Things That Go Bump in the Net: The Many Challenges of Debugging Network Application Performance

UNIVERSITY OF CALGARY

Carey Williamson

Department of Computer Science

University of Calgary

Introduction

- Network traffic measurement requires hardware or software measurement tools that attach directly to network
- Allows you to observe all packet traffic on the network (or a filtered subset for traffic of interest)
- Assumes broadcast-based network technology, superuser permission

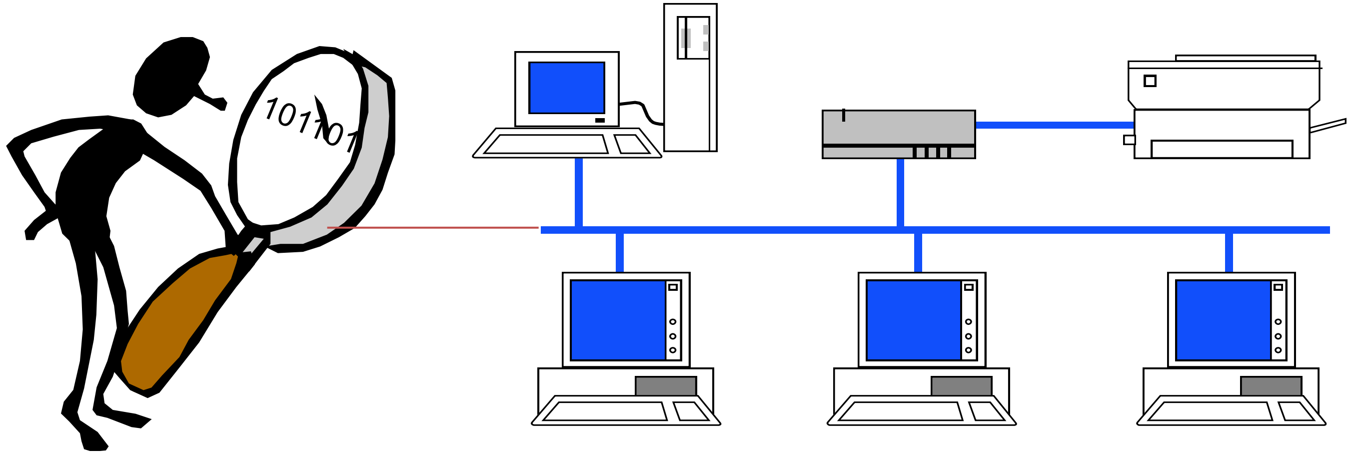

Example: tcpdump or Wireshark

```
Time       IP Source Addr     IP Dest Addr   Size  Prot  SPort DPort   TCP Data SeqNumber       TCP AckNum   Window   Flags

0.000000 192.168.1.201 -> 192.168.1.200    60 TCP   4105   80 1315338075 : 1315338075 0 win: 5840 S
0.003362 192.168.1.200 -> 192.168.1.201    60 TCP     80 4105 1417888236 : 1417888236 1315338076 win: 5792 SA
0.009183 192.168.1.201 -> 192.168.1.200    52 TCP   4105   80 1315338076 : 1315338076 1417888237 win: 5840 A

0.010854 192.168.1.201 -> 192.168.1.200   127 TCP  4105   80 1315338076 : 1315338151 1417888237 win: 5840 PA
0.014309 192.168.1.200 -> 192.168.1.201    52 TCP     80 4105 1417888237 : 1417888237 1315338151 win: 5792 A

0.049848 192.168.1.200 -> 192.168.1.201 1500 TCP     80 4105 1417888237 : 1417889685 1315338151 win: 5792 A
0.056902 192.168.1.200 -> 192.168.1.201 1500 TCP    80  4105 1417889685 : 1417891133 1315338151 win: 5792 A
0.057284 192.168.1.201 -> 192.168.1.200    52 TCP   4105   80 1315338151 : 1315338151 1417889685 win: 8688 A
0.060120 192.168.1.201 -> 192.168.1.200    52 TCP   4105   80 1315338151 : 1315338151 1417891133 win: 11584 A
0.068579 192.168.1.200 -> 192.168.1.201 1500 TCP     80 4105 1417891133 : 1417892581 1315338151 win: 5792 PA
0.075673 192.168.1.200 -> 192.168.1.201 1500 TCP     80 4105 1417892581 : 1417894029 1315338151 win: 5792 A
0.076055 192.168.1.201 -> 192.168.1.200    52 TCP   4105   80 1315338151 : 1315338151 1417892581 win: 14480 A
0.083233 192.168.1.200 -> 192.168.1.201 1500 TCP    80 4105 1417894029 : 1417895477 1315338151 win: 5792 A
0.096728 192.168.1.200 -> 192.168.1.201 1500 TCP    80 4105 1417896925 : 1417898373 1315338151 win: 5792 A
0.103439 192.168.1.200 -> 192.168.1.201 1500 TCP    80 4105 1417898373 : 1417899821 1315338151 win: 5792 A
0.103780 192.168.1.201 -> 192.168.1.200    52 TCP   4105   80 1315338151 : 1315338151 1417894029 win: 17376 A
0.106534 192.168.1.201 -> 192.168.1.200    52 TCP   4105   80 1315338151 : 1315338151 1417898373 win: 21720 A

0.133408 192.168.1.200 -> 192.168.1.201  776 TCP    80 4105 1417904165 : 1417904889 1315338151 win: 5792 FPA
0.139200 192.168.1.201 -> 192.168.1.200    52 TCP   4105   80 1315338151 : 1315338151 1417904165 win: 21720 A
0.140447 192.168.1.201 -> 192.168.1.200    52 TCP   4105   80 1315338151 : 1315338151 1417904890 win: 21720 FA
0.144254 192.168.1.200 -> 192.168.1.201    52 TCP    80 4105 1417904890 : 1417904890 1315338152 win: 5792 A
```

Flow summary (e.g., NetFlow record or Bro connection log entry):

0.000000 192.168.1.201 4105 192.168.1.200 80 0.144254 10 77 11 16654 SF

Example: Bro Connection Log

Time	IP Source Addr	Port	IP Dest Addr	Port	Duration	PS	PR	BS	BR	State
0.000000	192.168.1.201	4105	192.168.1.200	80	0.144254	10	77	11	16654	SF
0.237814	192.168.1.285	7336	192.168.1.200	80	2.765018	32	105	937	87932	SF
0.589206	192.168.1.141	1060	192.168.1.200	80	0.842541	15	26	361	37850	SF
0.837142	192.168.1.251	8109	192.168.1.200	80	0.713306	12	54	110	32768	RST
1.249788	192.168.1.281	7206	192.168.1.200	80	1.517842	81	340	1096	181654	SF
1.742355	192.168.1.271	4812	192.168.1.200	80	0.642311	15	71	82	3784	SF
2.168283	192.168.1.146	1090	192.168.1.200	80	5.254088	10	385	36	20176	SF
2.577825	192.168.1.285	7339	192.168.1.200	80	0.034217	7	46	18	9184	SF
3.492006	192.168.1.236	3607	192.168.1.200	80	0.594426	18	61	105	5408	SF
4.587426	192.168.1.141	1061	192.168.1.200	80	0.331344	11	20	28	12716	SF
5.824413	192.168.1.231	6022	192.168.1.200	80	0.680049	24	75	31	18533	SF
6.073508	192.168.1.104	8704	192.168.1.200	80	0.913426	27	37	88	14236	SF
7.198741	192.168.1.251	8122	192.168.1.200	80	1.744125	52	128	238	75890	SF
7.363601	192.168.1.281	7218	192.168.1.200	80	0.164425	12	8	22	6654	RST
8.597769	192.168.1.141	1063	192.168.1.200	80	0.517756	18	119	310	15024	SF
8.370944	192.168.1.271	4818	192.168.1.200	80	0.027399	6	30	45	18324	SF
9.127458	192.168.1.235	4093	192.168.1.200	80	2.044254	35	264	212	172654	SF
9.627145	192.168.1.281	7225	192.168.1.200	80	0.283158	15	46	53	18498	SF

U of C Monitoring & Analysis Infrastructure

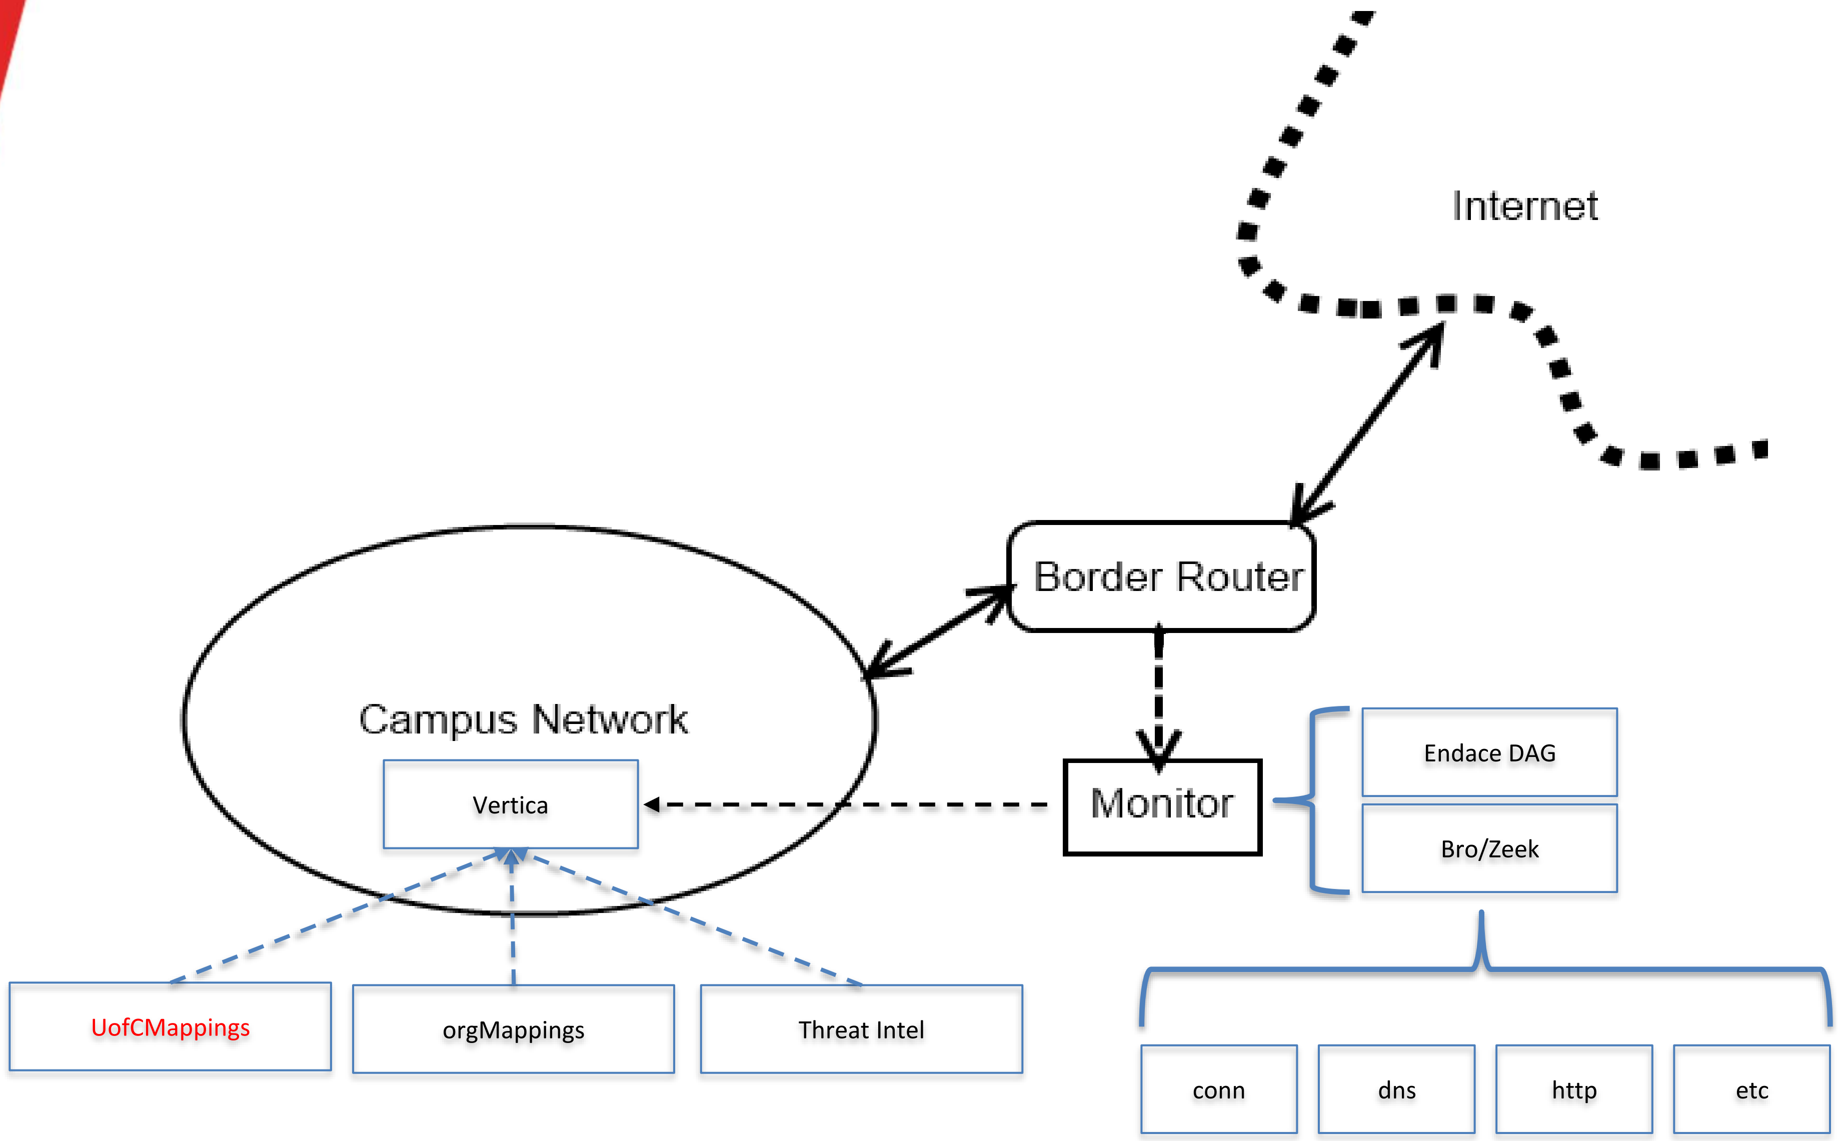

Network Traffic Data Volume (2014-present)

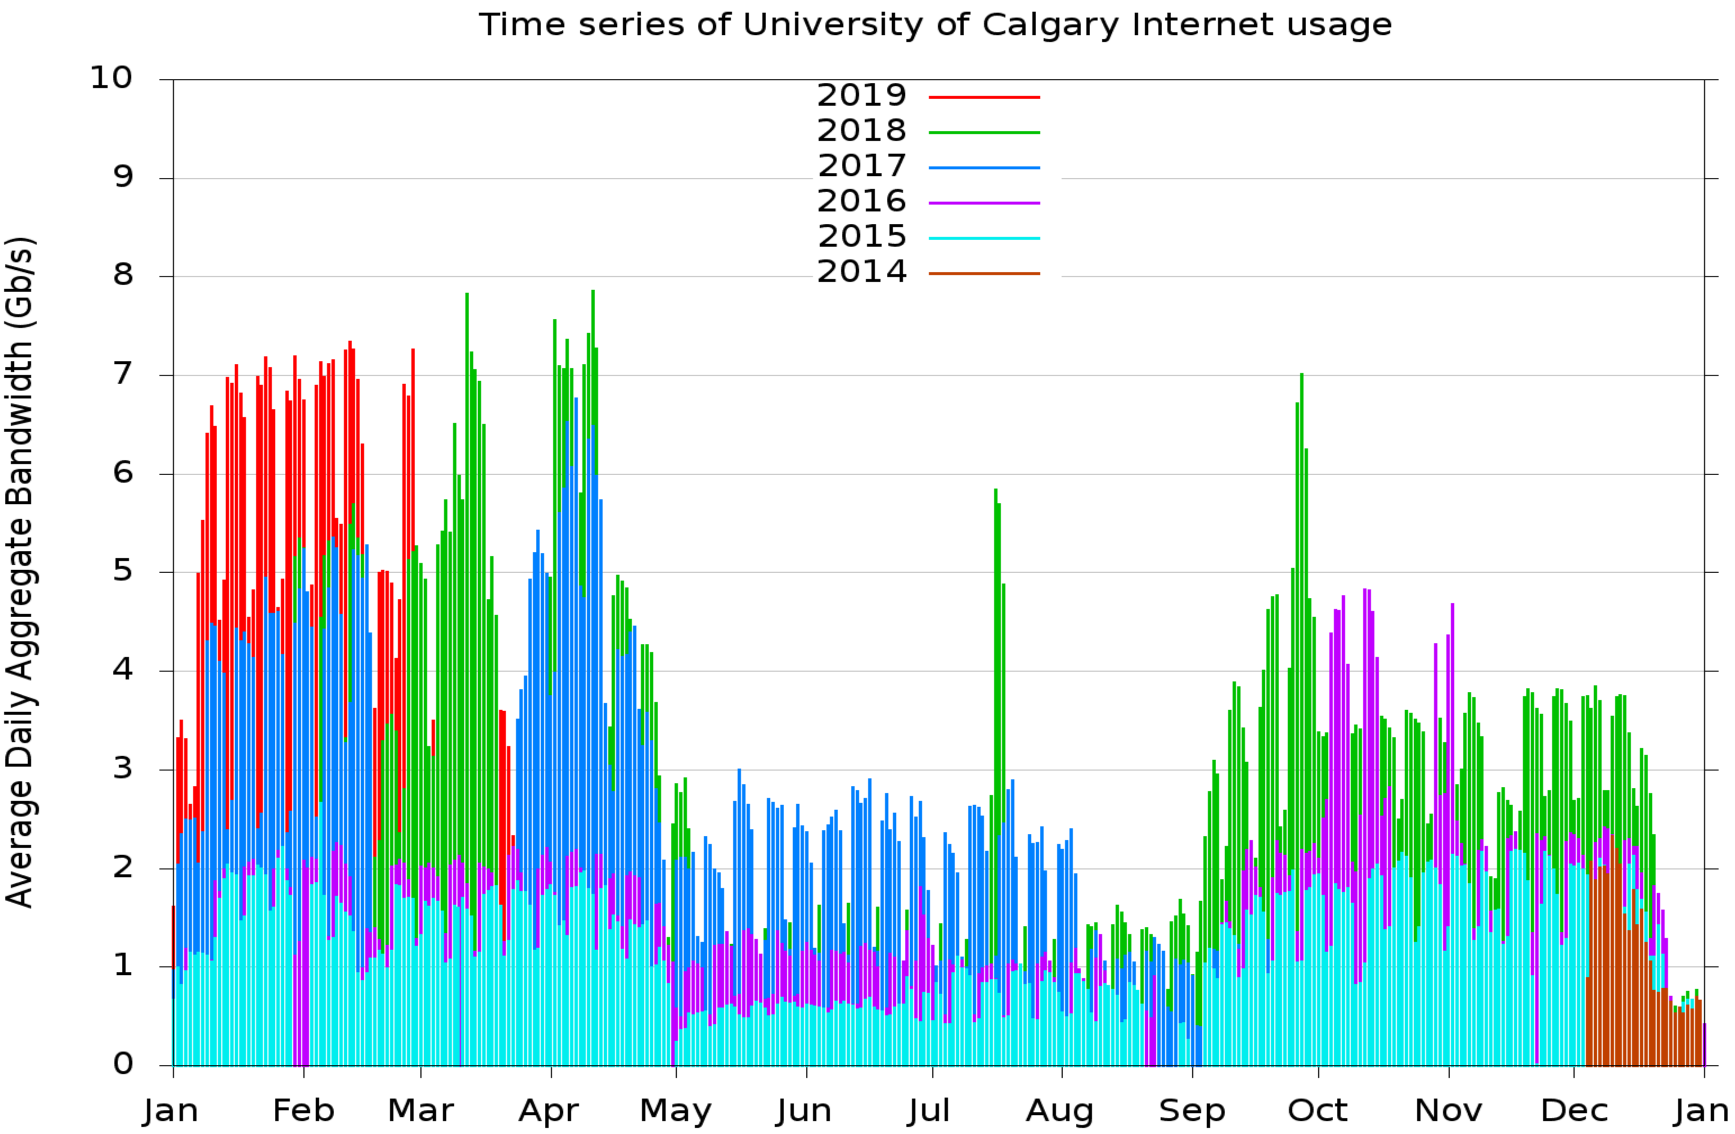

Top 20 Sites and Services (2019)

DstOrg	Srcs	Dsts	Services	Conns	OGB	RGB
Netflix Streaming Services Inc.	906	310	4	266308	59.3	6,089.6
Canarie Inc	1941	22	8	449694	13.4	1,551.8
Google LLC	3336	8413	250	7433788	170.2	1,547.0
Facebook, Inc.	29177	602	604	2217479	61.1	1,072.4
Apple Inc.	3129	2705	52	3087969	181.4	972.9
Amazon.com, Inc.	27610	54749	1578	6904241	128.2	717.3
Twitch Interactive Inc.	224	96	3	27215	19.1	693.2
Fastly	3055	728	9	822142	12.3	494.4
Akamai Technologies, Inc.	6295	9468	102	2603508	28.8	449.5
Microsoft Corporation	3498	3916	264	4354021	136.9	165.9
Shaw Communications Inc.	758	2410	2828	221592	55.1	147.0
Tencent Building, Kejizhongyi Avenue	1077	1850	548	918944	21.7	127.4
Dropbox, Inc.	1103	68	9	485147	90.6	119.0
No.31,Jin-rong Street	58861	48437	33480	811287	14.3	82.9
TELUS Communications Inc.	1099	2144	2048	311961	36.7	76.4
Bell Canada	16604	1880	1437	36086	15.7	34.3
Rogers Communications Canada Inc.	626	1637	1570	87703	9.8	23.3
Comcast Cable Communications, LLC	873	5992	3086	34203	12.8	5.5
PlusServer GmbH	125	136	6	1527	13.0	.3
Unwired	23	245	4	6434	12.7	.2

(20 rows)

Case Study Examples

- Learning Management System (LMS)
  - Desire-to-Learn (D2L) at University of Calgary
  - Moodle at University of Venice
- Video streaming applications
  - ASTRO 209
  - 360º video (Fri 10:00am at ICPE 2020)
- Online social networks
  - Instagram
- Electronic mail
  - IMAPS
  - Outlook (Office 365)
  - Spam filtering services
- Network services
  - Domain Name System (DNS)
  - Network Address Translation (NAT)

Summary and Conclusions

- If application performance debugging is an art, then network application debugging is a dark art!

- Many possible performance problems:
  - Client side
  - Network
  - Server side
- Protocol interaction effects are yet another factor

- The more you look, the more strange things you'll see!

The End!

- With many thanks to my students and colleagues:
  - Martin Arlitt, Xiaozhen (Jean) Cao, Mackenzie Haffey, Jennifer Harper, Mehdi Karamollahi, Sina Keshvadi, Steffen Berg Klenow, Michel Laterman, Rachel Mclean, Sean Picard, Masroor Syed, Zhengping Zhang, and UCIT

- For more information:
  - Email: carey@cpsc.ucalgary.ca
  - Web: http://www.cpsc.ucalgary.ca/~carey

- Thank you for listening!!
- Questions?

References

- M. Crovella and B. Krishnamurthy*, Internet Measurement: Infrastructure, Traffic, and Applications*, Wiley, 2006.
- M. Haffey, M. Arlitt, and C. Williamson, "Modeling, Analysis, and Characterization of Periodic Network Traffic", Proceedings of IEEE MASCOTS 2018, Milwaukee, WI, September 2018.
- M. Karamollahi and C. Williamson, “Characterization of IMAPS Email Traffic”, Proceedings of IEEE MASCOTS 2019, Rennes, France, pp. 214-220, October 2019.
- S. Keshvadi and C. Williamson, “MoVIE: A Measurement Tool for Mobile Video Streaming on Smartphones”, to appear, Proceedings of ACM/SPEC International Conference on Performance Engineering (ICPE), Edmonton, AB, April 2020.
- S. Klenow, C. Williamson, M. Arlitt, and S. Keshvadi, “Campus-Level Instagram Traffic: A Case Study”, Proceedings of IEEE MASCOTS 2019, Rennes, France, pp. 228-234, October 2019.
- J. Kurose and K. Ross, *Computer Networking: A Top-Down Approach*, 7th edition, Pearson, 2017.
- V. Paxson, "Bro: A System for Detecting Network Intruders in Real-time", *Computer Networks*, Vol. 31, No. 23, pp. 2435-2463, December 1999.
- V. Paxson, "Strategies for Sound Internet Measurement", Proceedings of ACM Internet Measurement Conference (IMC), Toarmina, Italy, October 2004.
- S. Roy, C. Williamson, and R. Mclean, “LMS Performance Issues: A Case Study of D2L”, ISCA Journal of Computers and Their Applications, Vol. 25, No. 3, September 2018.
- C. Williamson, "A Tutorial on Internet Traffic Measurement", *IEEE Internet Computing*, Vol. 5, No. 6, pp. 70-74, November/December 2001.
- Z. Zhang and C. Williamson, "A Campus-Level View of Outlook Email Traffic", Proceedings of the 7th International Conference on Network, Communication, and Computing (ICNCC 2018), Taipei, Taiwan, December 2018.

LMS Performance Study: U of Calgary

- Desire-to-Learn (D2L) is the official Learning Management System (LMS) at the University of Calgary (Spring 2014)
- Many faculty and students use D2L for their courses
- Context/Motivation:
  - Many universities use LMS (e.g., BlackBoard, D2L, Canvas, Moodle)
  - Few studies characterizing LMS usage and/or performance
  - Anecdotal reports suggest that D2L at U of C is “slow”
  - Network traffic measurement research provides a means to analyze, characterize, and understand D2L usage at U of C

Overview

- Network traffic measurement study of D2L usage (2015-2016)
- Combination of active and passive measurement approaches

- Research Questions:
  - How does D2L work?
  - How is D2L being used at the University of Calgary?
  - How can we improve the performance of D2L?

UNIVERSITY OF CALGARY

Results Preview

- Complex configuration of D2L setup at U of Calgary
  - Excessive HTTP redirection for session login and logout
- Long network RTT to access remotely hosted D2L content
  - Approximately 40 ms RTT to reach Kitchener, Ontario
  - No local CDN node at U of C; closest node is in Toronto
- Suboptimal configuration of TCP for D2L Web servers
  - Uploads and downloads are window-limited (64 KB per RTT)
  - D2L Web server seems very slow (IIS v7.5 on Windows 2000 R2)

Data Collection Methodology

- Data collection from Jan. 1, 2016 – April 30, 2016 (W2016)
  - Microscopic analysis was performed for this period
- Data collection from Jan. 1, 2015 – December 31, 2016
  - Longitudinal analysis was performed for this period
- Data was processed and stored in Bro logs
  - Records connection summaries for all TCP and UDP traffic
    - Connection logs provide inbound/outbound traffic information
    - HTTP logs provide user agent information for Web browsing
    - SSL logs provide server information and encryption details
- Active and passive measurement tools were used in this research

D2L Usage Patterns

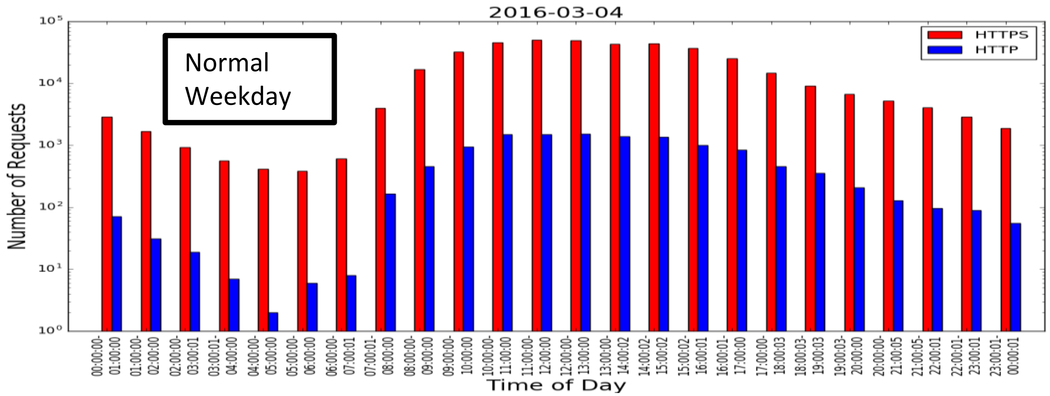

- This graph shows the number of requests made to D2L per hour over a one-day (24 hour) period
- Traffic pattern is diurnal
- Peak HTTPS traffic is 30x larger than that of HTTP traffic

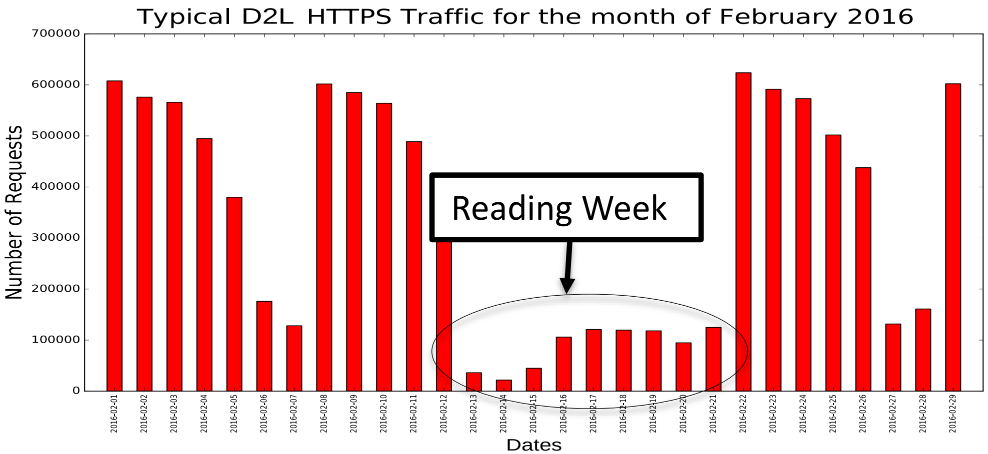

- This graph shows daily totals for D2L requests for one month
- Monday is the busiest day of the week for D2L traffic volume
- Request volume tends to decrease throughout the week
- Holidays have lower D2L traffic

D2L Configuration at U of Calgary

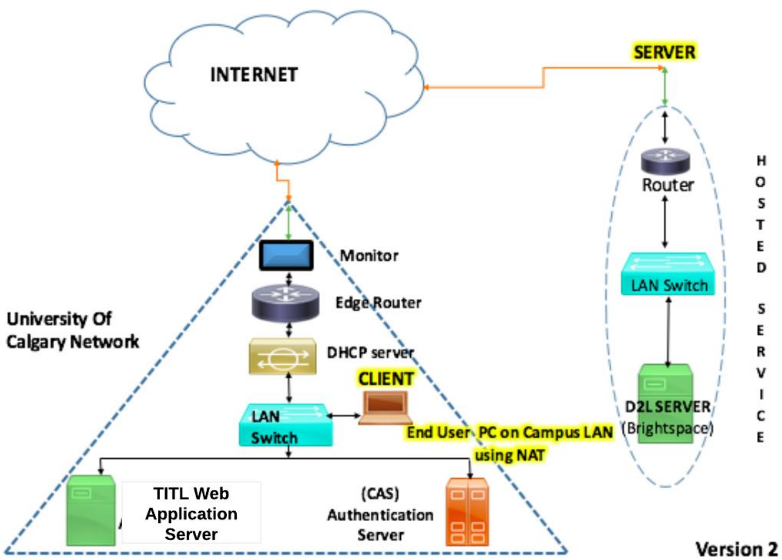

Typical Internet path for on-campus D2L users (including NAT, DHCP, wireless) spans 17 hops with 40 ms RTT

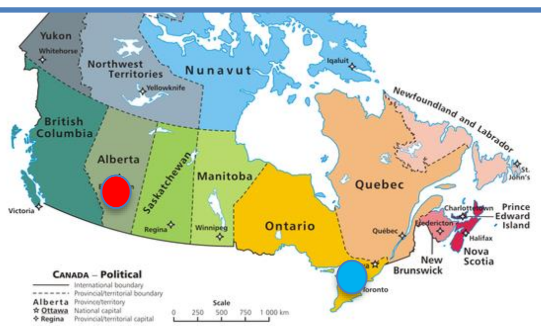

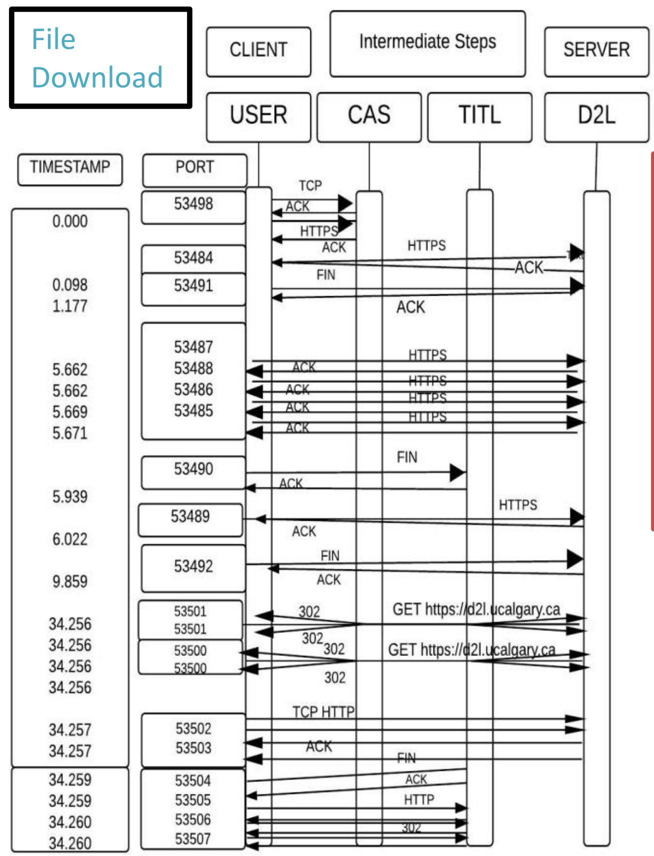

- Shows the role of intermediate servers
- Parallel connections seen when uploading or downloading files
- Persistent HTTP connections seen in D2L sessions

TCP Throughput: Downloads and Uploads

D2L File Download

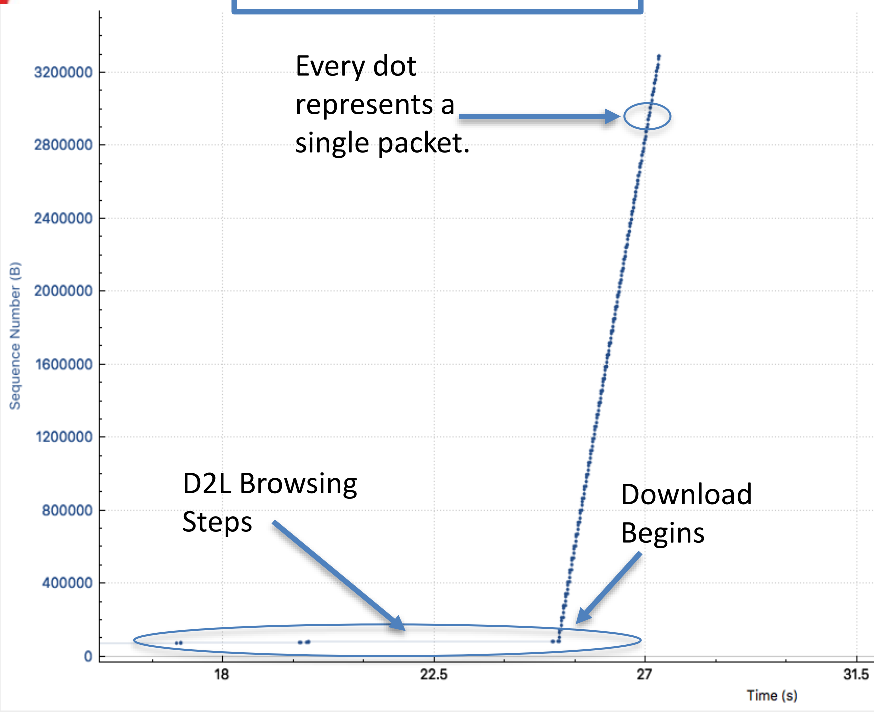

TCP Throughput: 14 Mbps
RTT Latency: 45 ms

D2L File Upload

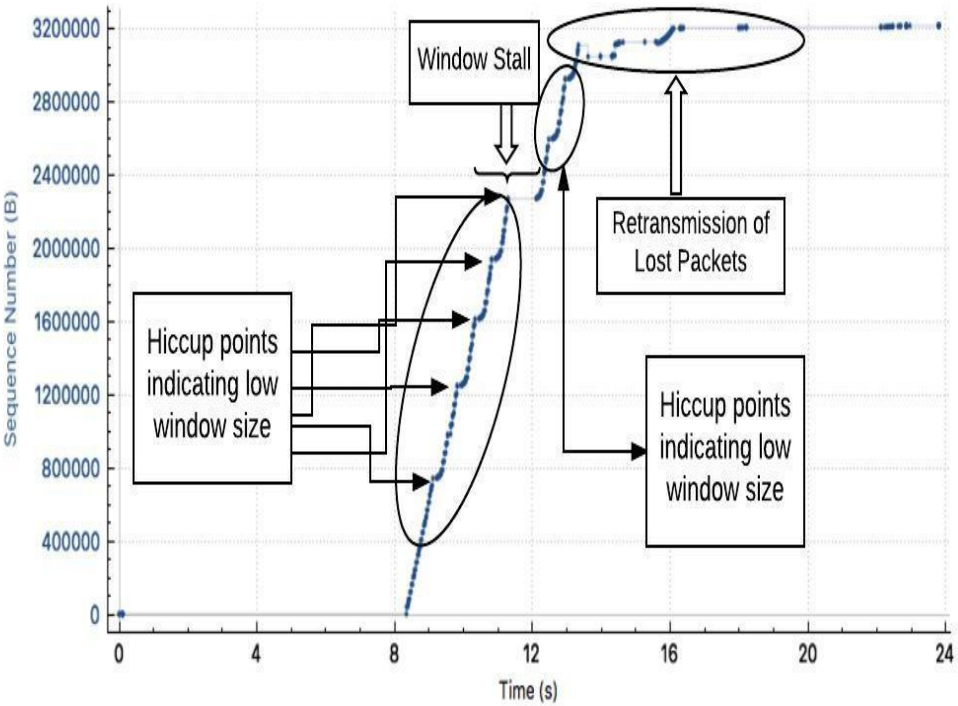

TCP Throughput: 7 Mbps

Results Summary

- Complex configuration of D2L setup at U of Calgary
  - Excessive HTTP redirection for session login and logout
- Long network RTT to remotely hosted D2L content
  - Approximately 40 ms RTT to reach Kitchener, Ontario
  - No local CDN node at U of C; closest node is in Toronto
- Suboptimal configuration of TCP for D2L Web servers
  - Uploads and downloads are window-limited (64 KB per RTT)
  - D2L Web server seems very slow (IIS v7.5 on Windows 2000 R2)

Lessons Learned

- Network traffic measurement can provide a better understanding of the usage and performance of LMS services like D2L
- D2L at the U of C has a rather complex delivery infrastructure, and several idiosyncracies that affect its user-perceived performance
- Long network latencies make remotely hosted content painful!

- Proposed solutions:
  - Having a local CDN node could improve D2L performance
  - Improving TCP configuration (version and/or socket buffer sizes) could improve throughput for D2L users
  - Faster servers (e.g., Amazon Web Services)

Return

Moodle Issue: University of Venice

- Moodle is the LMS at Ca' Foscari University (Venice)
- I was there as a Visiting Professor in November 2019
- What I noticed with their LMS:
  - Downloads were fine
  - Uploads unbelievably slow (about 70-75 sec per file)

- Network traffic measurement to the rescue!!

- Root cause: configuration error for virus scanning
- Reported and fixed! Uploads are 20x faster now ☺

Return

UNIVERSITY OF CALGARY

Instagram Case Study

- **University of Calgary** in Calgary, Alberta, Canada
  - 35,000 students (ugrad/grad)
  - 3,000 faculty/staff
- One week: Sunday **March 3**, 2019 to Saturday **March 9**, 2019

Active Measurement Results

- Over 90% of the Instagram-related requests go to a single IP: **157.240.3.63**
- All main features use the same IP address
- Monitoring this single IP address gives a good estimate (but slight underestimate) of the campus-level Instagram traffic!

Observed DNS host names:

- i.instagram.com
- platform.instagram.com
- instagram.c10r.facebook.com
- scontent-sea1-1.cdninstagram.com
- graph.instagram.com

Passive Measurement Results (1 Week)

Item Description	Sun Mar 3	Mon Mar 4	Tue Mar 5	Wed Mar 6	Thu Mar 7	Fri Mar 8	Sat Mar 9	Overall
TCP Connections	896,849	2,355,640	2,313,701	2,352,614	2,253,556	2,055,827	853,820	13.1 M
Mean Duration	78.7 s	72.1 s	71.9 s	72.0 s	72.3 s	73.4 s	76.7 s	72.3 s
Packets Sent	264.3 M	565.3 M	565.2 M	561.9 M	550.3 M	509.0 M	283.3 M	3.3 B
Packets Received	550.9 M	1,003 M	953.9 M	931.1 M	950.7 M	910.2 M	589.9 M	5.9 B
Bytes Sent	32.2 GB	63.4 GB	60.4 GB	60.2 GB	60.0 GB	57.3 GB	33.3 GB	367 GB
Bytes Received	695 GB	1,259 GB	1,196 GB	1,167 GB	1,193 GB	1,141 GB	744.5 GB	7.2 TB
Client IP Addresses	1,450	1,679	1,605	1,532	1,621	1,547	1,449	3,498
IP Subnets	31	60	53	49	59	52	49	81

U of C Instagram Traffic Profile

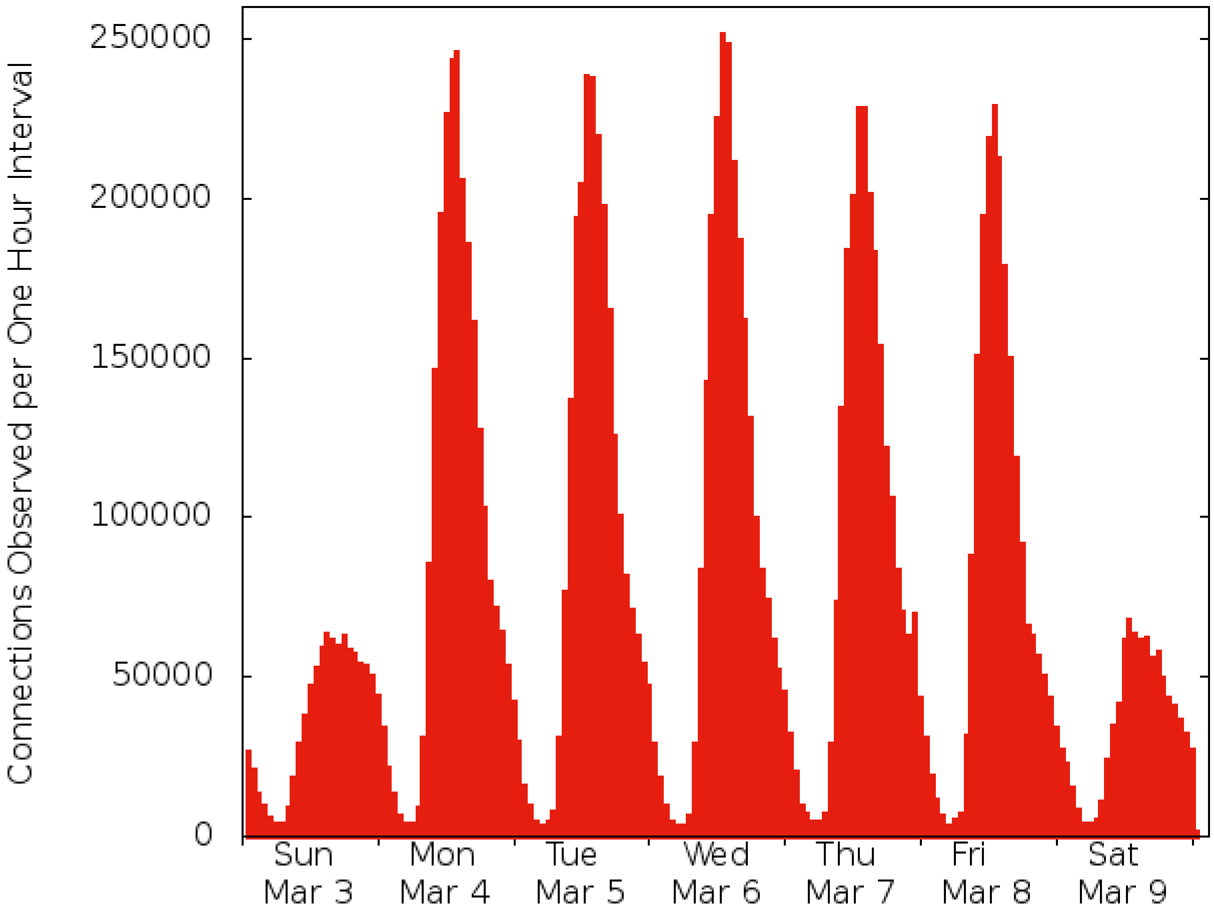

Observed TCP Connection States

State Description	Conns	%Conns	Bytes	%Bytes
SF: SYN-FIN	6,265,336	47.88%	3.78 TB	52.55%
RSTO: origin reset	2,487,505	19.01%	1.74 TB	22.91%
S3: no FIN seen	1,554,591	11.88%	879.9 GB	11.21%
S2: client FIN only	595,772	4.55%	340.1 GB	4.38%
S1: server FIN only	498,635	3.81%	189.7 GB	2.33%
RSTOS0: fail/RSTO	354,775	2.71%	222.9 GB	2.87%
RSTR: rcvr reset	335,304	2.56%	49.2 GB	0.63%
SH: no SYN-ACK	294,300	2.25%	107.1 GB	1.37%
SHR: no SYN seen	273,951	2.09%	57.3 GB	0.74%
OTH: other state	201,788	1.54%	71.3 GB	0.92%
S0: failed setup	166,822	1.27%	0.03 GB	$< 0.01\%$
REJ: rejected	37,455	0.29%	4.5 GB	0.06%
RSTRH: rcvr reset	20,329	0.16%	2.0 GB	0.03%
Total	13,086,563	100.0%	7.5 TB	100.0%

Time Series Illustration of TCP Connection States

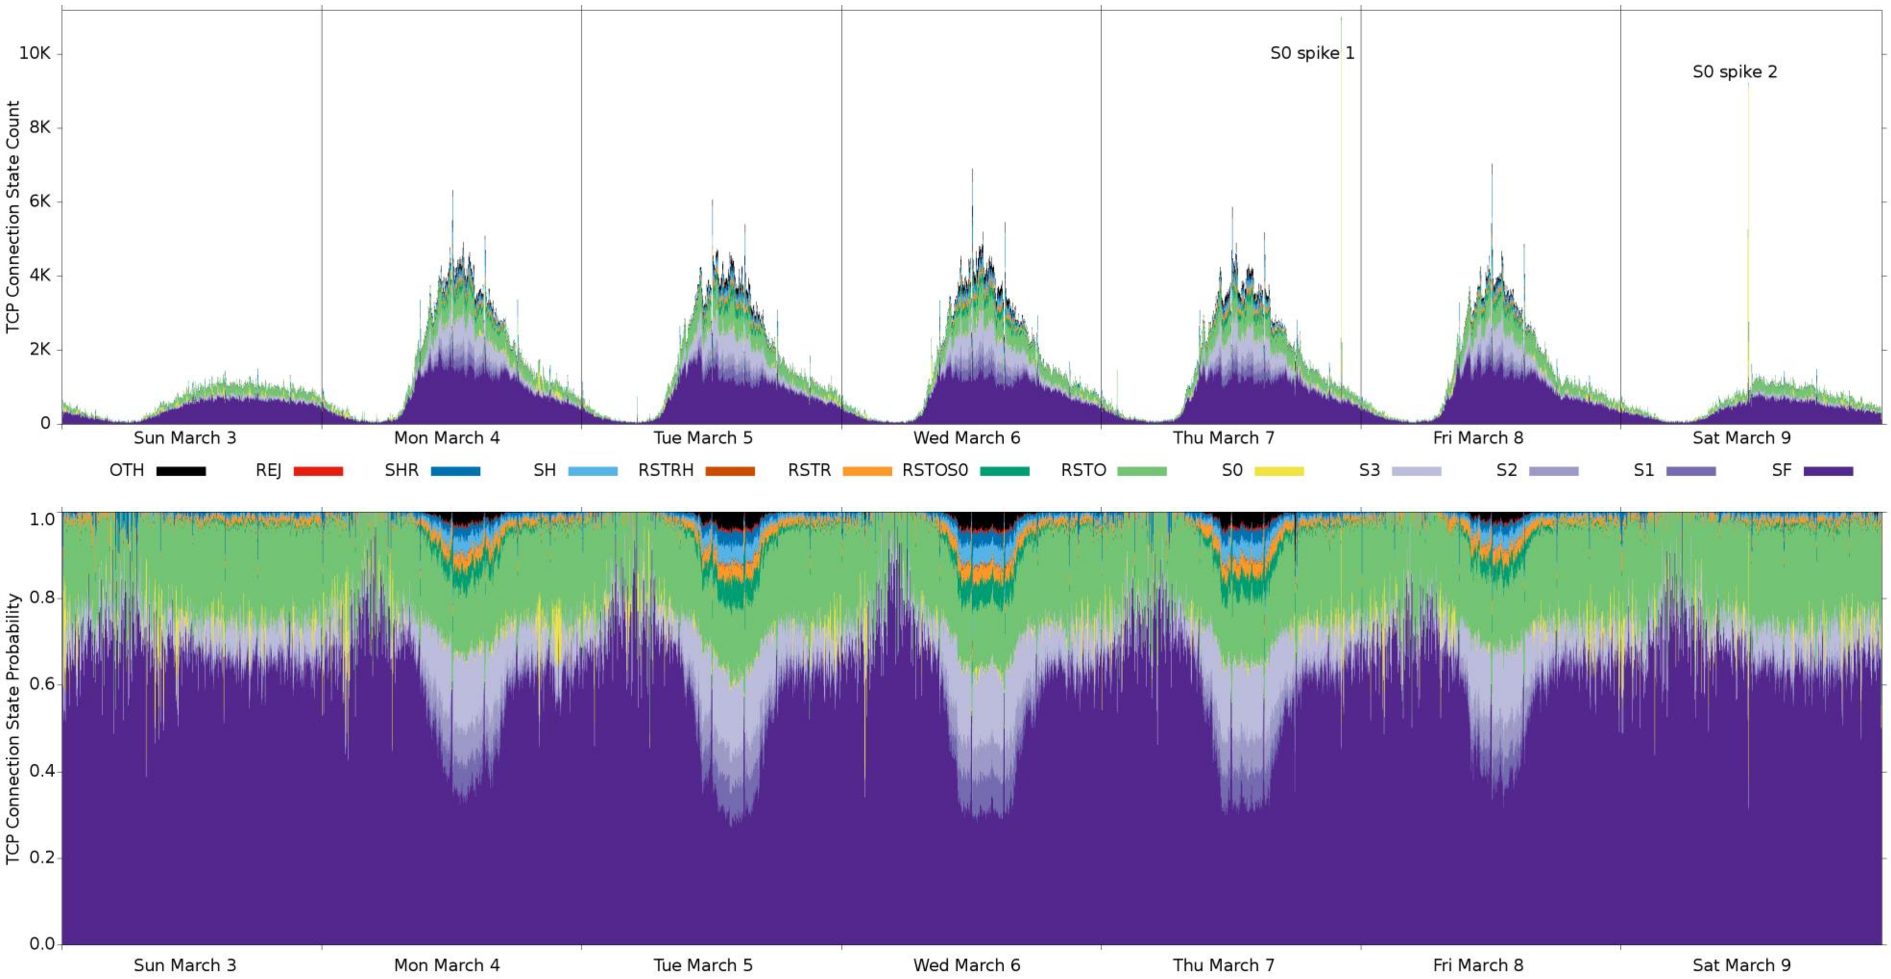

Instagram Traffic Anomaly (1 of 3)

A brief network router outage for about 80 seconds on Saturday March 9, which affected Instagram traffic, and other network services

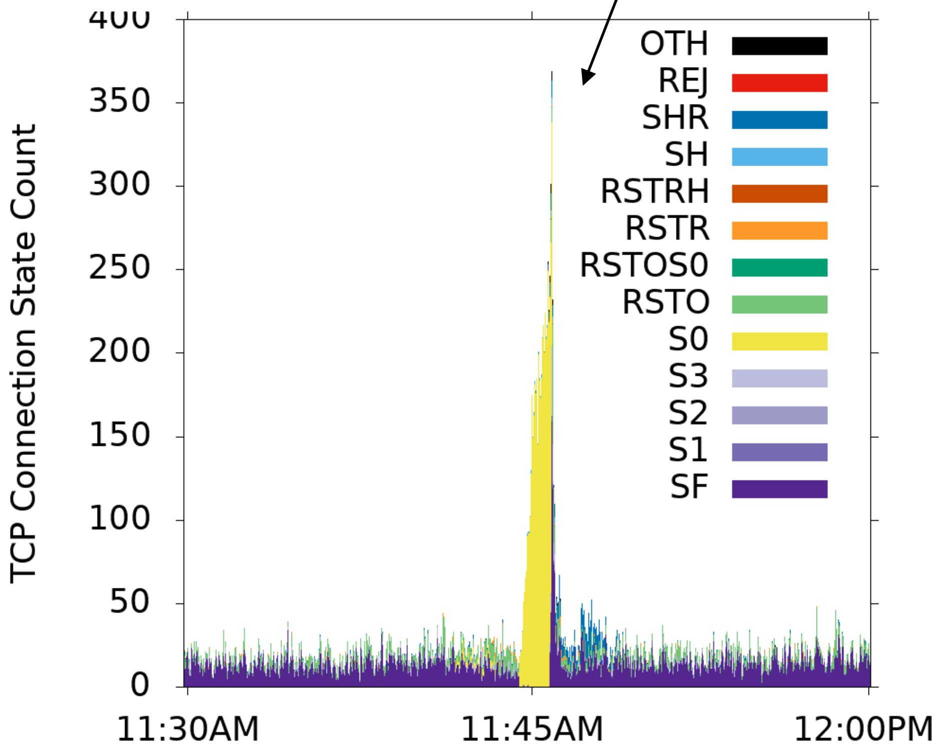

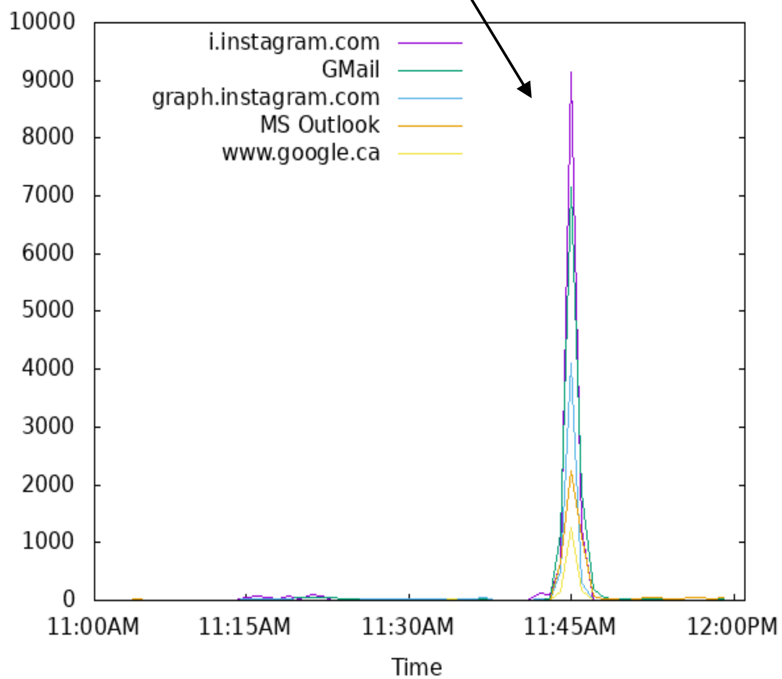

Instagram Traffic Anomaly (2 of 3)

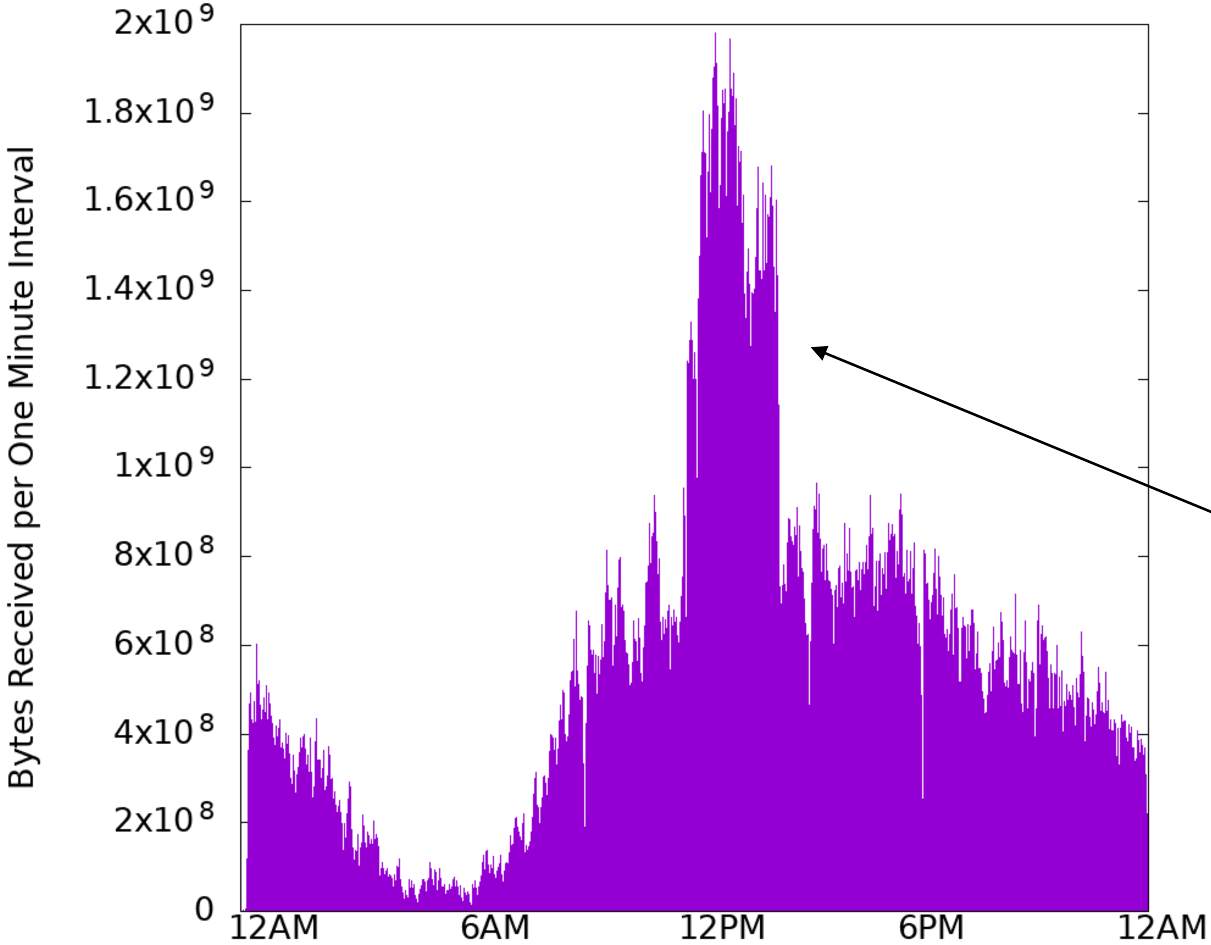

Multiple UCalgary IP addresses partaking in some Instagram video streaming event for about 2 hours on Sat March 9

Instagram Traffic Anomaly (3 of 3)

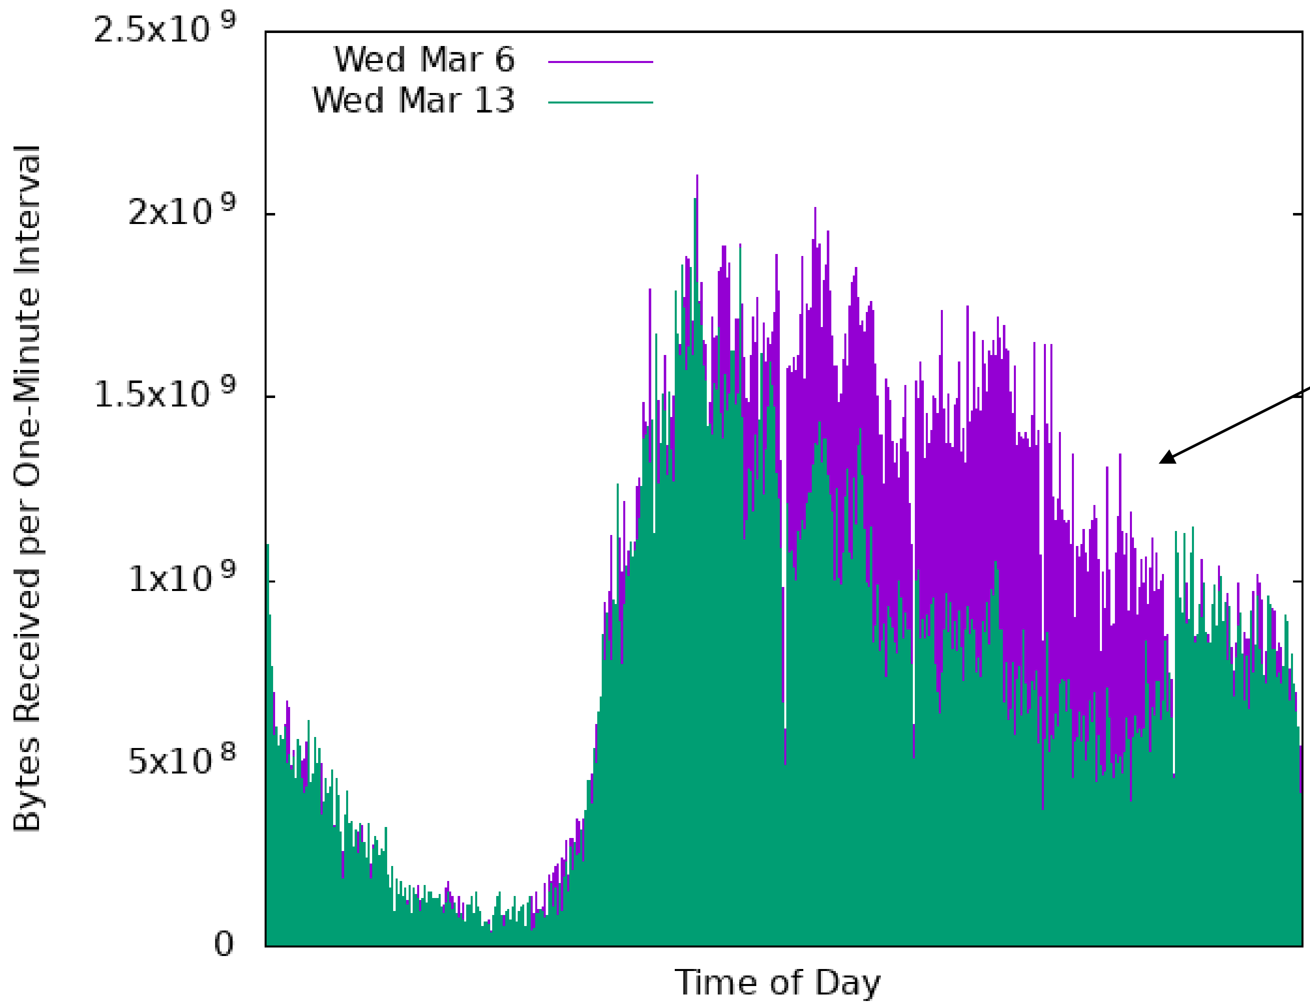

Partial outage for Facebook, WhatsApp, and Instagram for several hours on Wed March 13/19

Lessons Learned

- On our campus network, a typical weekday of Instagram traffic has:
  - **1 TB** of data **downloaded**
  - **60 GB** of data **uploaded**
- Third highest bandwidth consumption behind Netflix (6 TB per day) and YouTube (3 TB per day)
- Highly skewed distributions:
  - high variability (e.g., transfer sizes, throughputs)
  - heavy-tails (e.g., connection durations, transfer sizes)
- This traffic can have a large impact on a campus edge network!

Return

ASTRO 209 Experiment

Introduction

What? Astronomy: The Cosmos

First-year undergraduate course with 400 students
Taught in Winter 2015 (Jan-April)
Web site: notes, slides, linked rich media (70 GB/day)

Why? Workload Characterization

Understand how students use educational Web sites
Characterize network traffic and identify performance issues

How? Passive Measurement

ISM Server: CentOS , Apache Web Server, Port 80
Monitor: Dell, 2 Intel Xeon, Endace DAG 8.1SX card, Bro logs

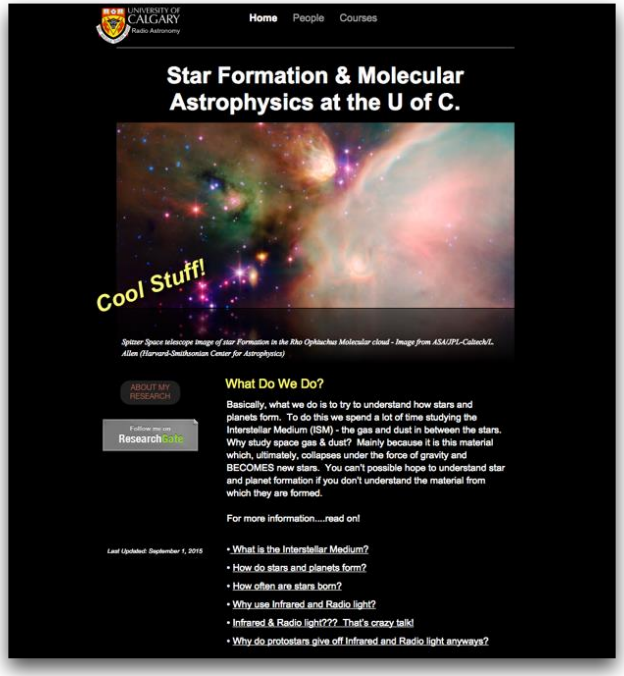

Measurement Results: Overview

HTTP Requests:
1,583,339
13,305 reqs/day

Unique IPs:
9,720

HTTP Method:

GET	99.5%
HEAD	0.5%

Data Volume:
8,483 GB
71.29 GB/day

Unique URLs:
10,563

Status Code:

200	32.04%
206	58.59%

Measurement Results: HTTP Requests per Day

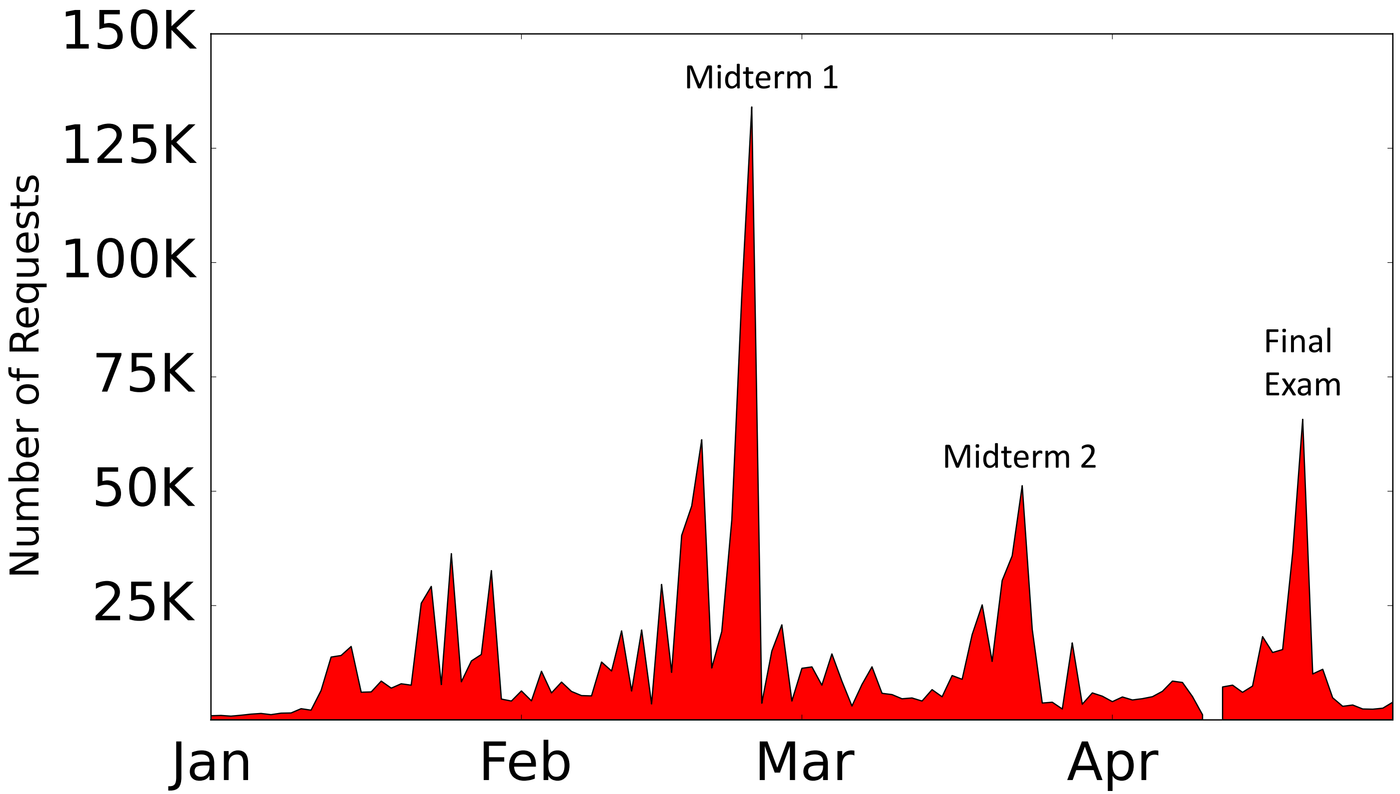

Measurement Results: Data Volume (GB/day)

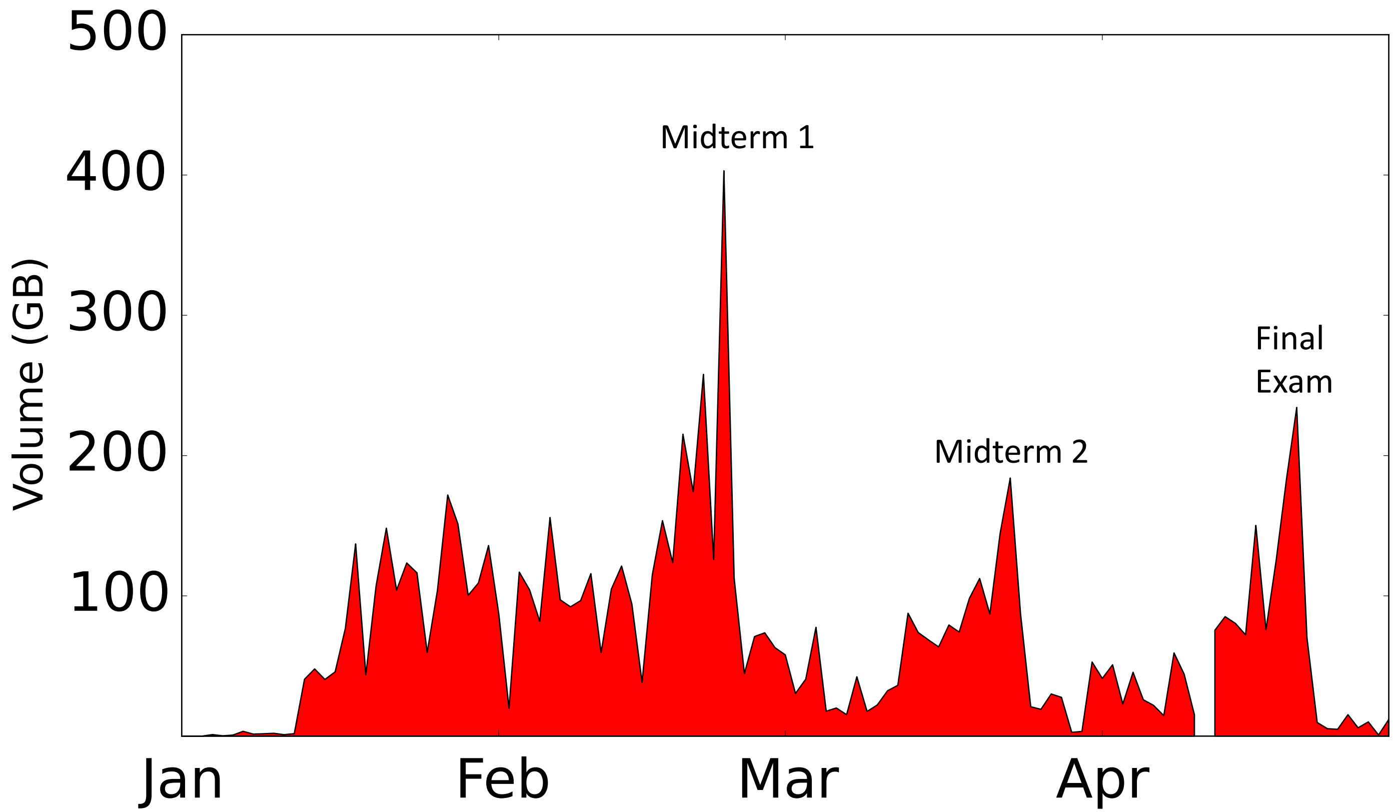

UNIVERSITY OF CALGARY

Measurement Results: HTTP Usage

HTTP Method

HTTP Method	Reqs	Pct.
GET	1,575,574	99.51%
HEAD	7,749	0.49%
OPTIONS	11	0.00%
POST	5	0.00%

HTTP Status Code

Status Code	Reqs	Pct.
206 Partial Content	927,733	58.59 %
200 OK	507,358	32.04 %
304 Not Modified	79,064	4.99%
404 Not Found	47,372	2.99%

Measurement Results: URL Analysis and File Types

Top 5 Requested URLs

URL	Total Reqs	Total GB
ASTR209 - Lec8 - Feb 5, 2015.mov	153,410	267.04
ASTR209 - Lec3 - Jan 20, 2015.mov	87,051	787.02
ASTR209 - Intro. & Lecture#1 - Jan 13,2015.mov	75,380	735.64
ASTR209 - Lec4 - Jan 22, 2015.mov	68,609	584.47
AST209 Podcast/rss.xml	56,293	0.71

Top 5 Requested File Types

File Type	Rank	Total Reqs	Pct.	Rank	Total GB	Pct.
Video/QuickTime	1	532,883	29.78%	1	5,159	60.35%
Application/PDF	2	250,244	13.99%	3	284	3.33%
Video/MP4	3	183,636	10.26%	2	3,082	36.06%
Text/HTML	4	177,506	9.92%	6	3	0.03%
Image/PNG	5	144,361	8.07%	5	4	0.05%

Measurement Results: Course-related Events

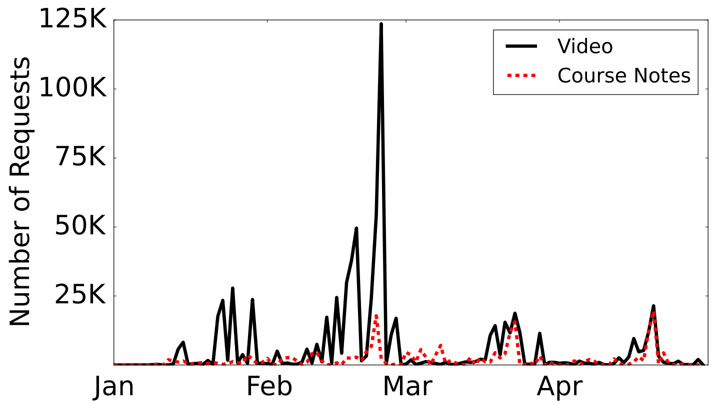

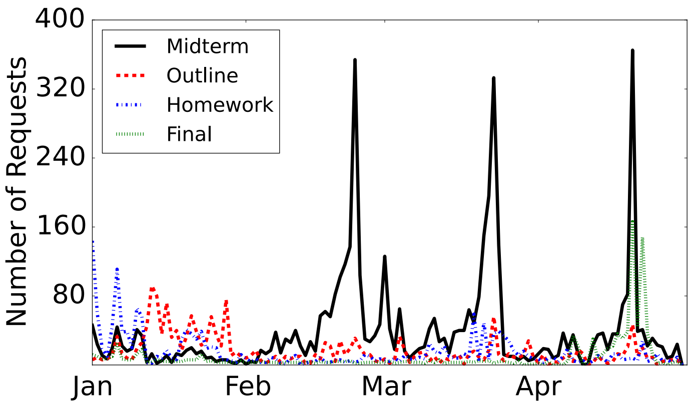

Requests

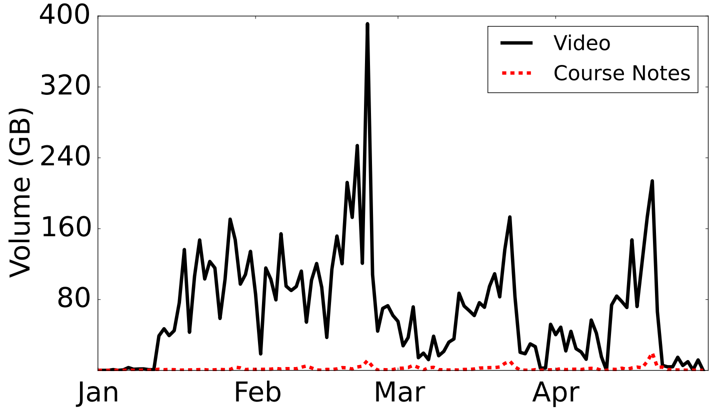

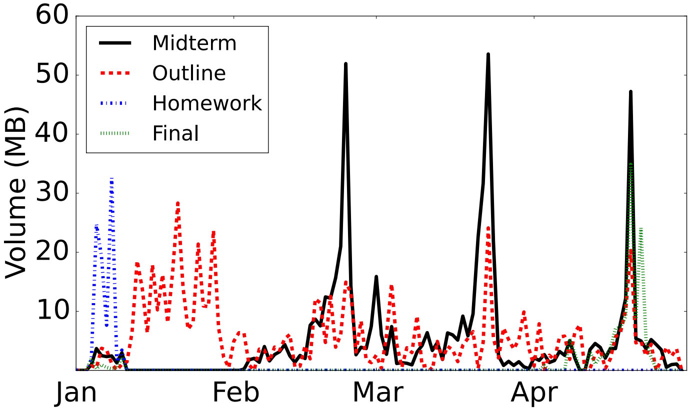

Data Volume

Measurement Results: User Agents

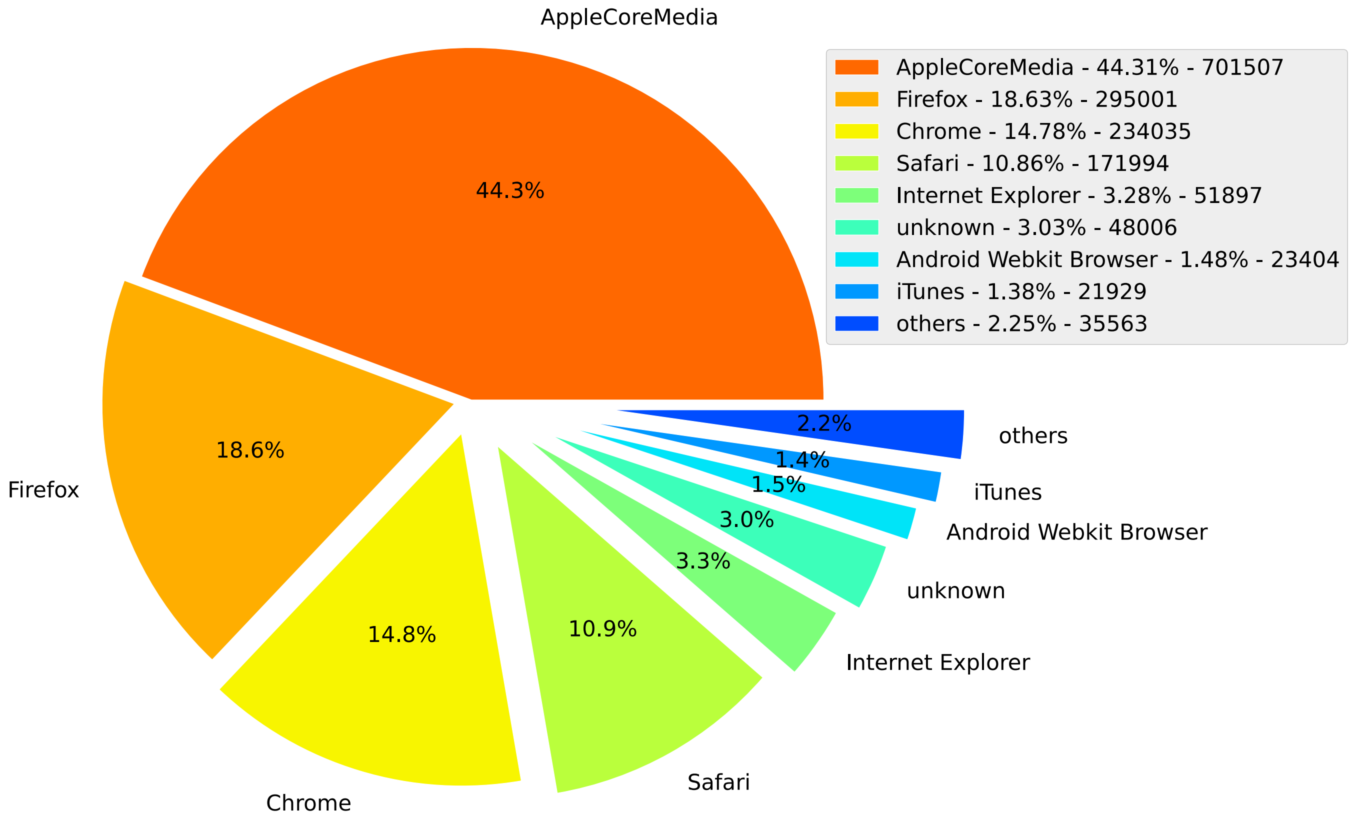

Active Measurements: Web Browser Experiments

Browser	Static File		<object> Element		HTML5 Video Tag	
	Play	Forward	Play	Forward	Play	Forward
Chrome (v44)	Yes	Yes	No	N/A	Yes	Yes
Safari (v8)	Yes	Yes	Yes	No	Yes	Yes
Firefox (v39)	Yes	Yes	Yes	No	Yes	Yes
IE (v11)	No	N/A	No	N/A	Yes	Yes

Lessons Learned

Learning-Related

1. First-year students are a technologically-savvy audience.
2. Study habits of students are reflected in their Web traffic.
3. Studying patterns changed for second midterm and final exam.

Technology-Related

1. Rich media Web sites can generate a LOT of network traffic.
2. Course-related events strongly influence the Web traffic.
3. Specific video configurations can adversely affect user experience and the network traffic.

Return

Case Study: Internet Mail Access Protocol (IMAPS) Traffic

- University of Calgary in Calgary, Alberta, Canada
- 35,000 students (grad/ugrad) and 3,000 faculty/staff
- Dates: Sunday April 14, 2019 to Saturday April 20, 2019

Statistical Summary of Email Traffic (1 week)

UNIVERSITY OF CALGARY

Protocol	Port	Dest	TCP Conns	Data Volume
HTTPS	443	Outlook	86,854,649	5.9 TB
IMAP	143	All	2,726,213	18.5 GB
IMAPS	993	All	11,901,742	530 GB
IMAPS	993	Outlook	791,746	7.3 GB
POP2	109	All	490,708	52.0 MB
POP3	110	All	2,479,096	10.9 GB
POPS	995	All	1,652,011	6.8 GB
SMTP	25	All	11,306,154	49.7 GB
SMTP	587	All	5,860,459	9.8 GB
SMTPS	465	All	5,015,557	7.0 GB

Main Observations: lots of TCP connections; high data volumes; IMAPS << HTTPS

IMAPS Destination IP Analysis

TCP Connections by Dest IPs

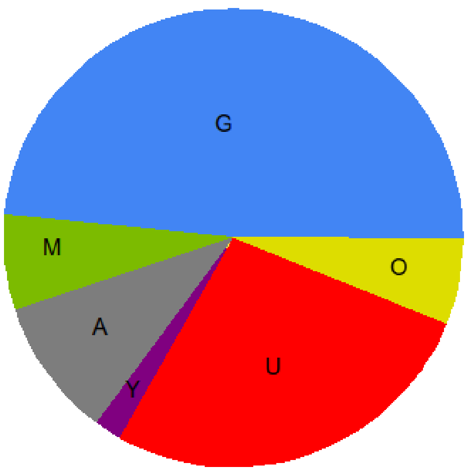

(G) Google 48.4%
(U) UofC 27.2%
(A) Apple 9.8%
(M) Microsoft 6.7%
(O) Others 27.2%
(Y) Yahoo 1.9%

Data Volume by Dest IPs

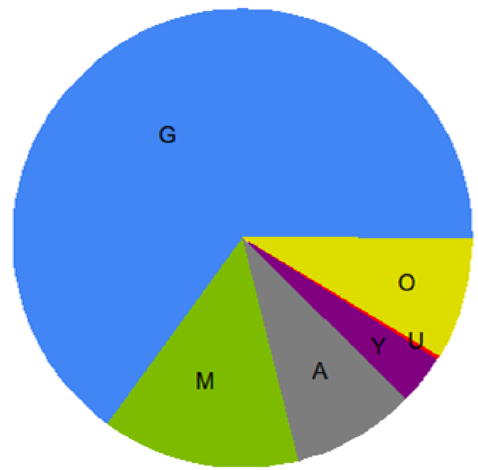

(G) Google 65.0%
(M) Microsoft 13.9%
(A) Apple 8.7%
(O) Others 8.6%
(Y) Yahoo 3.6%
(U) UofC 0.2%

Main Observations: lots of gmail traffic; many other email providers too

Transfer Sizes

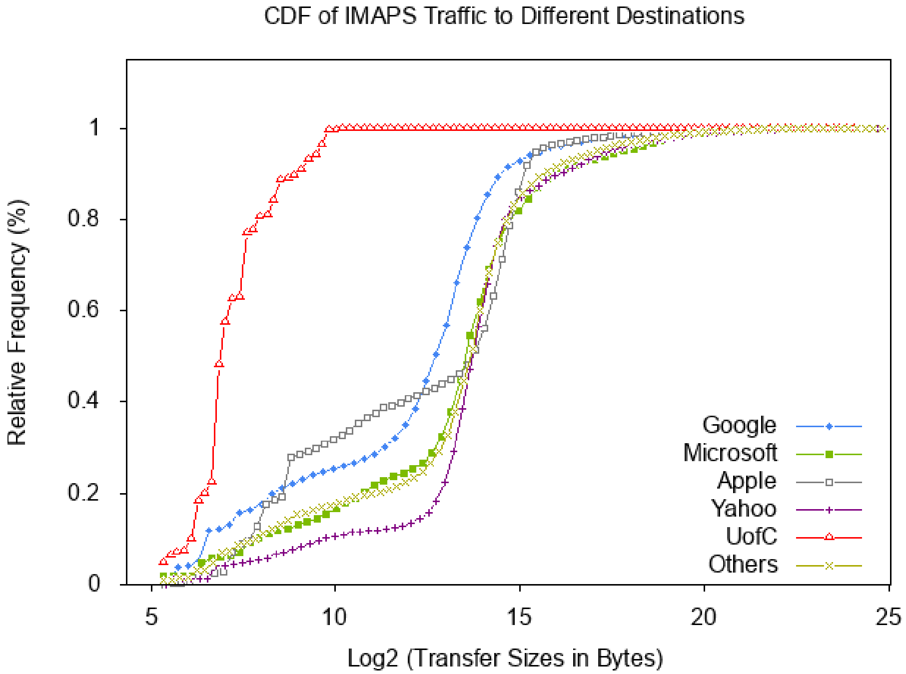

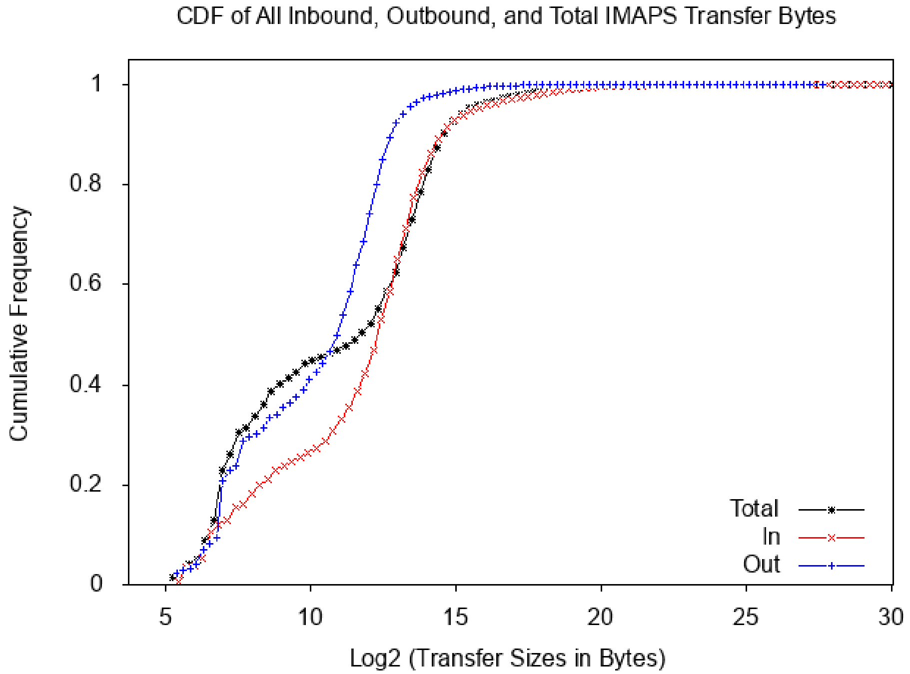

Main Observation: graphical evidence of heavy-tailed transfer sizes for IMAPS traffic

Throughput

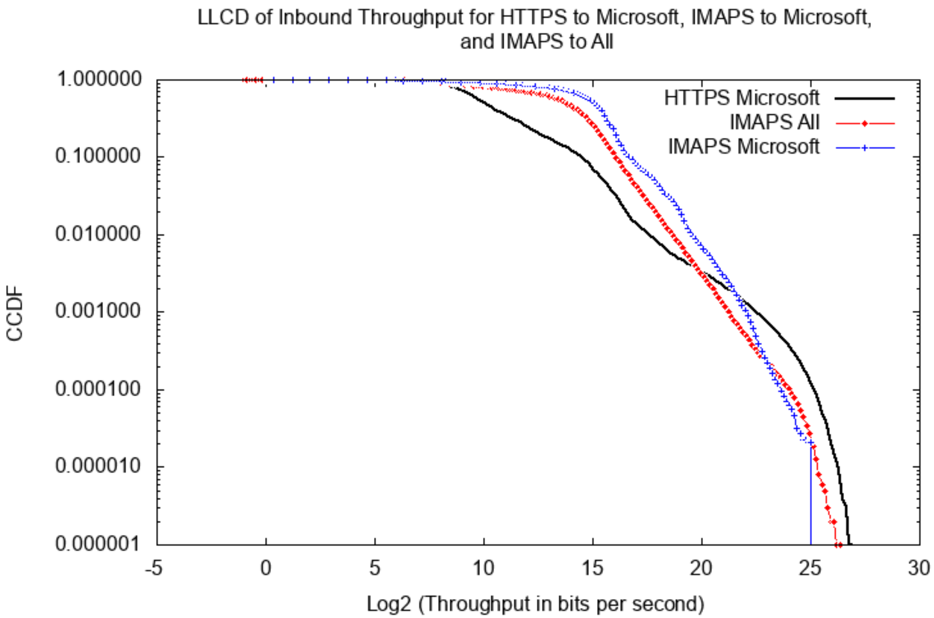

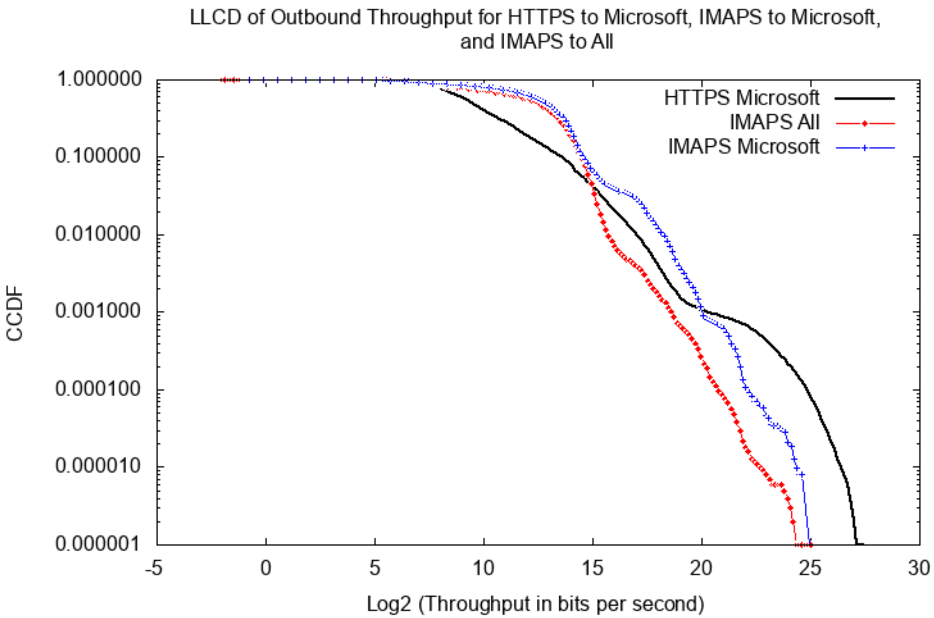

Observations: IMAPS throughput often higher than HTTPS, but varies with size and time of day

Lessons Learned

- There are strong diurnal and weekly patterns in email traffic
  - IMAPS has noticable spikes in connection traffic and data volumes
- Email protocols such as IMAPS and SMTP are highly asymmetric in their data transfers, while HTTPS is more symmetric
- High variability in transfer sizes, conn duration, and throughput
  - Evidence of heavy-tailed distributions in inbound/outbound transfer sizes
- Email traffic is highly complex: non-trivial workload models are required to capture these characteristics in network simulations

Return

Outlook (Office 365) Email Study

- Email Service
- Cloud Based Email Service
  - for economic and security reasons
- Outlook Email Service
  - powered by Microsoft
- 32,000 students + 3,000 Faculty/Staff
  - migrated to Outlook since 2014

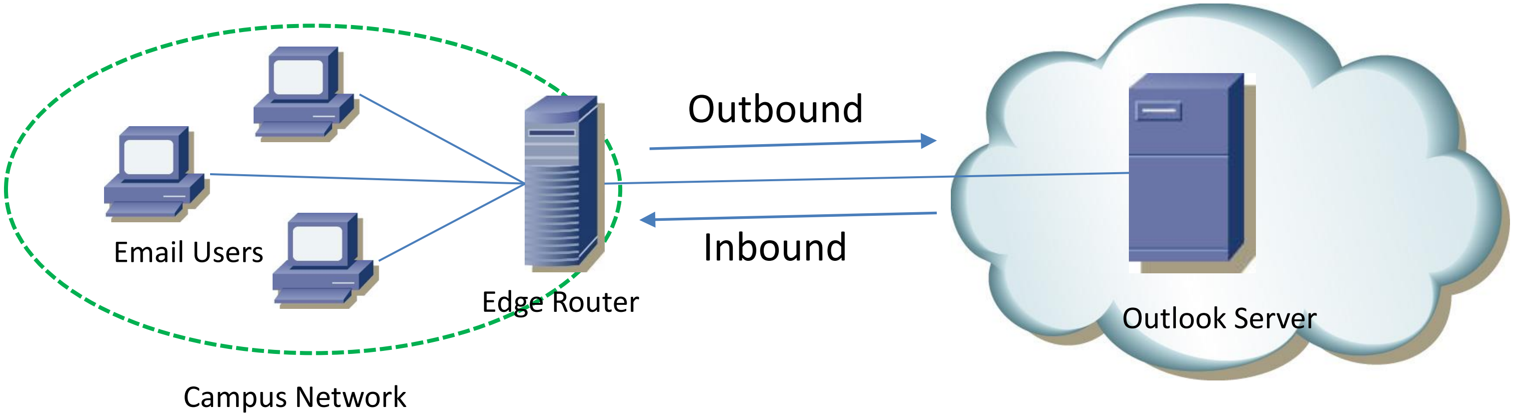

Server Classification

- Main Servers
  - Major servers, have several responsibilities
  - *outlook.office365.com*
- CDN Nodes
  - Deliver shared content such as icons, scripts, etc. <public access>
  - *r1.res* & *r4.res*
- Protection Servers
  - Spam filtering
  - Only talk with SMTP server
- Authentication Servers

Email Session Structure (1 of 3)

- 3 main approaches
  - Web/Client/Mobile Client
- 5 major steps
  - Login/Auth./Sending/Receiving/Logout

- Step 1. Login
  - *outlook.office365.com* or *outlook.office.com*
  - Several Parallel connections with different servers
  - Connections with other servers (*aria, nexus, skype,* etc.)
  - Central Authentication Server (CAS)

Email Session Structure (2 of 3)

- Step 2. Authentication
  - Microsoft Auth. <===> CAS at UofC *(fed.ucalgary.ca)*

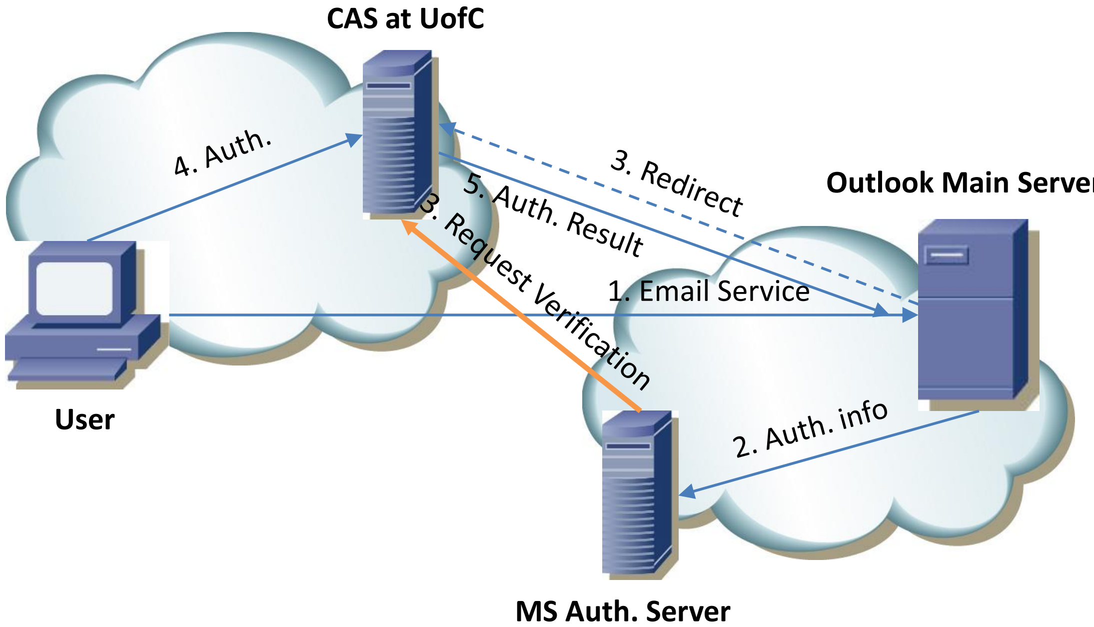

Email Session Structure (3 of 3)

- Step 3. Email Sending
  - HTTP POST
  - Attachment server for small attachments, OneDrive for large attachments (20 MB or more)
- Step 4. Email Receiving
  - **Periodical** HTTP POST
  - Similar Request Header
  - More frequently in web-based Outlook (10s)
- Step 5. Logout
  - sign out/close, same effect
  - FIN/ACK or **RST**
  - ~40% with RST (Main Server)

Performance Issues

- Extraneous TCP connections
  - Skype, Delve
  - Slows down the initialization step
- Improper use of RST
  - ~40% of Server connections
- Limited throughput
  - Different TCP window size for inbound and outbound
  - Maximum achievable throughput for inbound is 12 Mbps

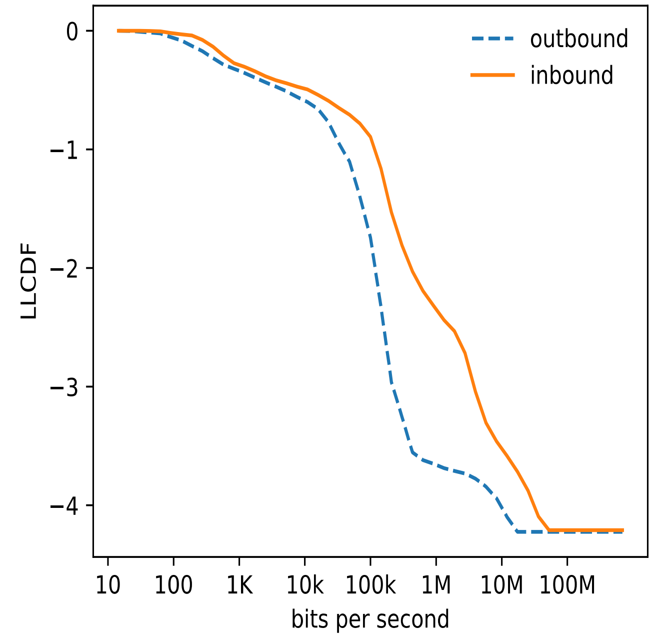

Lessons Learned

- Outlook email traffic measurement study at UofC
  - Four different servers
  - Five major steps
- Workload characterization of overall traffic
- Detailed analysis of session duration and data volume
- Potential performance issues with Outlook
  - Extraneous TCP connections
  - Excessive use of TCP RST
  - TCP window size issues

Return

Case Study: Email Delivery Delay

- Empirical observation: extremely long tail to the distribution of email delivery delay (i.e., elapsed time between "Send" and "Receive", as calculated from the SMTP headers in an empirical email dataset)

- My (incorrect) hunch: spam filtering service (Outlook)

- Student project: benchmarking spam filters

- Result: Main culprit is the mailman.ucalgary.ca service used for mailing lists on campus!

Empirical CDF for Email Delivery Delay

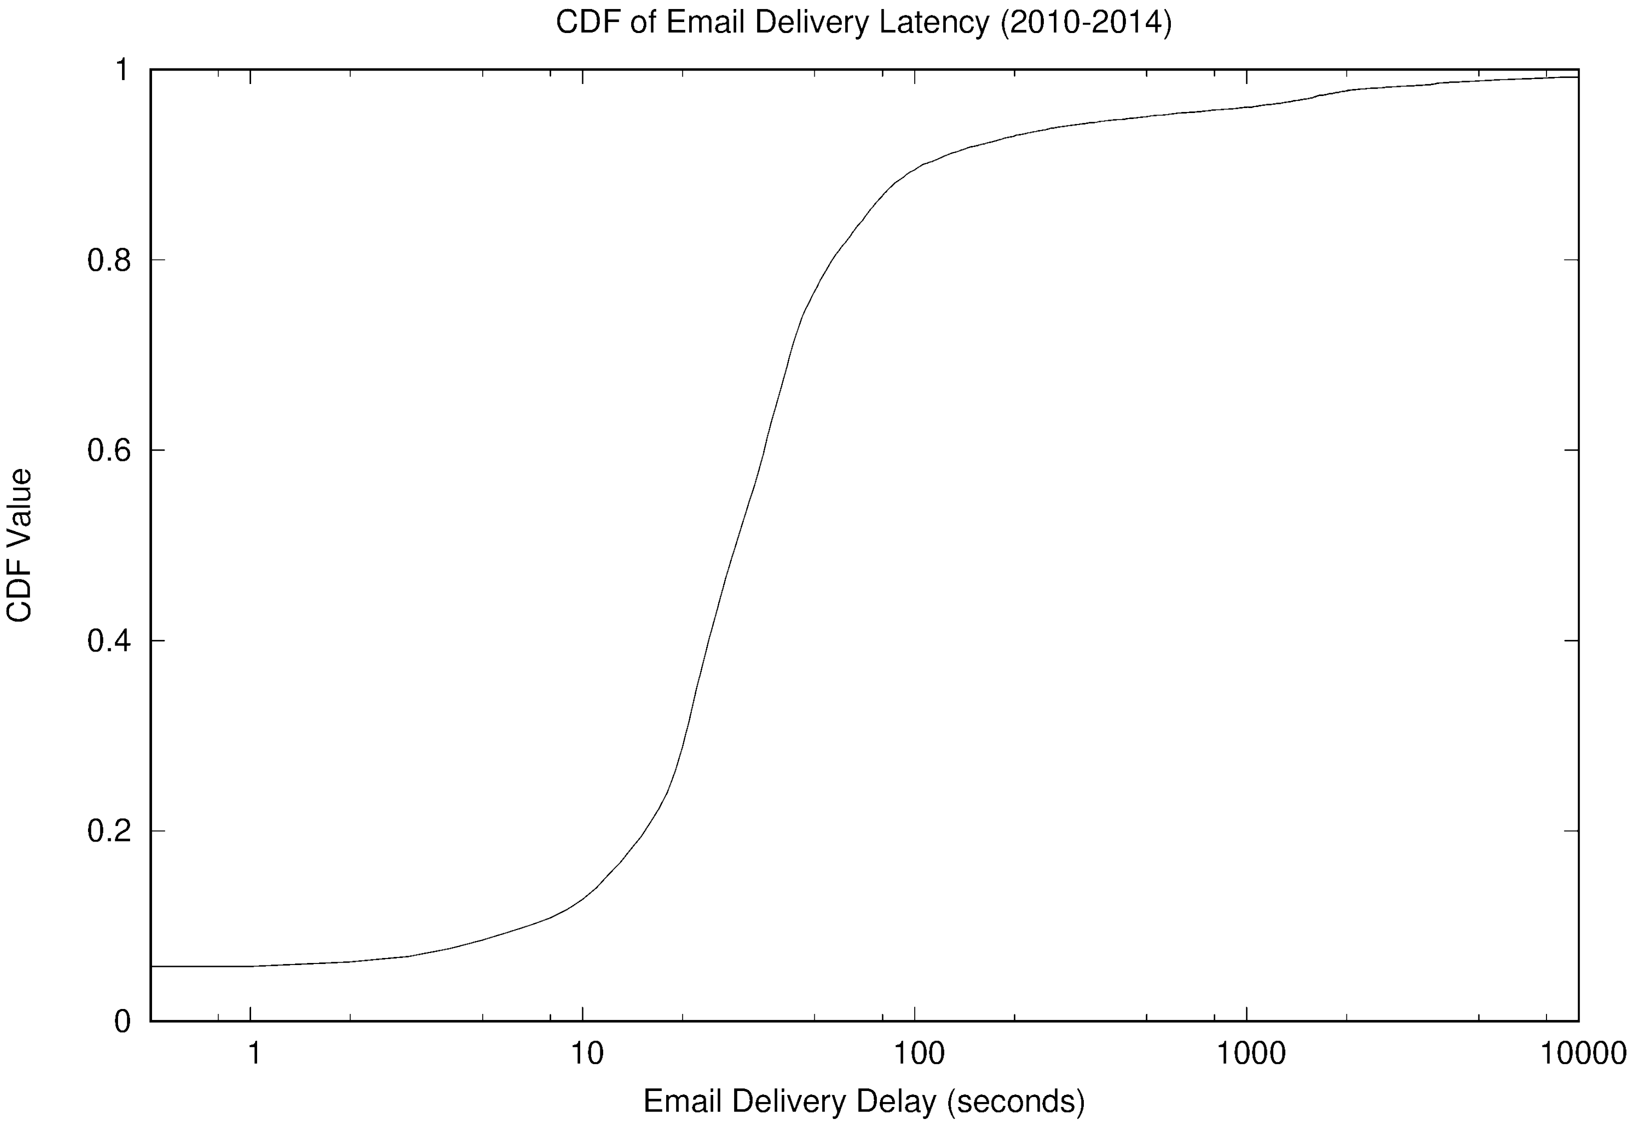

Email Delivery Delay with PDF Attachments

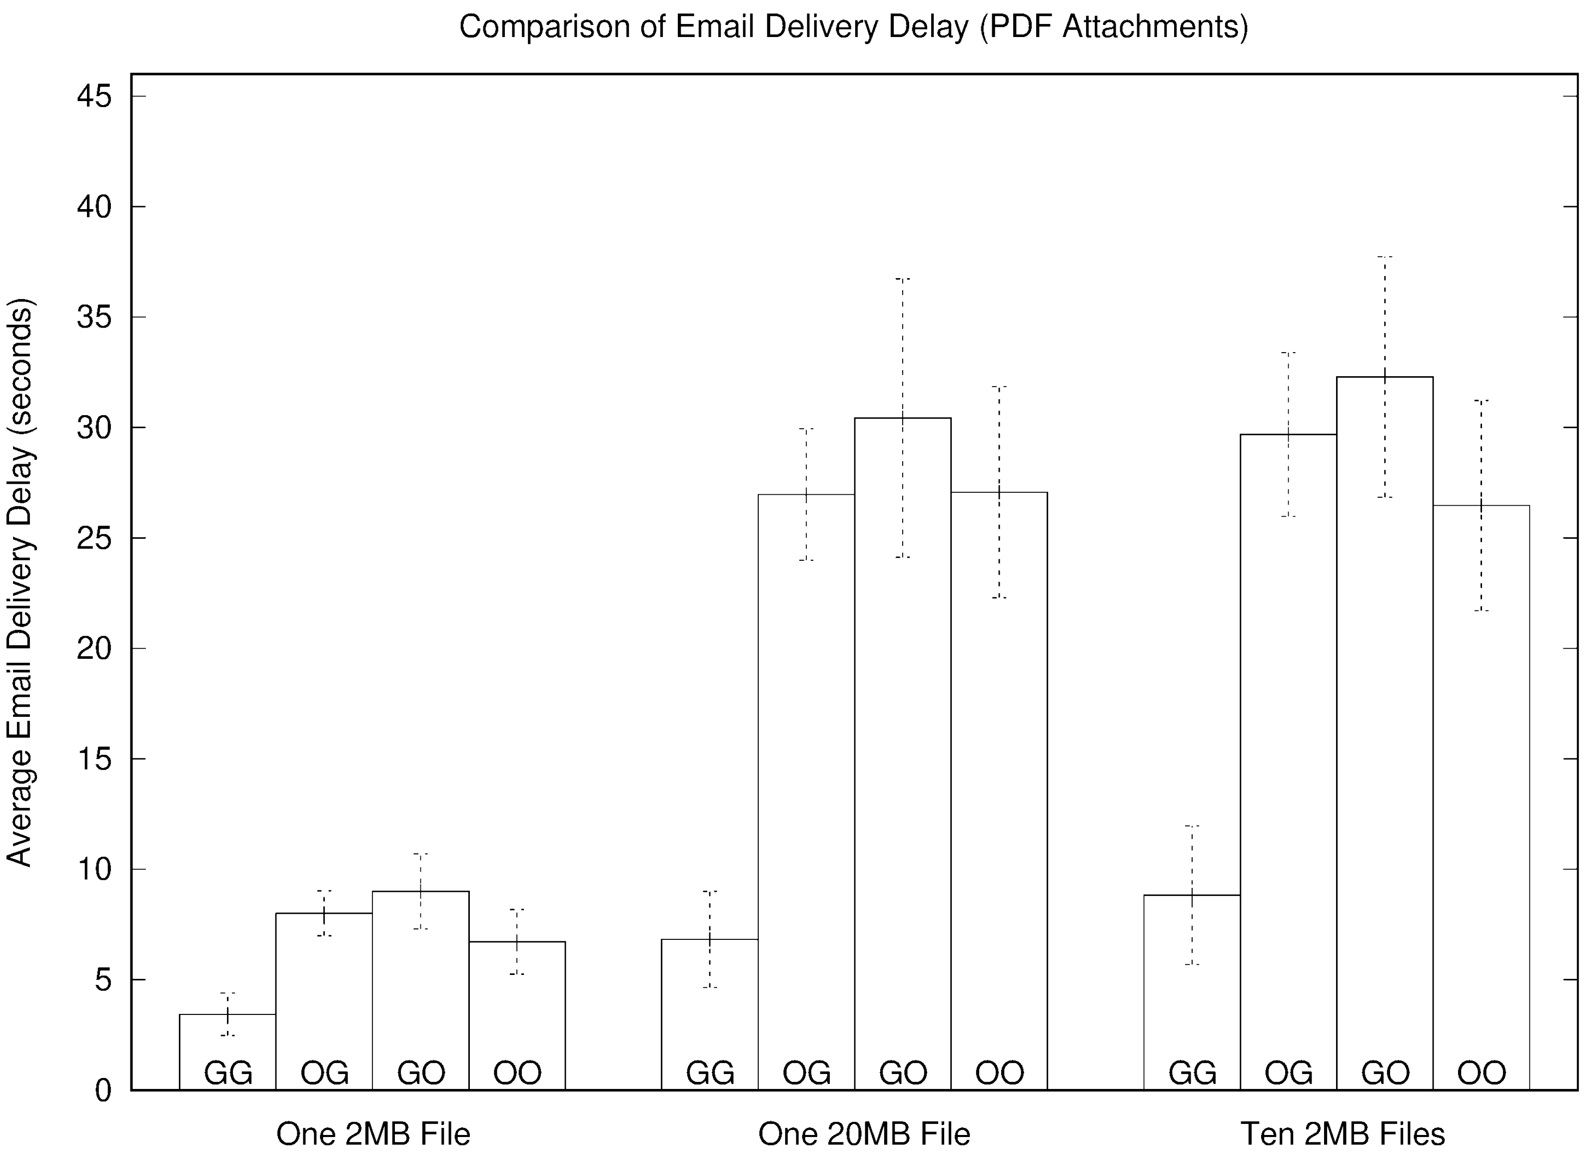

Email Delivery Delay with JPG Attachments

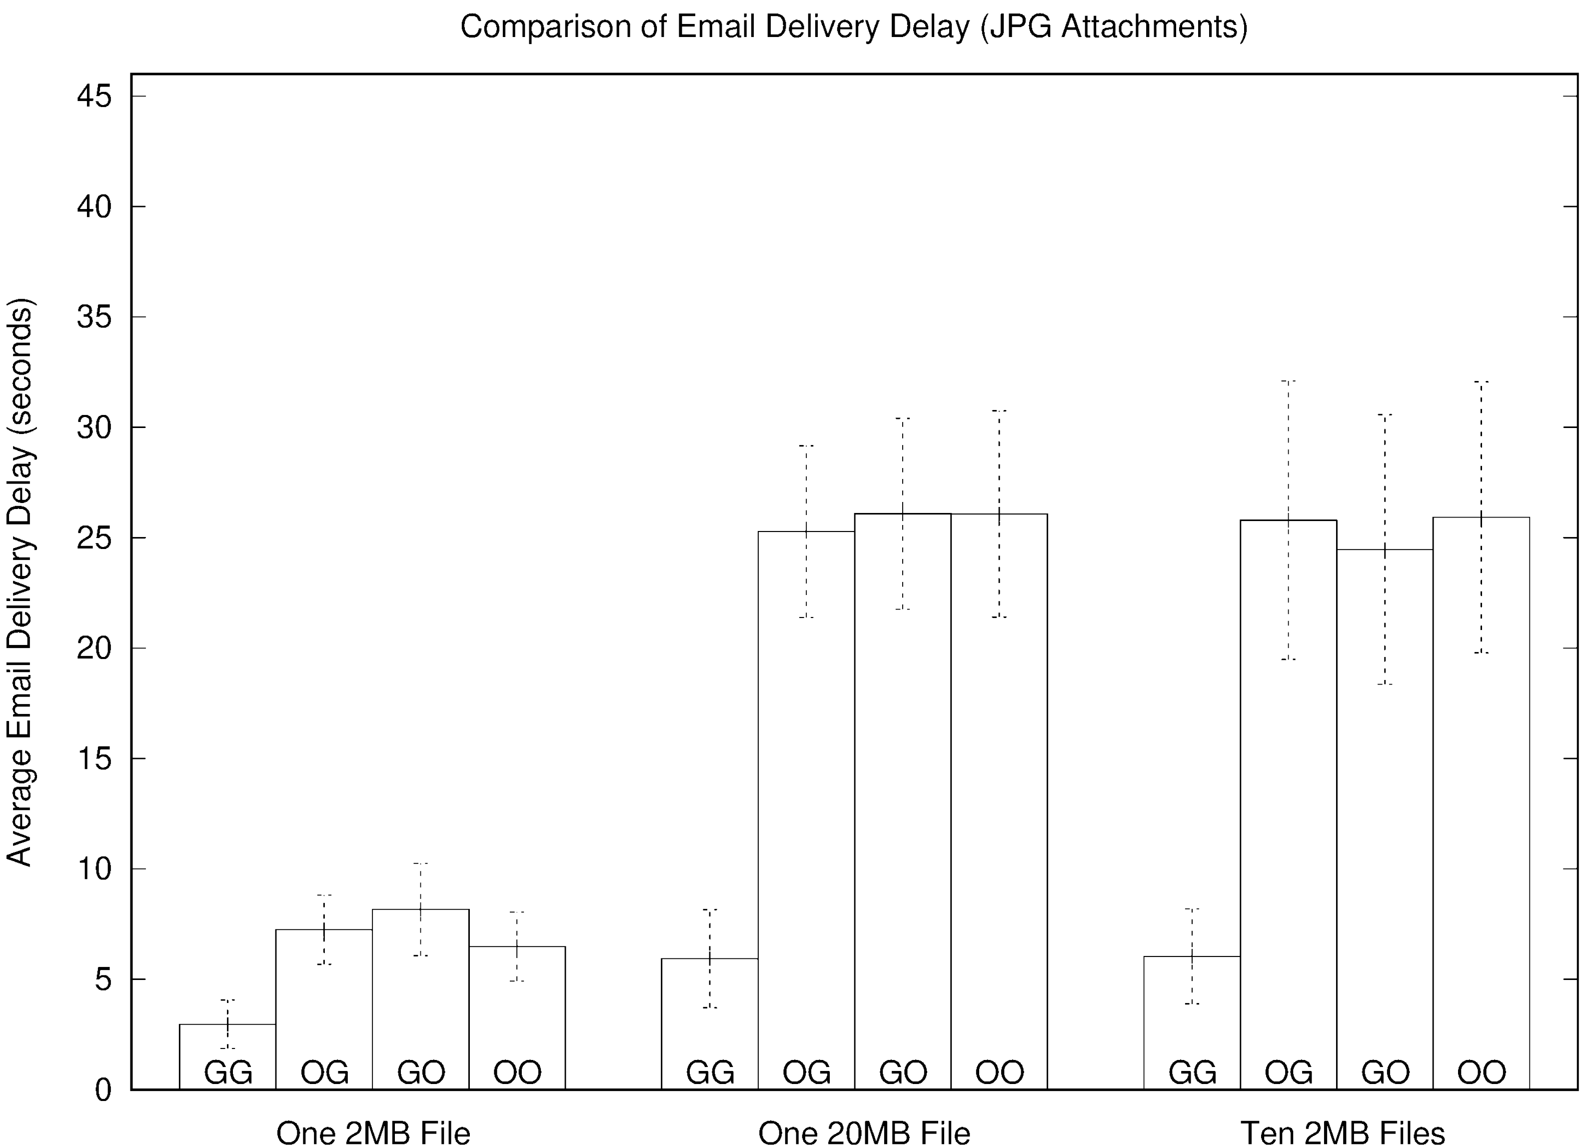

Email Delivery Delay with MP4 Attachments

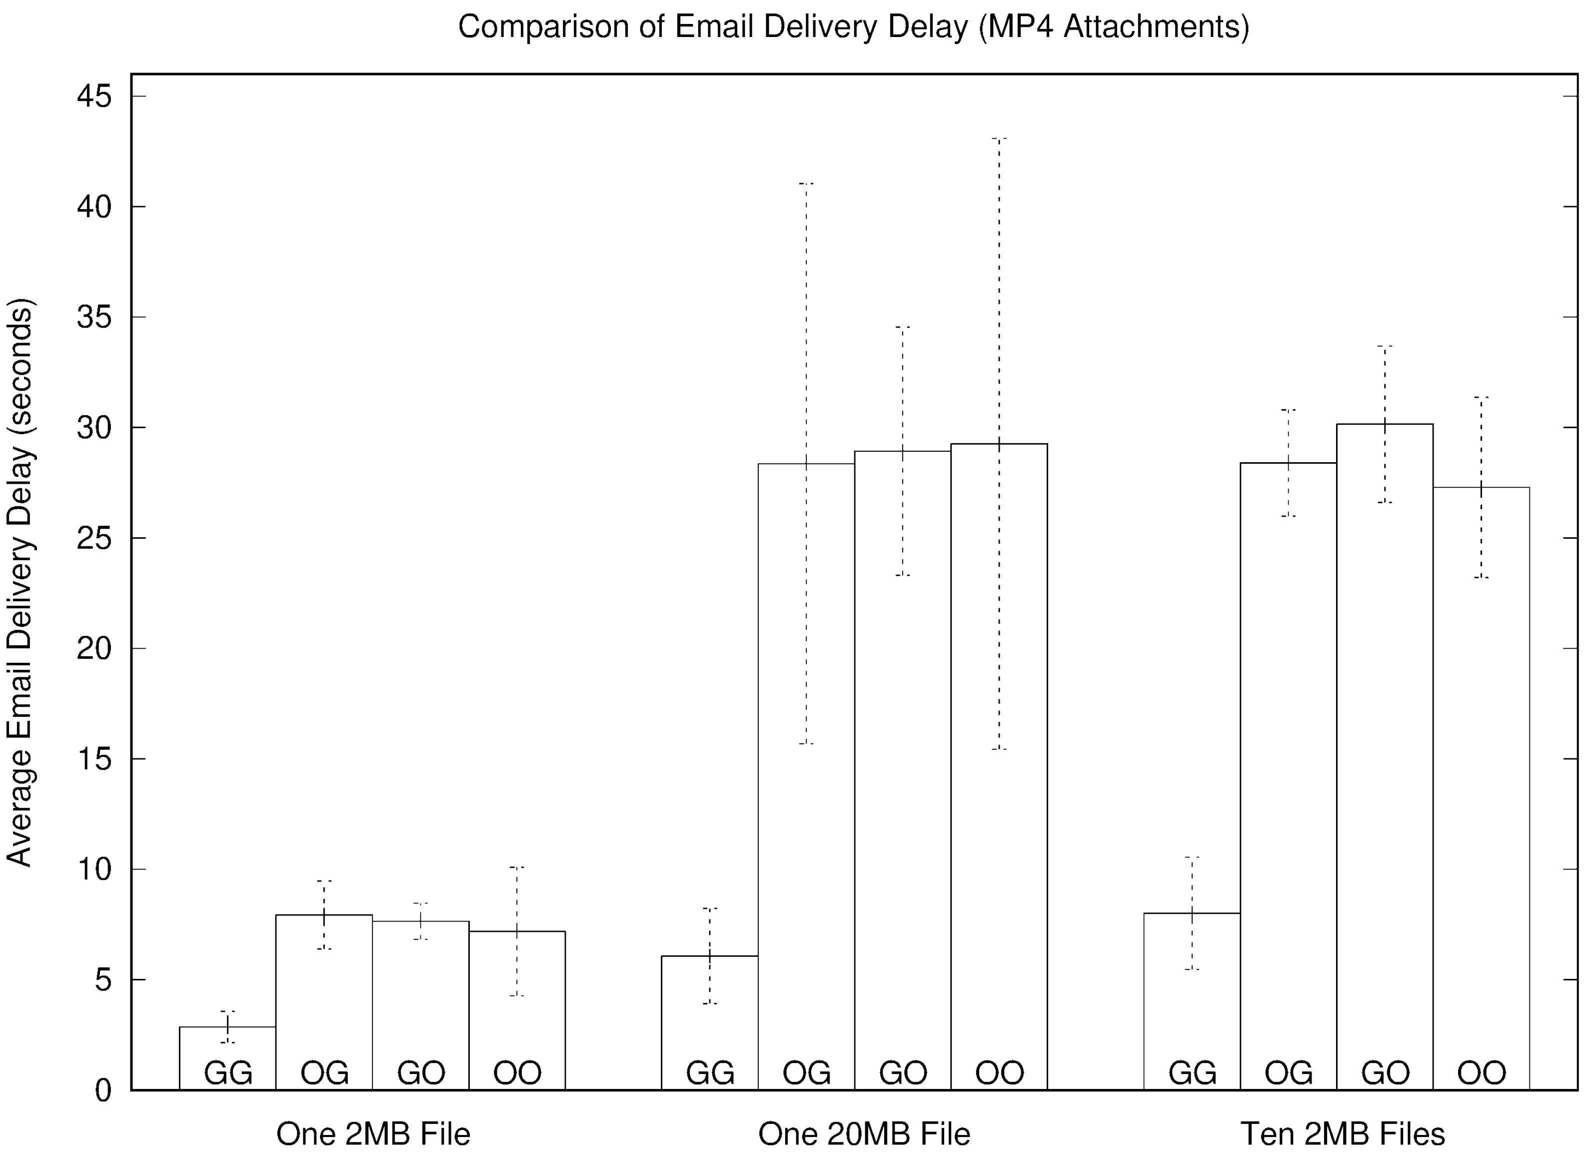

Email Delivery Delay with DAT Attachments

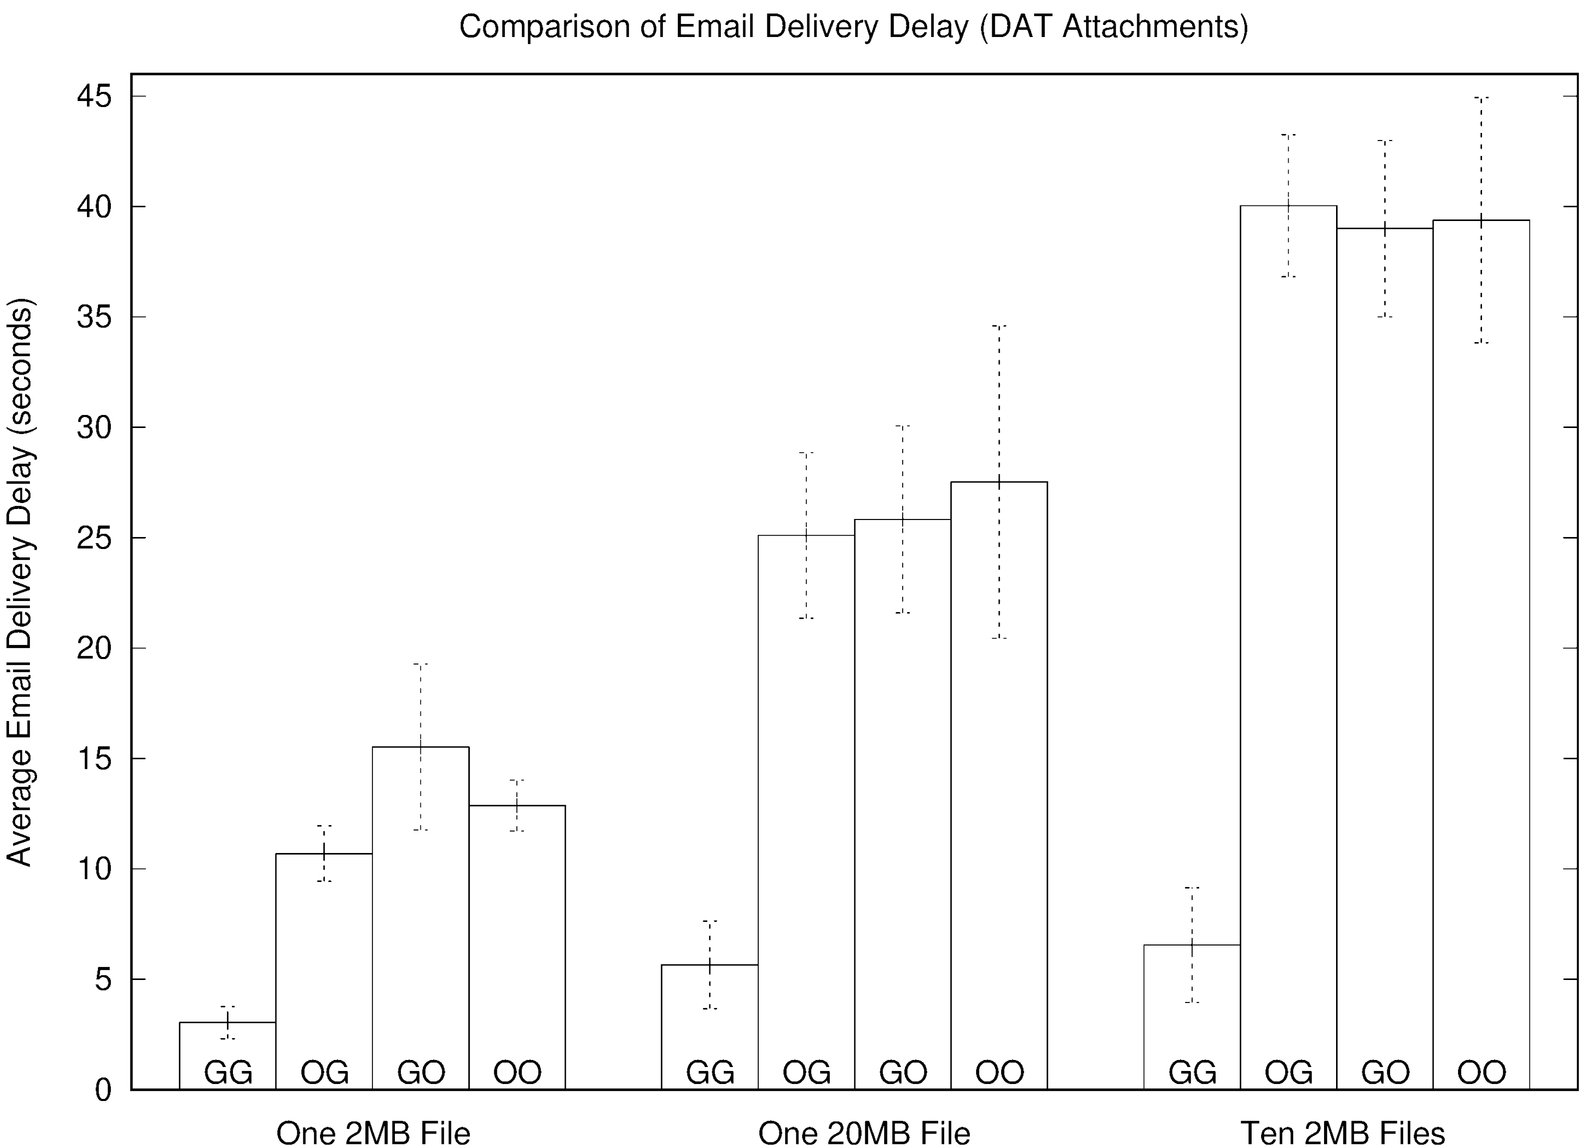

Lessons Learned

- Email ecosystem is extremely complex
- Outlook is not to blame for all email woes! ☺
- Differences in cloud-based email service providers
- Delays vary with size of attachments (obvious)
- Delays vary with type of attachments (less obvious)

Return

UNIVERSITY OF CALGARY

DNS: An Introduction

- Domain Name System (DNS) is one of the most prominently used infrastructure on the Internet.
- It provides a mapping between domain names and IP addresses.

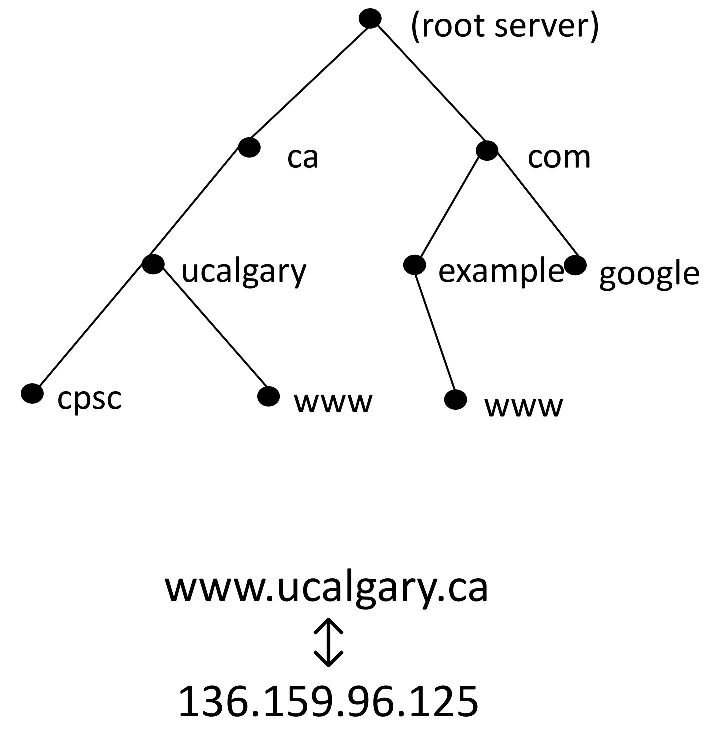

- Motivation:
  - Gain a better understanding of DNS traffic within our campus network.
  - Address potential performance issues and security vulnerabilities.

U of C DNS Traffic Characterization

- Outbound DNS Traffic
  - U of C is a client/user of DNS.
  - Four registered DNS resolvers: *OutCampusOne*, *OutCampusTwo*, *OutCPSC*, and *OutAkamai*.

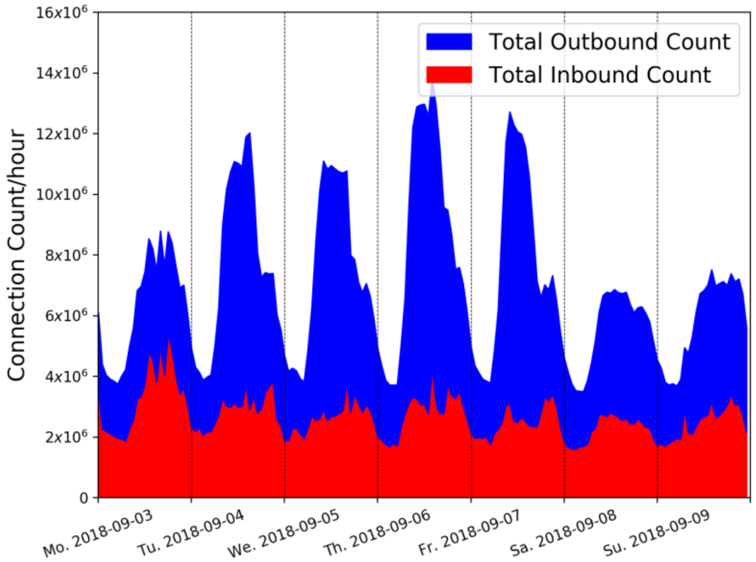

- Inbound DNS Traffic
  - U of C is a DNS service provider.
  - Five registered DNS servers: *InCampus*, *InCampusNew*, *InCPSC*, *InPHAS*, and *InAkamai*.

Type Analysis (Outbound)

- Query Types
  - A/AAAA/PTR/NS
  - Name-to-IP queries dominate.
  - Several *Other* types.

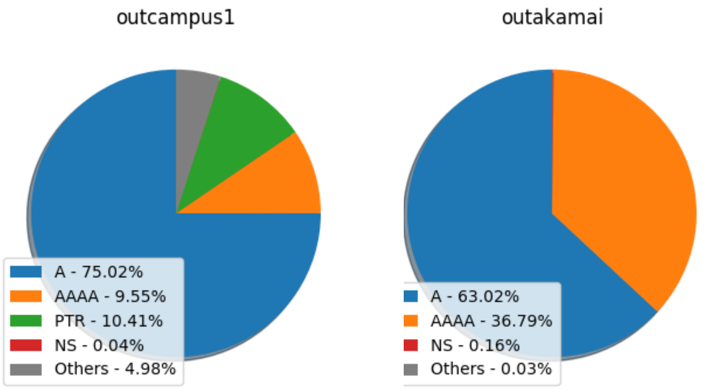

- Response Types
  - NoError/NXDomain/ServFail/Refused
  - Most are answered without error.
  - Others: NoError or NoResponse.

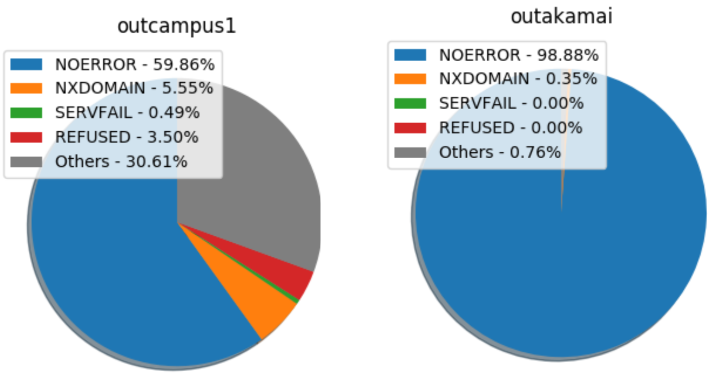

Type Analysis (Inbound)

- Query Types
  - A/AAAA/PTR/NS
  - PTR queries dominate.
  - Many *Others* types.

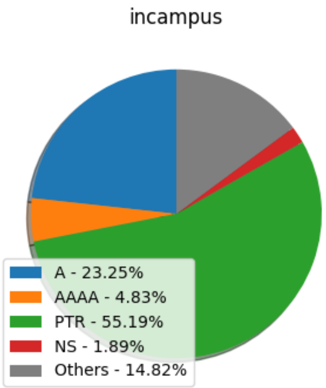

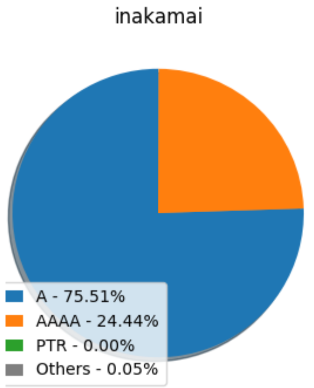

- Response Types
  - NoError/NXDomain/ServFail/Refuse
  - Most queries are answered.
  - Error rate is very high.

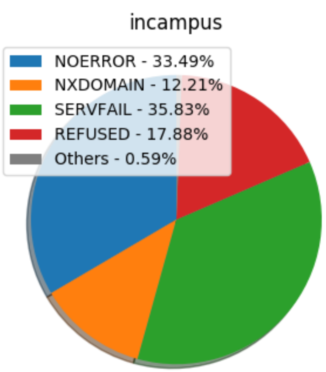

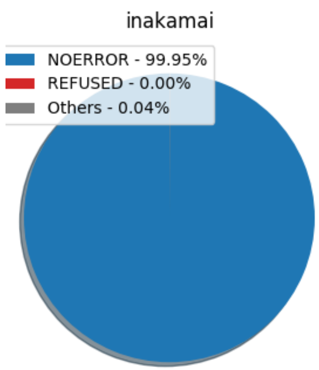

Inbound DNS Scanning

- DNS-based scans are prevalent within inbound traffic.
- More than 4 million scan connections in one week.
- Scanning for DNS service and recursive DNS support.
- All the open resolvers on campus are detected.

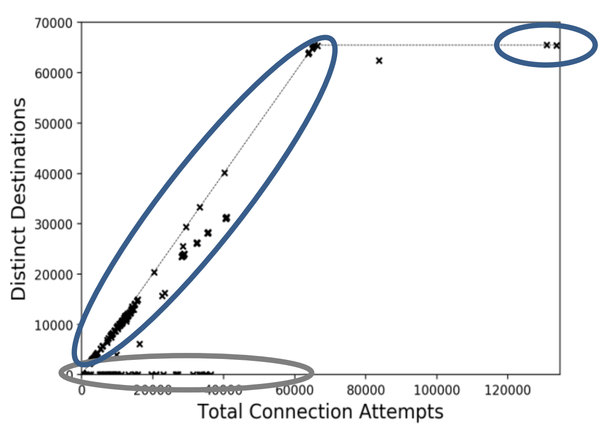

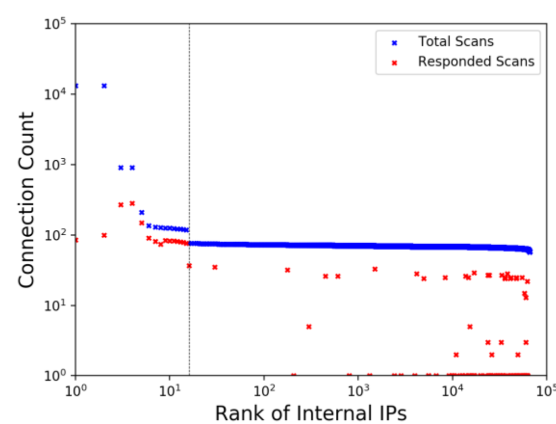

UNIVERSITY OF CALGARY

Lessons Learned

- Outbound traffic seems “normal”, but inbound traffic seems quite “abnormal”.
- Short TTLs or misconfigurations can generate a large number of extraneous DNS requests.
- Reverse DNS queries are prevalent in inbound traffic.
- Four main types of DNS anomalies are observed.
- Efficiency of DNS service can be improved.

- Future Work
  - Longer measurement period.
  - Deeper investigation of misconfigurations.

Return

- NAT: Network Address Translation

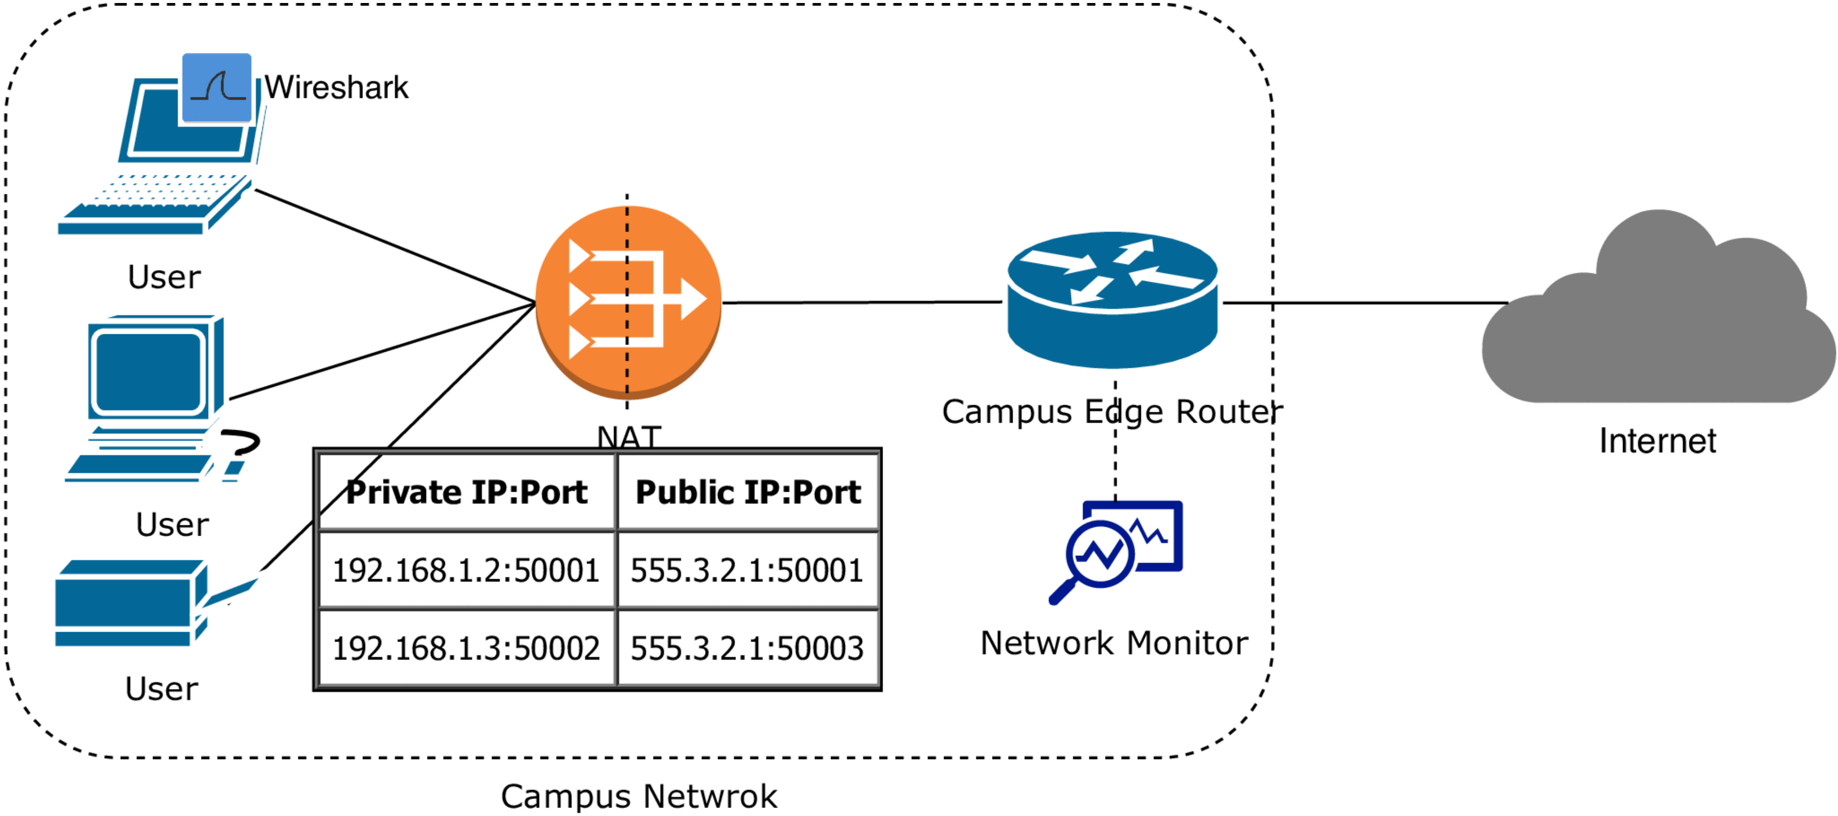

Private IP:Port	Public IP:Port
192.168.1.2:50001	555.3.2.1:50001
192.168.1.3:50002	555.3.2.1:50003

- Unusually many rejected TCP connections (REJ)

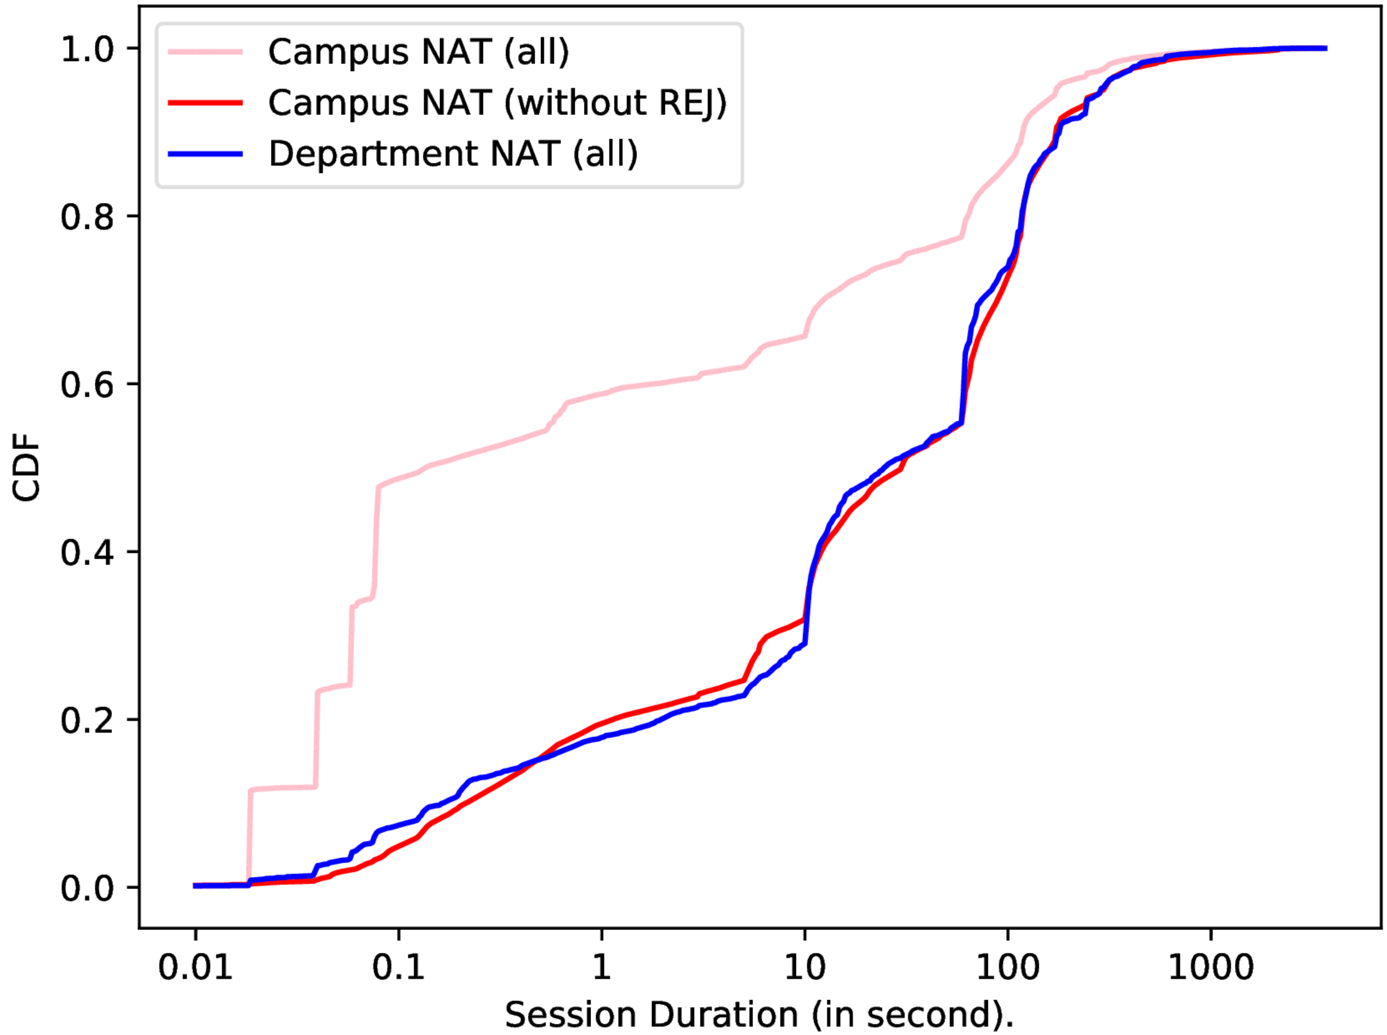

- Example of CPSC department-level NAT (30 minutes)

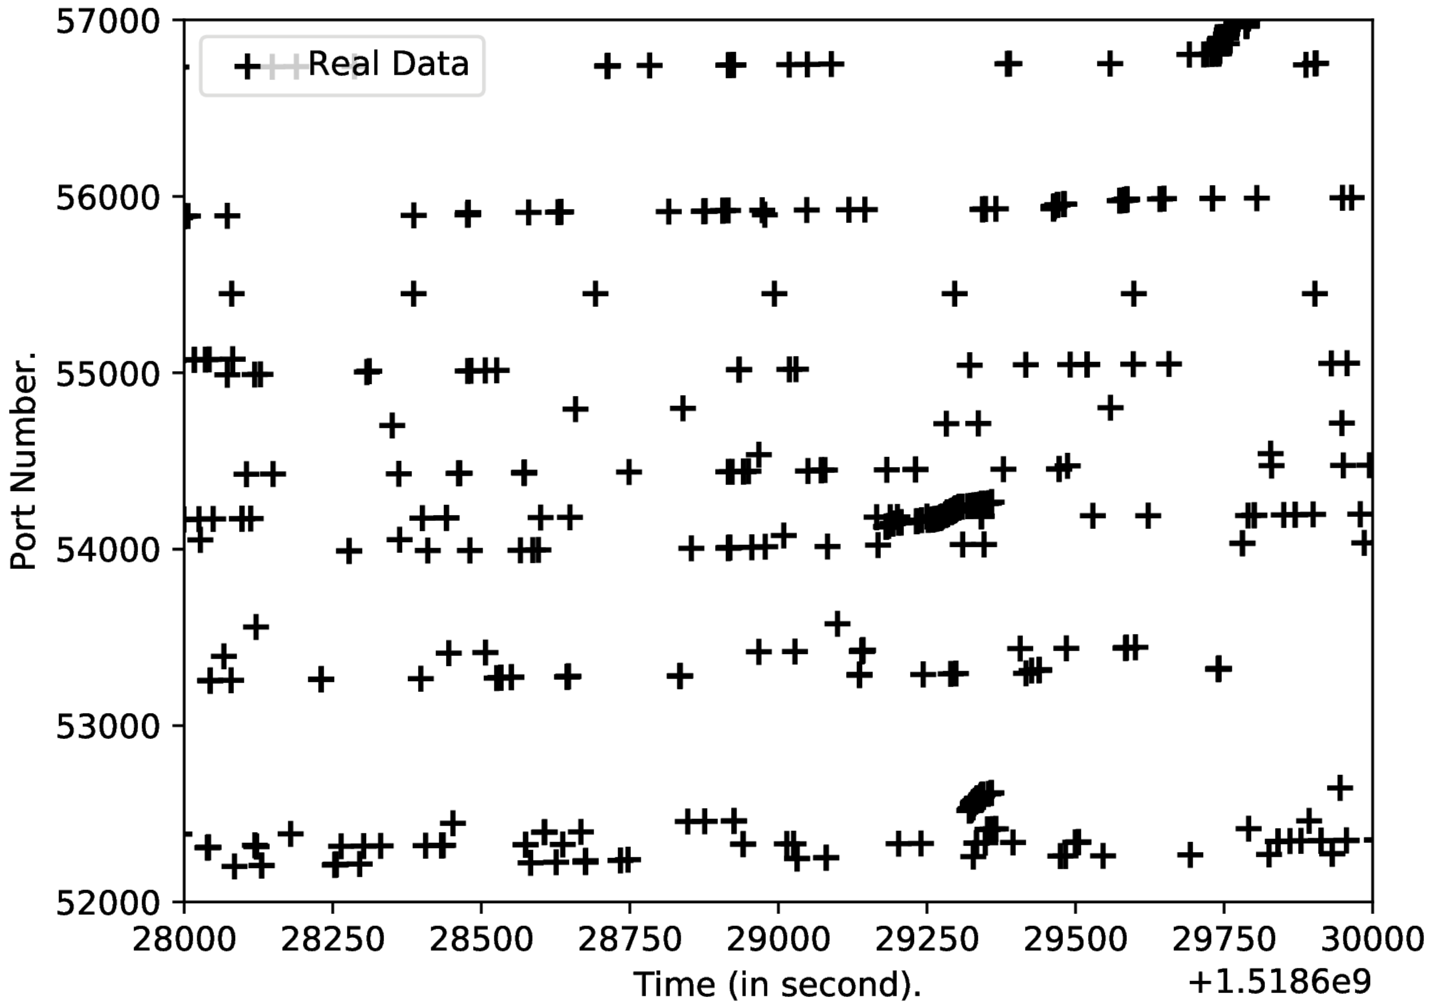

- Example of university-level NAT (6 minutes)

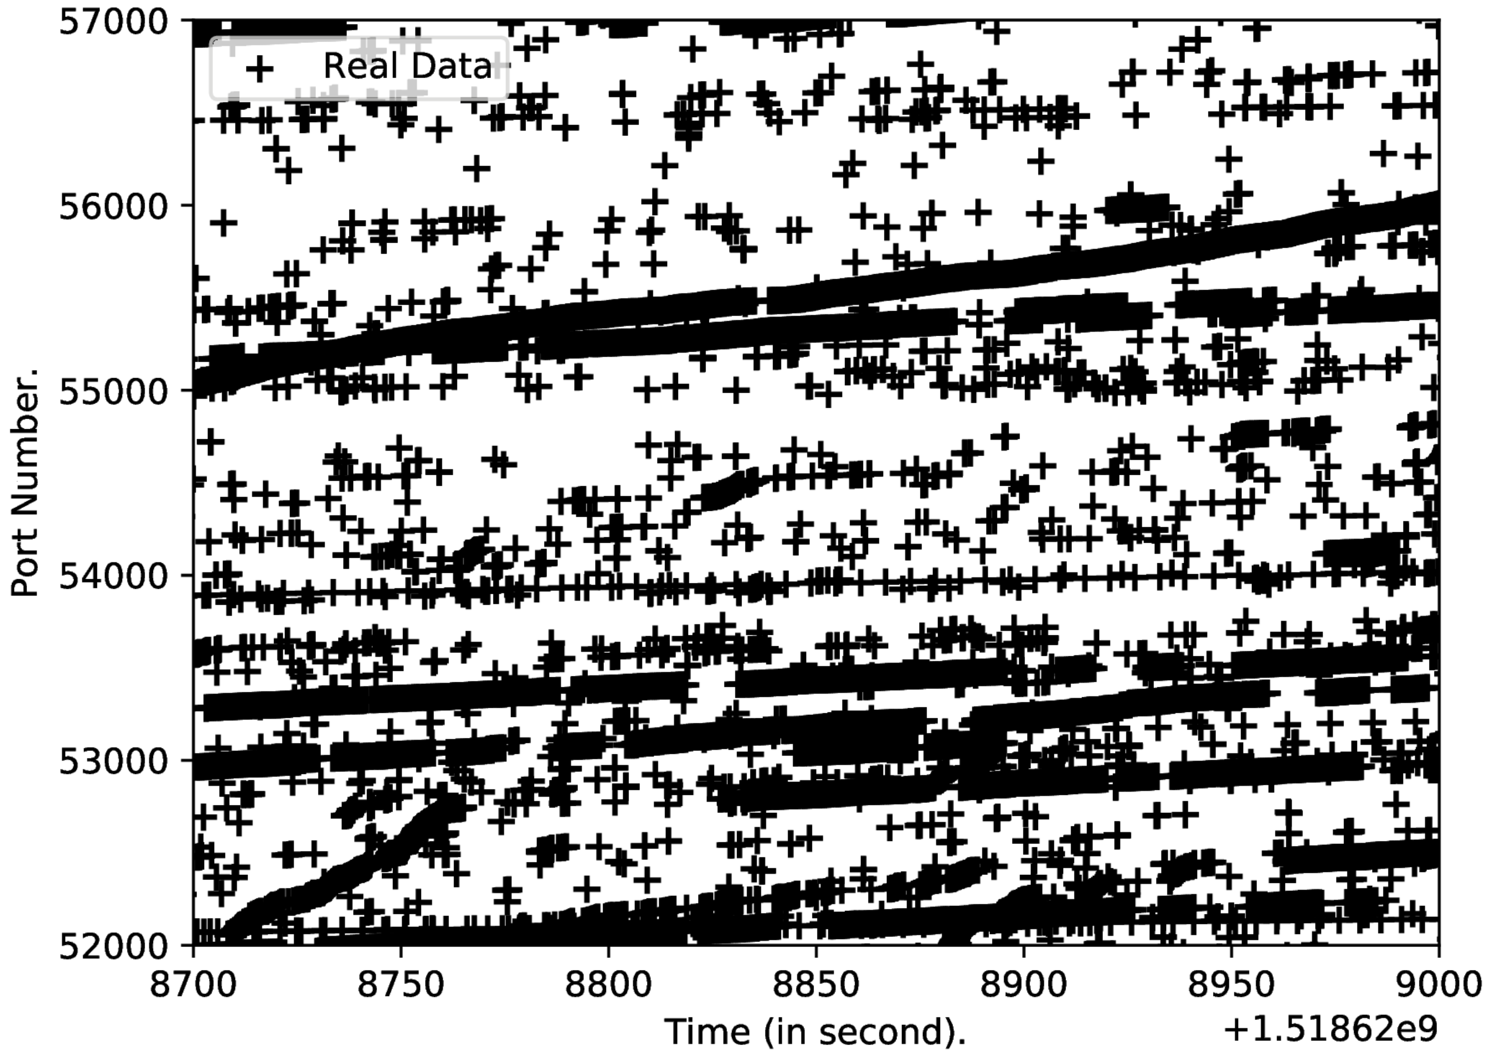

- Unusual reject (REJ) behaviours in Outlook traffic

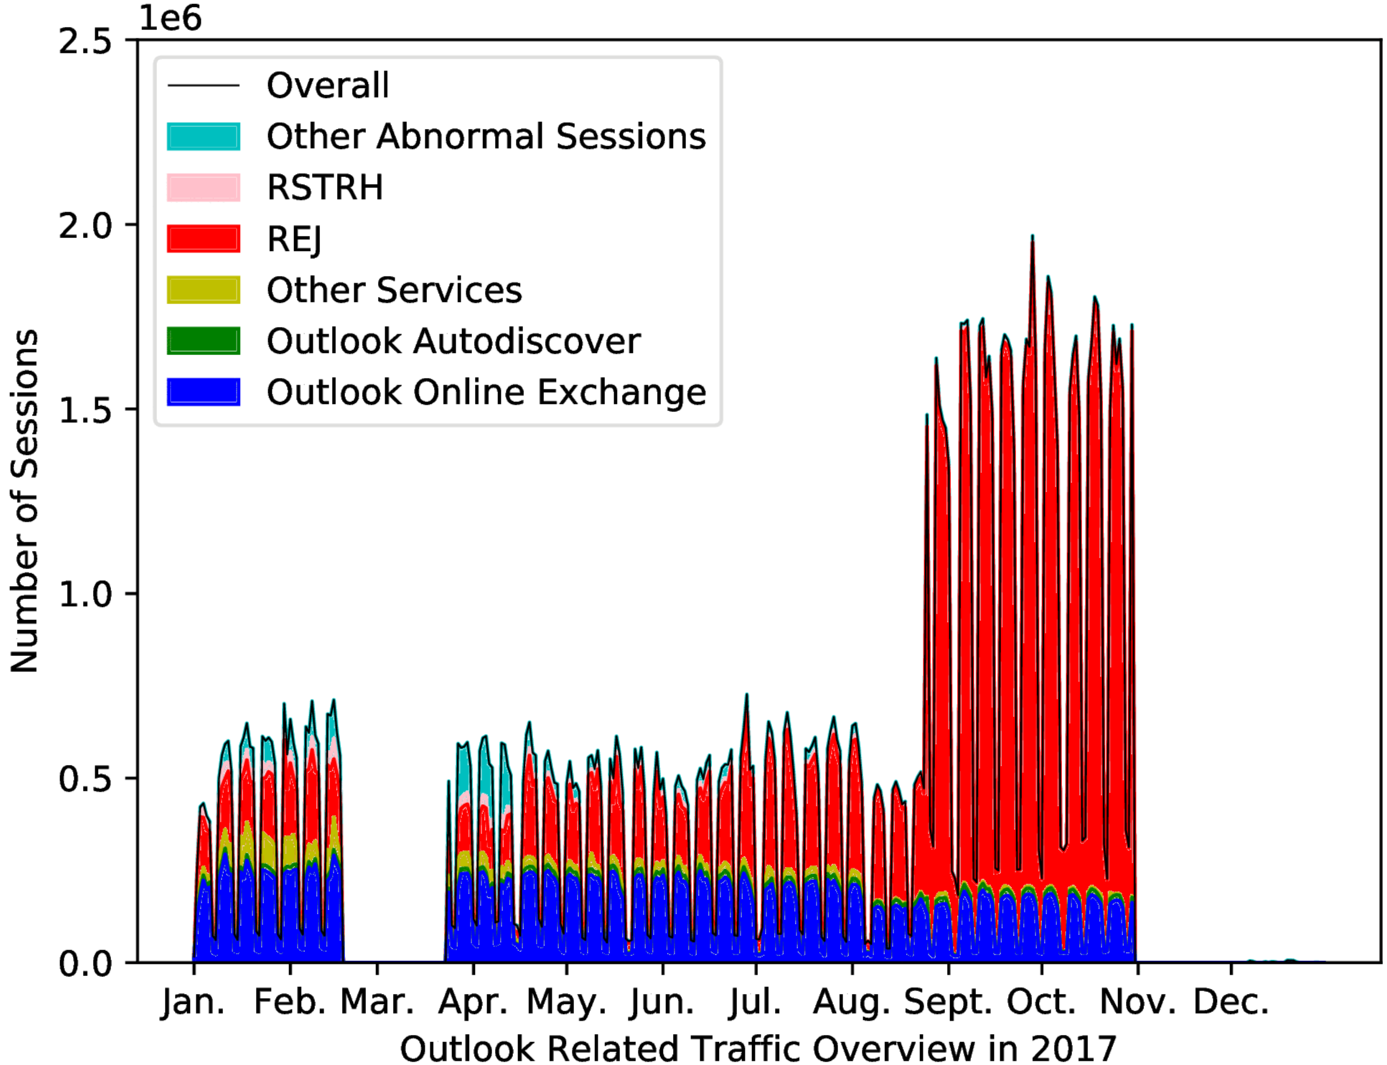

Lessons Learned

- Something definitely wacky in our NAT traffic
- Root cause unknown
- Prime suspect: misconfigured Outlook server
- Still trying to sort this one out with UCIT

Return