

# NETFLIX TRAFFIC CHARACTERIZATION

---

Michel Laterman  
Department of Computer Science  
University of Calgary

Supervisors: Carey Williamson and Martin Arlitt

# Introduction

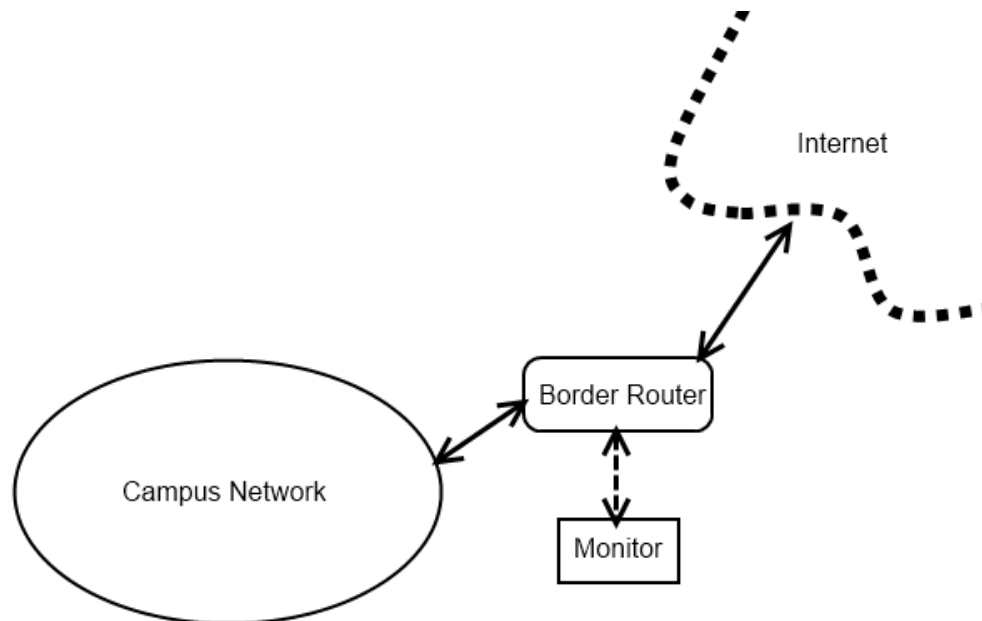
- Video streaming traffic constitutes a large (and growing!) proportion of modern Internet traffic
- Popular video streaming services include:
  - YouTube – user-generated content, short-clips (well-studied)
  - NetFlix – on-demand video, TV shows, movies (some studies)
  - Twitch – live streaming of video game play (few studies)
  - Vimeo – video-sharing site with High-Definition videos
  - Hulu – on-demand video, not in Canada
  - Yahoo Screen – professionally produced content, limited availability in Canada
- On the University of Calgary network, the top video streaming sites observed are YouTube, NetFlix, Twitch

# Research Objectives

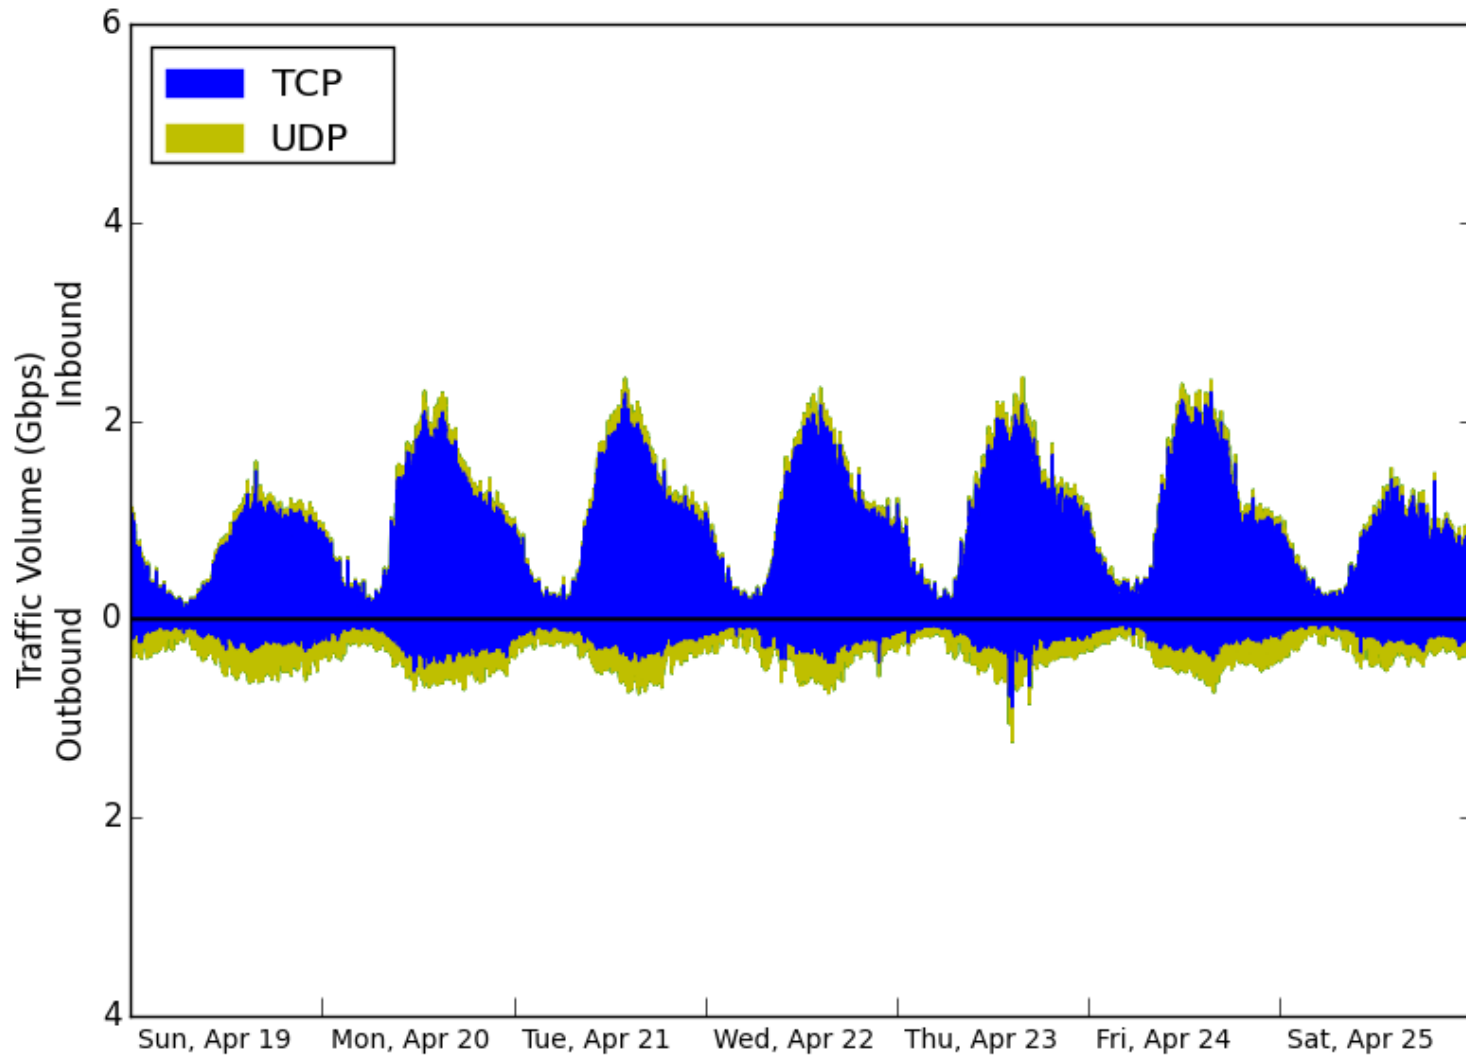
- General
  - Improve understanding of U of C network traffic
  - Identify network performance problems and anomalies
- Specific
  - Characterize video streaming services on U of C network
  - Understand similarities/differences between NetFlix and Twitch

# Methodology

- Passive network traffic measurement
- Hardware: Endace DAG packet capture card
- Software: Bro network security monitor
- 5 months of data (December 1, 2014 to April 29, 2015)
- Analysis of TCP connection and HTTP transaction logs



# Example: Traffic Overview (April 2015)



# HTTP Traffic Overview

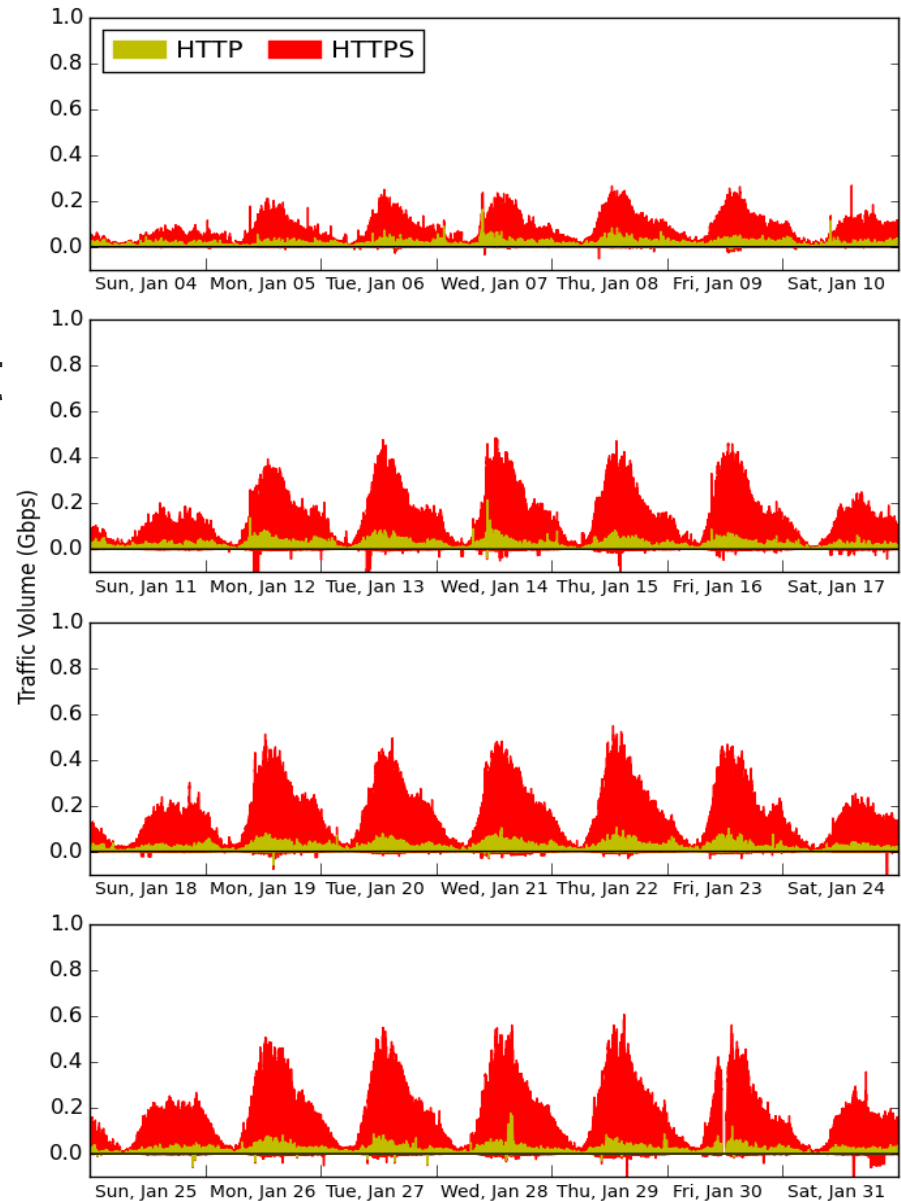
Host	Req. Percent	Volume
netflix.com	33.81%	217.1 TB
apple.com	8.37%	53.75 TB
googlevideo.com	2.43%	15.59 TB
steampowered.com	2.14%	13.79 TB
twitch.tv	2.04%	13.12 TB

# HTTPS Traffic Overview

Host	Connections	Percent	Volume
google.com	314 million	7.91%	27.3 TB
apple.com	179 million	4.51%	2.8 TB
majuwe.com	168 million	4.23%	106.7 GB
akamaihd.com	151 million	3.80%	32.7 TB
googlevideo.com	131 million	3.30%	230.1 TB

# YouTube Traffic

- January 2015
- Uses HTTPS by default
- HTTP for some embedded clips
- Outbound traffic is for video uploads

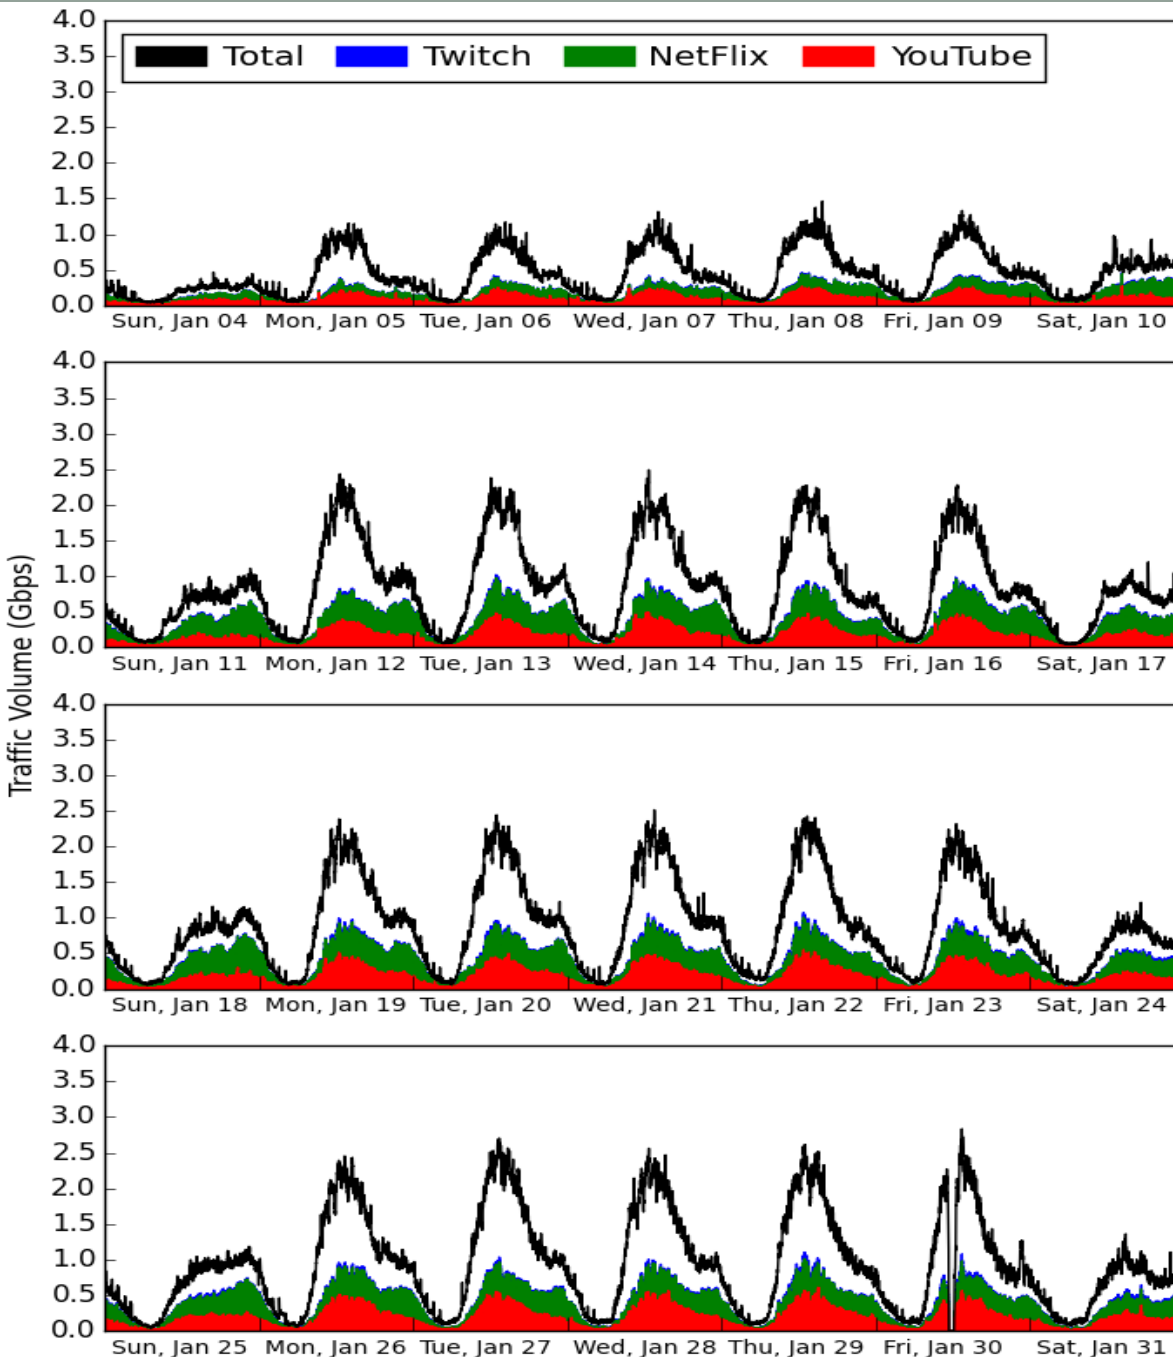




# Video Traffic Volume

- Outbound traffic to NetFlix and Twitch is negligible.

	YouTube - HTTP		YouTube - HTTPS		NetFlix	Twitch
	Inbound	Outbound	Inbound	Outbound	Inbound	Inbound
December	1.93 TB	0.14 TB	36.22 TB	0.89 TB	30.77 TB	2.82 TB
January	1.89 TB	0.12 TB	36.31 TB	1.06 TB	44.41 TB	3.14 TB
February	1.79 TB	0.05 TB	45.47 TB	1.14 TB	43.83 TB	3.74 TB
March	2.08 TB	0.05 TB	59.63 TB	1.36 TB	54.29 TB	4.79 TB
April	1.51 TB	0.05 TB	52.43 TB	1.08 TB	43.85 TB	3.74 TB



# Video Traffic

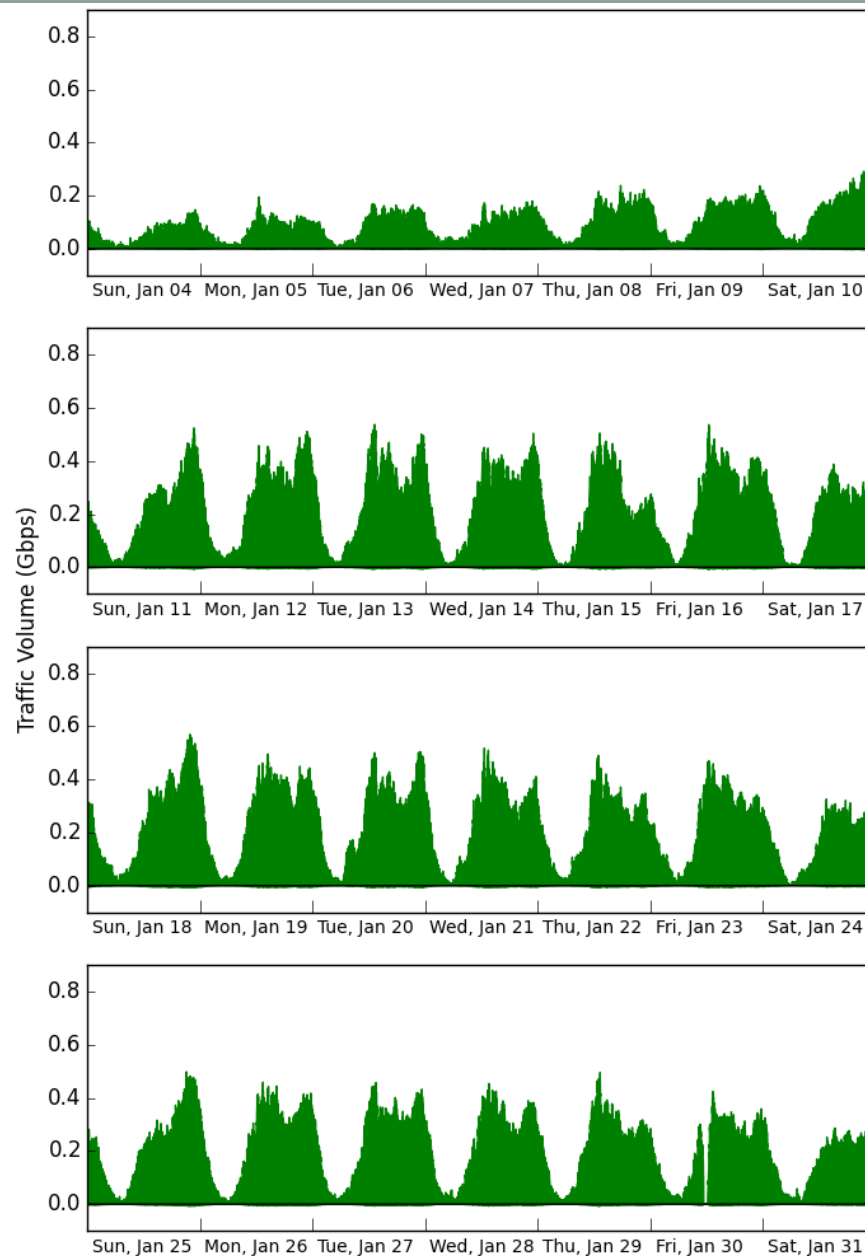
- January 2015
- Top line (Total) is HTTP+HTTPS
- Red is (HTTPS) YouTube
- Green is NetFlix
- Blue is Twitch

# NetFlix

- 305 million request-response pairs on 14.3 million connections generating 217.1 TB of volume
- 62.9% of requests had code 200 (OK), 29.9% had 206 (Partial content), 6.09% had no code.
- 35 different content-type headers
  - Application/octet-stream 216.7 TB
  - Text/html 328.8 GB

# NetFlix Traffic

- Video content is served from several unnamed servers with NetFlix IP addresses
- 217.1 TB total traffic
- Connections average 26 MB in, 370 KB out
- Average duration 150 seconds



# NetFlix – Video Delivery

- HTML5 Player (transitioned away from Silverlight)
- Requests to the Web interface player include a parameter called movieID
- Desktop and Mobile devices use different request paths
  - Can't see movieid from mobile requests
- 162.6 TB of traffic was responses to content requests from desktop devices, 54.01 TB mobile
- Multiple connections are used to transport video (7-9 for a 22 min episode, 14-16 for 42 min)


# NetFlix – What are people Watching?

Title	Dec	Jan	Feb	Mar	Apr
1. Friends	-	1	1	1	1
2. Grey's Anatomy	1	2	2	3	2
3. House of Cards	20	16	3	2	9
4. Gilmore Girls	2	4	9	10	5
5. Gossip Girl	3	3	7	7	7
6. That 70's Show	42	49	4	4	6
...					
18. Daredevil	-	-	-	-	3

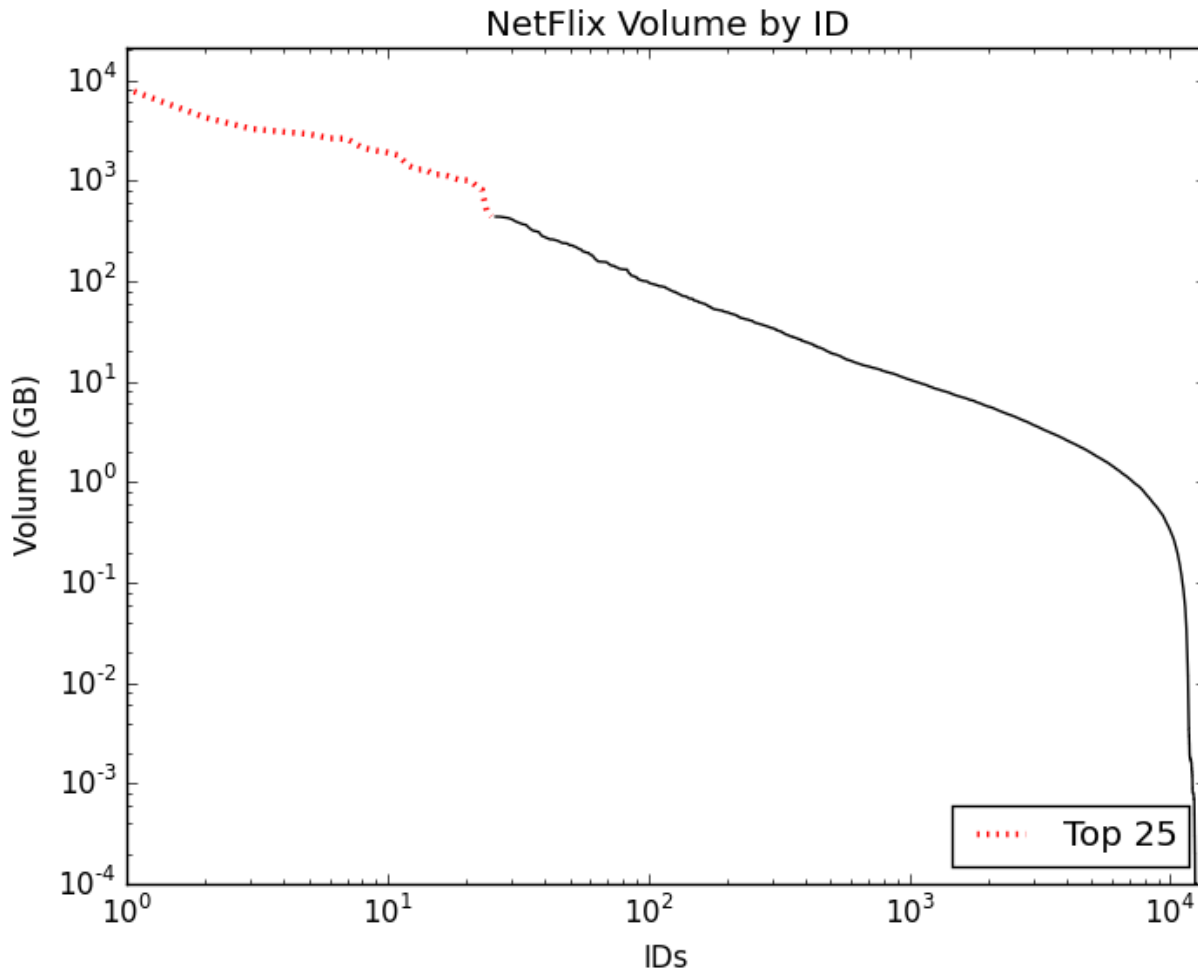
Long-term popularity

Short-term popularity

# A Week of NetFlix Traffic – Top Content



# NetFlix movieID Traffic Volumes

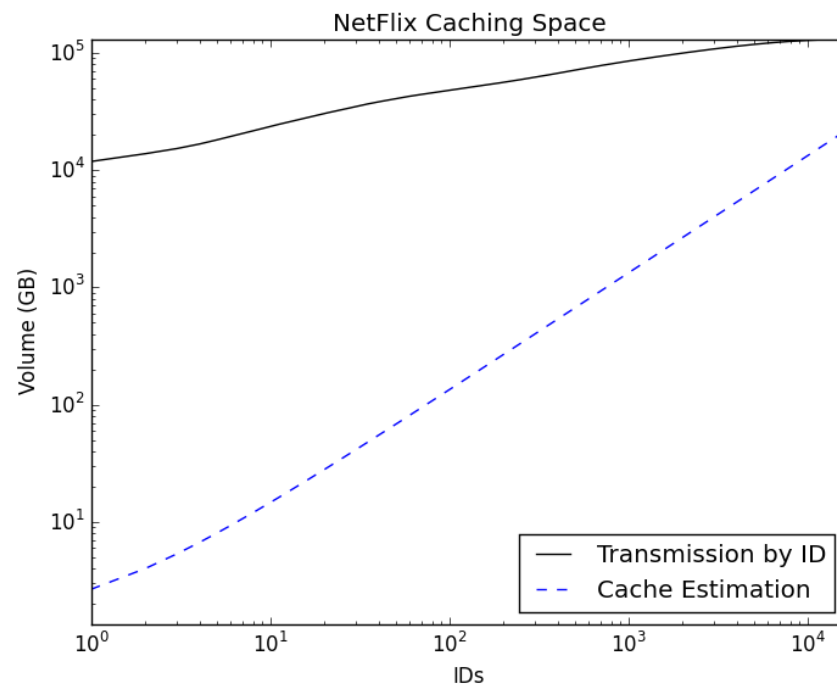
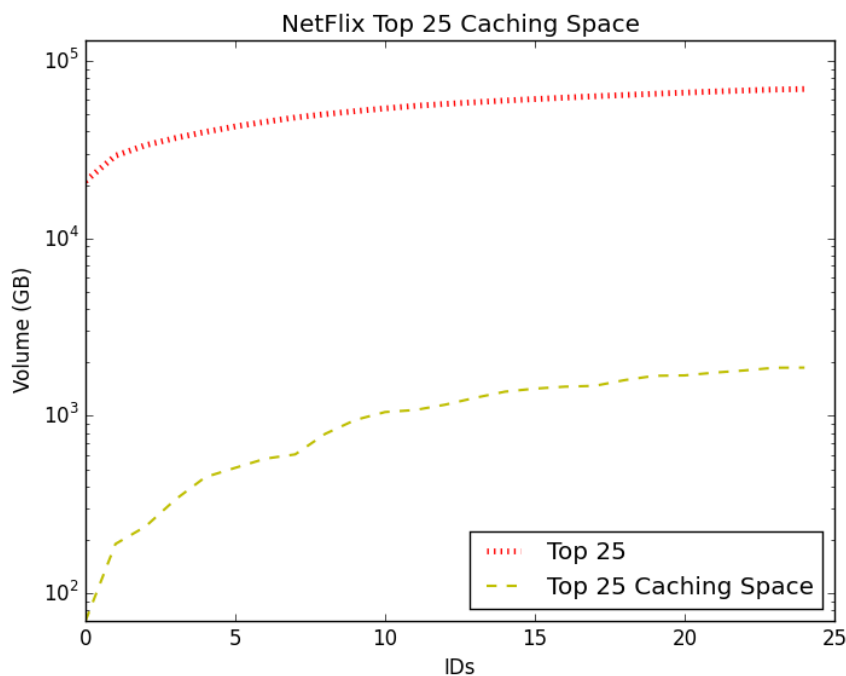


- Top 25 shows (2,801 IDs)
  - 50% of traffic volume
- Friends: 21 TB
- Grey's Anatomy: 8 TB
- House of Cards: 4 TB



# Caching NetFlix

- File sizes: 13.23 MB/minute (SD) or 22.58 MB/min (HD)
- 70 GB to cache Friends (21 TB transmission)
- 120 GB to cache Grey's Anatomy (8.2 TB)
- 40 GB to cache House of Cards (4.25 TB)



# Conclusions (Netflix)

- Video streaming services constitute a large proportion of inbound traffic on the U of C network
- YouTube and NetFlix are the most popular currently
- Caching NetFlix could greatly reduce network traffic
  - Caching “Friends” (70 GB) would reduce traffic by 20 TB
- Studies like this will be much more difficult once Netflix moves to HTTPS for all content delivery (mid-2015)