connection characteristics for Netflix and Twitch. The average bit rates are approximately 1.2-1.5 Mbps for these connections. The Netflix connections are slightly larger and last longer than Twitch. Responses, detailed in the next section, show similar characteristics.

B. HTTP Response Characteristics

In this section, we characterize HTTP responses from Netflix and Twitch. As stated in Section IV, we characterize mobile and desktop responses from Netflix separately. The Twitch responses that we characterize are live-stream responses. For both Netflix and Twitch, the requests from the client browsers use the GET method, with a body length of zero (i.e., no data).

Figure 3 shows response characteristics for video content from Netflix (desktop and mobile) and Twitch. This figure shows that Netflix and Twitch behave differently at the application level. The distribution of response sizes, measured by body-length, are shown in Figure 3(a). Netflix shows some minor differences between desktop responses (blue line) and mobile responses (dashed blue line). Mobile Netflix responses are slightly smaller than desktop responses; this is likely due to smaller screen sizes on mobile devices. The step behavior that Netflix responses show is due to the different response sizes Netflix uses to transport content. Median Netflix response sizes are 790 KB and 303 KB, for desktop and mobile, respectively. At the 99th percentile, response sizes are 3.44 MB and 3.04 MB, for desktop and mobile. A majority of responses from Twitch are smaller than Netflix responses. Twitch responses show a step behavior that is caused by the two types of responses (`Video/mp2t` and `Application/vnd.apple.mpegurl`) issued by Apple HLS and by users requesting different qualities for their stream. Twitch's median response length is 460-470 bytes; this is due to the second response type used by Apple HLS. At the 99th percentile, Twitch has a response length of 2.05 MB.

Response durations are shown in Figure 3(b). Response durations are measured by taking the timestamps on the first and last segments of the response. The median response time for both Netflix and Twitch was 0 seconds; that is, the entire

TABLE V
TWITCH TOP VIDEO CONTENT RANK BY MONTH (DECEMBER 2014 TO APRIL 2015)

Rank	Twitch Stream	Description	Dec	Jan	Feb	Mar	Apr
1	riotgames	Video gaming and eSports company, Los Angeles	338	1	1	1	1
2	beyondthesummit	Broadcasting org and tournament sponsor for DoTA 2	2	2	2	14	5
3	imaqtpie	Player for international eSports Team Dignitas	13	5	3	4	4
4	lirik	Partnered live streamer on Twitch since 2011	7	3	13	13	8
5	nl_kripp	Canadian Twitch streamer with Team SoloMid	5	8	5	22	2
6	esltv_lol	Electronic sports league TV for League of Legends	1	27	-	-	-
7	trumpsc	Professional video game player, US, Team SoloMid	6	7	8	10	9
8	summit1g	Popular CSGO video game streamer based in US	8	44	28	6	3
9	tsm_theoddone	Retired LoL player streaming for Team SoloMid	4	11	12	7	22
10	destiny	Online-only first-person shooter video game	3	9	21	20	17
11	esl_lol	Multi-player online battle arena game (LoL, ESL)	-	-	-	2	1618
12	faceittv	Automated platform for access to eSports video games	53	6	9	25	19
13	dotastarladder	Defense of the Ancients tournament (Warcraft III mod)	35	24	-	5	6
14	amazhs	Professional video game player for Team NRG	9	10	17	38	20
15	clgdoublelift	American professional LoL player for Team SoloMid	20	31	18	21	10
16	forsenlol	Swedish professional video game player	17	18	16	16	31
17	mushisgosu	Canadian LoL player and streamer for Team SoloMid	42	15	15	9	104
18	flosd	Competitive LoL player based in USA	22	12	6	33	65
19	esl_csgo	eSports league for Counter Strike Global Offensive	-	-	-	3	61
20	riotgames2	Promotional channel with free videos for Riot Games	-	37	19	15	16

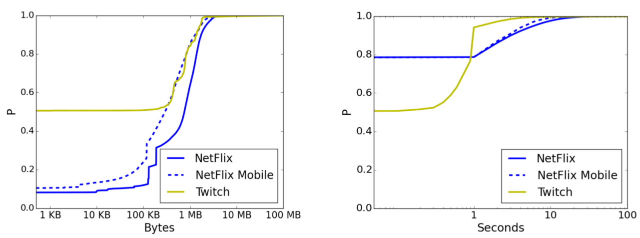

(a) Response Size CDF (b) Response Duration CDF

Fig. 3. Response Characteristics: (a) Size CDF; (b) Duration CDF

response was delivered with a single packet. Netflix responses again show slightly different behavior between mobile and desktop devices. Responses sent to mobile devices are slightly faster than desktop. This is likely a result of the smaller responses that mobile devices receive. As the figure shows, Twitch's responses often take less than one second. Quick responses are extremely important for Twitch since it is a live-stream service. When responses take too long, the viewer does not receive a live-stream and their experience is greatly diminished. The 99th percentile response durations were between 4-5 seconds for Twitch, 17-18 seconds for Netflix, and 10-11 seconds for Netflix mobile.

VII. DISCUSSION OF PERFORMANCE IMPLICATIONS

This section addresses our third and final research question, by discussing the implications of our observations.

As we have shown, Netflix and Twitch can consume significant amounts of network bandwidth. Since both of these services are expected to grow [8], we next discuss how one might alter delivery infrastructure to eliminate redundant traffic and simultaneously improve user experience.

The service access patterns show that the traffic is human-generated, and contributes significantly to the daily peak traffic volume. Lessening the impact of this traffic would improve network performance.

The content being accessed shows highly skewed popularity, with a majority of bytes generated from a relatively small number of titles or streams. Caching popular content from Netflix locally (e.g., by hosting a CDN node on campus) would greatly reduce the amount of traffic that is transmitted over the network at a low cost to the network operator. For example, the TV show *Friends* generated over 20 TB of traffic in five months; we estimate that caching this series in its entirety would require only 70 GB of hard disk space.

For campus networks, Twitch provides an obvious use case for native IP multicast, using either IPv4 or IPv6. However, IP multicast support is rarely enabled in enterprise networks. The next best solution is rebroadcasting streams locally (e.g., application-level multicast via a local CDN node). This could have similar network savings, and could improve service by providing faster response times. Streams broadcasting eSports events increase traffic on Twitch when they occur. Rebroadcasting these live streams, whose scheduled times are known months in advance (providing ample time to set up a local CDN node), would reduce network utilization.

Another important observation is that in an environment where there are no data limits on mobile devices, we saw non-negligible use of Netflix and Twitch on mobile devices. This shows that when connection costs are not a concern, users are interested in viewing video content on a variety of devices. This behavior may differ from what is observed on cellular networks today, where there is a high cost to use data-intensive services on mobile devices. If data becomes cheaper on cellular networks, we anticipate video use will increase well above current levels. The increasing use of mobile devices is something that both service providers and network operators need to consider in their plans for the future.

While both Netflix and Twitch are based on DASH, we saw some differences in the connection and response-level characteristics between the two services. Netflix provides video-on-demand content, which exhibits different behaviors than Twitch's specialization in live-stream content. Multiple connections are used by both of these services when delivering content. Netflix used concurrent connections throughout delivery, while Twitch only used concurrent connections when a new stream was requested. Responses also showed different behaviors. Netflix responses were larger, and differed between mobile and desktop devices. Twitch responses were smaller and faster than Netflix, allowing the user to view a live stream. Since the use of mobile devices is non-negligible, protocol designers and service providers need to consider diverse platforms for streaming video.

VIII. Conclusions and Future Work

In this study, we characterized Netflix and Twitch traffic using data collected from December 2014 through April 2015. These were the two largest unencrypted video services accessed from the university network during that time period.

Despite the differences in the services and their target audiences, we see many similarities in traffic characteristics. Traffic for both services is significant and expected to increase in the future. Access to the services are driven by humans, and put additional strain on network links during peak usage. Mobile devices are commonly used to access both services. Content accessed from both services shows skewed access patterns amenable to caching. Popular content on both services also shows similar behaviors, making it easier to predict viewing trends. Finally, connections and responses on both of these services showed many of the same characteristics.

Even though Netflix and Twitch are rather different services, both could benefit from improvements to their (similar) transport and delivery mechanisms. Such solutions could include campus-level CDNs to shorten network RTTs, improve TCP performance, and enhance the video streaming experience for users. With such a solution, future edge networks would be better able to accommodate the ongoing growth in video streaming traffic (e.g., users, bandwidth, content availability), especially on mobile/wireless networks.

Our study provides a final look at unencrypted Netflix traffic. Future studies will be challenged by the ongoing trend toward application-level encryption.

Acknowledgment

The authors thank the anonymous reviewers for their constructive feedback on our work. Financial support for this research was provided by Canada's Natural Sciences and Engineering Research Council (NSERC), and the Queen Elizabeth II scholarship program at the University of Calgary.

References

[1] V. Adhikari, Y. Guo, F. Hao, V. Hilt, Z-L. Zhang, M. Varvello, and M. Steiner, "Unreeling Netflix: Understanding and Improving Multi-CDN Movie Delivery", *Proceedings of IEEE INFOCOM*, Orlando, FL, pp. 1620-1628, March 2012.

[2] V. Adhikari, Y. Guo, F. Hao, V. Hilt, Z-L. Zhang, M. Varvello, and M. Steiner, "Measurement Study of Netflix, Hulu, and a Tale of Three CDNs", *IEEE/ACM Transactions on Networking*, Vol. 23, No. 6, pp. 1984-1997, December 2015.

[3] Y. Borghol, S. Mitra, S. Ardon, N. Carlsson, D. Eager, and A. Mahanti, "Characterizing and Modelling Popularity of User-generated Videos", *Proceedings of IFIP Performance*, Amsterdam, Netherlands, pp. 1037-1055, October 2011.

[4] M. Cha, H. Kwak, P. Rodriguez, Y. Ahn, and S. Moon, "Analyzing the Video Popularity Characteristics of Large-Scale User Generated Content Systems", *IEEE/ACM Transactions on Networking*, Vol. 17, No. 5, pp. 1357-1370, October 2009.

[5] Cisco Systems Inc.: Cisco Visual Networking Index: Forecasting and Methodology, 2014-2019. Tech. rep., Cisco Systems Inc., May 2015.

[6] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "YouTube Traffic Characterization: A View from the Edge", *Proceedings of ACM Internet Measurement Conference*, San Diego, CA, pp. 15-28, October 2007.

[7] W. Hamilton, O. Garretson, and A. Kerne, "Streaming on Twitch: Fostering Participatory Communities of Play within Live Mixed Media", *Proceedings of ACM CHI*, Toronto, ON, pp. 1315-1324, April 2014.

[8] J. Hirschhorn, "7 Deadly Sins: Where Hollywood is Wrong about the Future of TV". `https://www.linkedin.com/pulse/7-deadly-sins-where-hollywood-is-wrong-about-the-future-of-tv-jason-hirschhorn`, 2015.

[9] M. Kaytoue, A. Silva, L. Cerf, W. Meira, Jr., and C. Raïssi, "Watch Me Playing, I Am a Professional: A First Study on Video Game Live Streaming", *Workshop on Mining Social Network Dynamics (MSND), Companion Proceedings of WWW 2012*, Lyon, France, pp. 1181–1188, April 2012.

[10] M. Laterman, *Netflix and Twitch Traffic Characterization*, MSc Thesis, Department of Computer Science, University of Calgary, September 2015.

[11] J. Martin, Y. Fu, N. Wourms, and T. Shaw, "Characterizing Netflix Bandwidth Consumption", *Proceedings of the IEEE Consumer Communications and Networking Conference (CCNC)*, Las Vegas, NV, pp. 230–235, January 2013.

[12] G. Nascimento, M. Ribeiro, L. Cerf, N. Cesario, M. Kaytoue, C. Raïssi, T. Vasconcelos, and W. Meira, "Modeling and Analyzing the Video Game Live-Streaming Community", *Proceedings of the Latin America Web Congress (LA-Web)*, Ouro Perto, Brazil, pp. 1-9, October 2014.

[13] Netflix, `http://www.netflix.com`, 2015.

[14] V. Paxson, "Bro: A System for Detecting Network Intruders in Real-Time" *Computer Networks*, Vol. 31, No. 23-24, pp. 2435–2463, December 1999.

[15] A. Reed and J. Aikat, "Modeling, Identifying, and Simulating Dynamic Adaptive Streaming over HTTP", Best Poster award, *Proceedings of IEEE International Conference on Network Protocols (ICNP)*, Goettingen, Germany, pp. 1-2, October 2013.

[16] T. Stockhammer, "Dynamic Adaptive Streaming over HTTP: Standards and Design Principles" *Proceedings of the Second Annual ACM Conference on Multimedia Systems (MMSys)*, San Jose, CA, pp. 133–144, February 2011.

[17] J. Summers, T. Brecht, D. Eager, and A. Gutarin, "Characterizing the Workload of a Netflix Streaming Video Server", *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC)*, Providence, RI, pp. 43-54, October 2016.

[18] Twitch, `http://www.twitch.tv`, 2015.

[19] K. Xu, F. Wang, L. Gu, J. Gao, and Y. Jin, "Characterizing Home Network Traffic: An Inside View" *Personal Ubiquitous Computing*, Vol. 18, No. 4, pp. 967–975, April 2014.

[20] C. Zhang and J. Liu, "On Crowdsourced Interactive Live Streaming: A Twitch.tv-based Measurement Study", *Proceedings of ACM NOSSDAV*, Portland, OR, pp. 55–60, March 2015.