Zoomiversity: A Case Study of Pandemic Effects on Post-Secondary Teaching and Learning

Mehdi Karamollahi, Carey Williamson, and Martin Arlitt

University of Calgary

Abstract. The first wave of the COVID-19 pandemic hit North America in March 2020, disrupting personal and professional lives, and leading to work-from-home mandates in many jurisdictions. In this paper, we examine two years of empirical network traffic measurement data from the University of Calgary's campus network to study the effects of the pandemic on a post-secondary education environment. Our study focuses on the online meeting applications and services used, as well as traffic volumes, directionality, and diurnal patterns, as observed from our campus edge network. The main highlights from our study include: changes to inbound and outbound traffic volumes; reduced traffic asymmetry; significant growth in Zoom, Microsoft Teams, and VPN traffic; structural changes in workday traffic patterns; and a more global distribution of campus network users.

Keywords: Network traffic measurement · Workload characterization.

1 Introduction

The COVID-19 pandemic affected everyone's daily life, both personally and professionally. Lockdowns, travel restrictions, and stay-at-home orders were in effect in most parts of the world in 2020, and they had many consequences on people individually and socially. The way that people work and study changed quite drastically, with many people relying much more extensively upon the Internet and online tools for their daily tasks [10].

In a broader sense, the pandemic has provided a glimpse into the possible Future of Work (FoW) [3, 24, 31], a term used to describe a flexible work-from-home society enabled by digital connectivity, telepresence, and computer networks. With the continuous move to the cloud infrastructures within the organizations and universities around the world [11], most people were aware of the possibility of working remotely before the pandemic occurred, but relatively few have done it. At the very least, the pandemic accelerated the transition to FoW and made it real for many more people, changing some mindsets and possibly influencing remote work and online learning technologies for the future.

In this paper, we study the effects of the pandemic within the context of a post-secondary education environment. We do so from a network-level viewpoint, by studying the changes in the Internet traffic patterns into and out of our campus network. Specifically, we examine two years of empirical connection-level network traffic data to identify changes in the volume, timing, and directionality of traffic, as well as the application mix. Doing so offers insights into how the pandemic affected the work and study habits of our campus community.

A main focus in our study is on the use of Zoom video-conferencing software, which was adopted by University of Calgary (UCalgary) as the preferred solution for remote teaching and learning. Zoom has been adopted by many universities, companies, and other organizations for remote communication purposes.

Zoom is a popular and easy-to-use video-conferencing solution. Zoom offers a free account with some limitations, such as a maximum meeting duration of 40 minutes. However, many organizations (including UCalgary) purchased the corporate license for Zoom so that their members could use it for teaching and learning, as well as meetings and conferences, without the duration limit.

Zoom, of course, is not the only video-conferencing solution on the market. Microsoft Teams and Google Meet are two other online conferencing applications used by our campus community for meeting purposes. Some features are free for the public to use, and the rest are accessible only to licensed organizations. Other popular solutions include FaceTime, Skype, Vidyo, and Webex.

In our work, we study the network traffic of three online meeting applications on our university campus network. We focus on characterizing the network traffic from these applications, as seen on our campus. As a baseline, we provide pre-pandemic traffic measurements from 2019 and the early months of 2020, and compare 2020 traffic with this baseline.

Our main objectives are to answer the following questions:

- How has the campus network traffic changed during the pandemic, and why?
- What are the usage patterns for Zoom as the most prominent online video conferencing application on our campus?
- What other network applications and services are used to support remote work and learning?
- What are the potential implications of these changes on the future usage of our campus network?

The main contributions of this paper are as follows:

- We compare empirical network traffic data from 2019 (pre-pandemic) and 2020 (pandemic) to identify structural changes in traffic patterns.
- We identify the emergence of Zoom and Teams as popular applications for teaching and for meetings, respectively, and characterize Zoom usage.
- We identify temporal and geo-spatial changes in how our research and education community accesses and uses campus network resources.

The results from our work should be of value not only to networking researchers, but also to educators, academic administrators, and IT professionals. Using longitudinal data analysis, we provide several key insights on the growth and evolution of network traffic for online learning, and the performance implications of such traffic on a campus edge network.

The rest of this paper is organized as follows. Section 2 discusses prior related work on network traffic characterization and the effects of the COVID-19 pandemic. Section 3 describes the methodology for our study, focusing on our network environment, our network traffic measurement infrastructure, and our data analysis tools. Section 4 presents the main high-level results from our study, while Section 5 provides detailed results regarding Zoom traffic. Finally, Section 6 concludes the paper.

2 Related Work

Researchers in academia and in industry rely on network traffic measurement as an increasingly important methodology to obtain data, analyze Internet traffic, assess network performance, identify network security issues, and investigate different features of new protocols and applications. The book by Crovella and Krishnamurthy [8] provides the technical underpinnings of this discipline.

The usage of network traffic measurement and workload characterization techniques is broad and extensive. Classic examples include the characterization of wide-area TCP connections [26], Web traffic [2, 4], and email traffic [29]. More recent works have studied video streaming services [1, 12], as well as the growth and evolution of online social networks [15, 23, 32]. Such studies offer insights into the changing nature of Internet traffic, and its potential effects on network performance. We follow a similar approach in our work.

The COVID-19 pandemic has affected Internet usage dramatically. Since the onset of the pandemic, several research works have noted changes in the timing, volume, and directionality of traffic, as many people switched to work-from-home scenarios. One of the first was the weekly blog by Labovitz [16], analyzing data from several networks in Western Europe. As of March 9, 2020, this report noted 20-40% increases in traffic during the evening peak hours, 3x growth in teleconferencing apps (e.g., Skype, Zoom), and 4x growth in gaming traffic [16]. A later report [17] indicated that aggregate traffic was up by over 25%, and that the normalized peak traffic was 25-30% above pre-pandemic levels. Also, DDoS attacks increased by 40-50% after the pandemic [17]. Similar observations arise in our work, along with insights that are specific to Zoom traffic.

The Broadband Internet Technical Advisory Group (BITAG) produced a detailed report on how Internet traffic changed, and how network operators and providers managed the unprecedented circumstances [14]. Though focused on the US, this report provides valuable insights into network operations from many different vantage points, including core, edge, and ISP locations. The report states that the Internet, in general, was robust during the pandemic, and continued to perform well. Several performance issues experienced by some users were attributed to end-user system configurations and outdated wireless equipment. Dramatic growth in VPN usage by campus networks is also reported in this document, along with the notable asymmetry in traffic growth between upstream and downstream. The busy hours for the downstream were in the evening with 12-25% growth, while the upstream peak hours start in the morning and run most of the day until about midnight. Our results also confirm VPN traffic growth, and noticable shifts in network usage patterns.

Feldmann *et al.* [10] provided a multi-perspective look at pandemic effects on Internet traffic, using datasets from ISPs, IXPs, and mobile network operators. The main highlights were shifts of 15-20% in Internet traffic within a week of lockdown. Their paper noted the emergence of non-hypergiants among the contributors to traffic growth, and identified a plethora of network applications being used in work-from-home environments. Our work confirms Zoom as a new potential hypergiant.

Lutu *et al.* [20] presented an analysis of the changes in user mobility patterns and how this affected the cellular traffic of a UK mobile network operator. They observed an overall decrease in mobility (i.e., roaming) by 50%, with non-uniform geographical changes. They reported a 150% increase in voice traffic, a 20% overall decrease in download traffic, and a 10% increase in uplink traffic. Nonetheless, the network operator was able to maintain service quality standards. Our work does not address cellular traffic at all, but we do see reduced WiFi usage from having fewer people on campus.

Liu *et al.* [19] studied how several US providers responded to changes in Internet traffic demands during the pandemic. They also identified some differences between rural versus urban users, which can affect QoS/QoE for online learning applications. The shift to using online meeting applications and platforms for learning and collaboration is also well documented in the literature [27, 28]. Our work indicates potential performance problems when Zoom is used for teaching and learning on a large campus edge network.

The closest study to our own so far is by Favale *et al.* [9], who studied traffic on the campus network of the Politecnico di Torino (PoliTO). They analyzed the changes in traffic patterns due to the restrictions in place in Turin, Italy, and the switch to online learning solutions. They observed that the campus inbound traffic drastically decreased, since fewer students were on campus, while outbound traffic more than doubled, due to the remote learning platform installed at the campus to support all online classroom instruction. Furthermore, they provided insight into the growth of online collaboration platforms, VPN, and remote desktop services. Compared to their work, our research spans an entire calendar year of pandemic-related network traffic data (2020), rather than just a few months, with the previous calendar year (2019) as a baseline. Furthermore, our campus uses the widely-adopted Zoom platform for remote teaching and learning, rather than a custom in-house solution. We provide observations on how the usage of online learning and meeting applications has changed in terms of connection counts and traffic volume, and offer insights into these changes.

A recent paper [5] studied three major videoconferencing systems: Zoom, Webex, and Google Meet. They used a cloud-based emulated framework to generate videoconferencing sessions on these applications and then measure, study, and compare them. They measured streaming delay (lag), as well as a range of well-known objective QoE metrics, including PSNR (Peak Signal-to-Noise Ratio), SSIM (Structural Similarity Index Measure), and VIFp (Pixel Visual Information Fidelity). They found that these systems vary in terms of geographic location, resulting in different QoE. For example, Webex sessions created in US-west are subject to artificial detours via relays in US-east, inflating their lags. They saw that high-motion video feeds experience non-negligible QoE degradation on all three systems compared to low-motion video streaming. Finally, systems react differently under bandwidth constraints.

Another recent study [21] compared Zoom, Microsoft Teams, and Google Meet in an experimental testbed. They tried to find the baseline level of Internet performance needed to support common videoconferencing applications for remote learning. Under simulated conditions, they measured the bandwidth utilization, time to recovery from interruptions, and fairness under competitive circumstances.

Our work differs from these papers too, as we study empirically-captured network traffic data from thousands of users on our campus network. Our companion paper [6] developed tools to analyze Zoom sessions and meetings from these captured data and provides a microscopic view of Zoom traffic. This paper, on the other hand, provides a longitudinal (macroscopic) view of Zoom, Teams, Meet, VPN, and other applications involved in remote working and learning during the pandemic.

3 Data Collection and Methodology

This section discusses the methods and technologies used for our network traffic characterization study, as well as some of the limitations of our approach. We start with some brief contextual information about the university environment that we studied.

3.1 University Environment

UCalgary is a medium-sized university with about 30,000 students. The academic schedule is semester-based, with the Fall (September to December) and Winter (January to April) semesters each having a full set of course offerings. There are also Spring (May to June) and Summer (July to August) semesters, each with reduced course offerings.

During the COVID-19 pandemic, the switch from in-person learning to remote online learning took place quite abruptly on March 13th, 2020, during the Winter semester. Online learning remained the norm throughout the rest of the calendar year, though a small number of students (20%) were allowed back on campus in Fall 2020, mainly in capstone and/or lab-based experiential learning courses with limited enrollments.

Videoconferencing. With the shift to remote learning, the students, staff, and faculty members started to use online meetings and screen-sharing applications to continue with the courses, academic tasks, and regular or occasional meetings. The University officially advised its community to use Zoom for teaching and learning, and it has been the dominant way of teaching classes since the lockdown. Microsoft Teams is offered for internal or external meetings. An organizational license was purchased for Zoom, and Microsoft Teams is an integrated application within the Office365 suite available via a campus-wide licence.

Remote Access. Being physically away from campus raised access issues for almost everyone. For example, many faculty, staff, and graduate students needed to access computers in their offices or labs to proceed with their work or research. Even undergraduate students using systems in different labs before the lockdown needed to connect remotely to those systems. Furthermore, certain services require access from a university IP address, augmenting the demand to connect to the campus network. Three different remote access solutions were offered to resolve these issues: Secure Shell (SSH), Virtual Private Network (VPN), and Remote Desktop Protocol (RDP).

Authentication. Our campus network uses an authentication service that checks user credentials before accessing enterprise resources, such as the wireless network, learning management system, email, and Office365 applications.

3.2 Passive Measurement

Passive measurement involves capturing ambient network traffic and analyzing it either online or offline. With this technique, no additional traffic is produced, and the ordinary network traffic is not altered in any way. We collected two years of empirical network traffic data using this approach.

Our monitor uses an Endace DAG (Data Acquisition and Generation) packet capture card. The monitor is installed in the main data center on campus, and receives from the edge routers a mirrored copy of every packet entering or leaving the campus network. Those packets are then sent to a Zeek (formerly known as Bro) worker node [25]. For privacy purposes and to reduce storage requirements, Zeek aggregates all the packets of the same connection and stores a summarized entry for that connection. This summary consists of many fields, including a unique identifier of the connection, the connection's 4-tuple of endpoint addresses/ports, the time of the first packet, duration of the connection, and the number of packets and bytes sent by both the originator and responder.

We use ARC (Advanced Research Computing), an existing HPC (High Performance Computing) cluster at UCalgary, for storage, management, and script-based processing of our traffic data. We also use Vertica, an SQL-based big data analytics platform, to analyze the captured data. Using Vertica is fast and convenient for network traffic analyses, since it supports parallel execution of SQL queries on structured data [18].

3.3 Active Measurement

Active measurement refers to establishing connections and sending data packets to identify entities in the network, characterize traffic, or measure different metrics. In this study, we used active measurement techniques judiciously to identify hosts and servers associated with organizations and autonomous systems under study, and their traffic attributes, such as port numbers. This information is most often essential in network traffic measurement and workload characterization studies. For this purpose, we mainly conducted simple experiments using basic network tools like `nslookup` and `traceroute` and used Wireshark to capture packet-level traffic. We then analyzed the captured logs and extracted the required fields, such as IP addresses associated with the target organizations and the port numbers used by applications. This information may also be utilized in the passive measurement when required.

3.4 Challenges and Limitations

As with any network traffic measurement study, there are challenges and limitations that affect the completeness of our data, and hence the interpretation of results. We discuss these issues here.

First and foremost, it is important to note that our monitoring infrastructure is set up to observe packet traffic that is strictly *between* the university and the Internet. Specifically, the monitor does not see traffic that stays completely within the campus network (e.g., a student in residence connecting to an internal server), nor traffic that is completely external (e.g., a home residential user directly accessing Netflix). The pandemic has thus changed the visibility into Internet usage by our campus community. Some traffic that was not visible previously (e.g., accessing a university Web server while at work) is now visible when people work and learn from home. Conversely, some traffic that was visible previously (e.g., YouTube accessed from the campus WiFi network) is no longer visible when these users directly access the Internet from home. For VPN, however, remote users actually obtain a campus IP address from the BYOD subnet, which is then used to connect to the Internet. Therefore, a connection to the campus VPN contributes to both incoming and outgoing connection counts as seen by the monitor.

A second challenge, as in any longitudinal traffic study, arises from unexpected events that disrupt data collection. Several such incidents occurred during the 2020 year under study. The most pernicious of these were aggressive scanning attacks (horizontal and vertical) that exhausted the memory resources on our monitor, and crashed the system. These outages in data collection are visible in several of the time-series graphs presented in the paper.

To mitigate the foregoing problem, we disabled the scanning module in Zeek, and reconfigured our monitor to do a software restart every 3 hours. While this strategy avoids crashes that lose substantial amounts of data, it does limit visibility into long-duration connections. We subsequently experimented with shorter (1 hour) and longer (6 hours) restart intervals as well, prior to settling on 6-hour intervals since July 2020. The effects of these configuration changes are also apparent in several of our traffic plots.

Another challenge regarding the videoconferencing applications is that (unlike the on-site proprietary solutions such as the case for Favale *et al.* [9]) we have limited information available about their infrastructure and how the applications behave. In many cases, we had to reverse engineer their behavior based on a few documents. Furthermore, their deployments may have been changed during the pandemic. However, due to the wide adoption and availability of these applications across the globe, our analysis and results should be generalizable to other environments with similar contexts.

Despite these issues, we still believe that our empirical dataset offers great research value. Where appropriate, we exercise caution in our interpretations of results, and contextualize them accordingly.

3.5 Ethical Considerations

Permission to capture network traffic data was authorized via the ethics review process at UCalgary and was carried out with the cooperation of the IT center. Our network monitor is mounted in a secure data center with restricted physical access. A limited amount of traffic data is stored on the monitor at a time, with data summarization and transfer to a secure storage server happening on a daily basis. All data is stored in logs as per-connection summaries. Data analysis is done at an aggregate level, and not individually. Furthermore, most users get transient IP addresses from DHCP and/or NAT when connected to the campus network. Any identification process in the active or passive measurement is limited to the hosts and servers associated with organizations and applications under study, not individual users. Access to the log data is restricted to those specifically authorized to conduct networking and security research. Any security-related vulnerabilities (e.g., compromised machines, amplification attacks) detected in these summarized data are reported to the campus IT team for subsequent follow-up.

4 Measurement Results

This section presents the results from our empirical network traffic study. We start with an overview of the traffic on a year-to-year basis, and then focus on specific applications and services, including authentication, learning management system, and VPN.

4.1 Traffic Overview

Figure 1[1] provides a high-level overview of our campus network traffic volume, in Terabytes (TB) of data per day for calendar years 2019 and 2020. The horizontal axis shows the time in months, while the vertical axis shows inbound data on the upper part of the plot, and outbound data on the lower (negative y axis)

[1] This figure uses the direct information from DAGstats and not the logs captured by Zeek. Therefore, it is not affected by the monitor restarts and the reconfiguration in mid-July. However, it is affected by the monitor crash in late March.

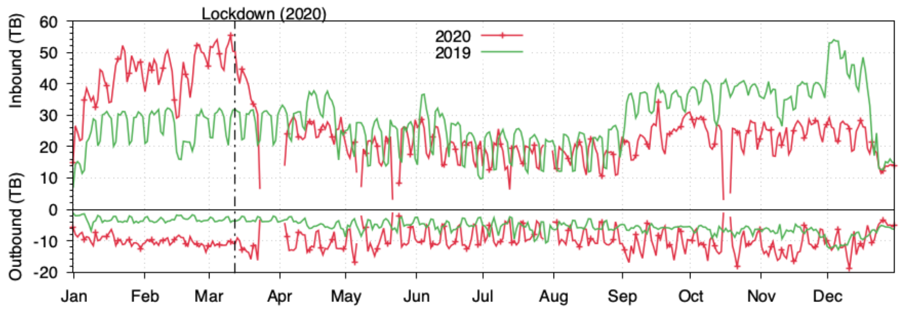

Fig. 1. Daily network traffic data volume in TB for 2019 (green) and 2020 (red). Upper axis is inbound traffic; lower axis is outbound traffic.

part of the plot. The green lines are for the baseline pre-pandemic year 2019, while the red lines are for the pandemic-affected year 2020.

There are several structural patterns evident in Figure 1. First, there is generally higher network activity during the main semesters (Jan-Apr and Sept-Dec) than during the Spring/Summer period (May-Aug). Second, there are distinctive weekly cycles. Third, our campus is a net consumer of data, with inbound traffic far exceeding outbound traffic. Fourth, there is a general decline in network traffic in late December when the university is closed for the holiday season, and few people are on campus.

There are also several pandemic-related effects evident in Figure 1. The most prominent of these is the sharp decline in traffic volumes in mid-March 2020, when classes were cancelled, people were asked to stay home, and remote learning began. Also notable is how the inbound traffic for Fall 2020 (Sept-Dec) is about 25% below that observed for Fall 2019. In over a decade of monitoring our campus network, this is the first time that we have observed a reduction in network traffic volume from one academic year to the next. Furthermore, this decline differs starkly from the Winter semester (Jan-Apr), in which the 2020 traffic prior to the lockdown exceeds that of 2019, for both inbound and outbound (with about 84% growth in overall traffic).

These dynamics in Fall 2020 reflect the fact that most people were still away from campus, working and learning from home. This observation is further supported by the increases in outbound traffic volume in Fall 2020 (almost 50% increase with respect to the prior year). As explained earlier, a connection to the campus VPN generates an incoming connection to the VPN server, as well as outgoing connections to the user's target hosts on the Internet. It results in the symmetry between outbound and inbound traffic volume due to a VPN connection. Therefore, VPN connections have no net effect on the overall asymmetry of the campus traffic observed.

The campus is still a net consumer of Internet traffic, and there are two main explanations for it. First, a large subset of campus services is being hosted in the cloud, such as Microsoft Office365 and Learning Management System (LMS).

Table 1. Top 10 External Organizations by Traffic Volume on Day2019 (2019-09-24)

Rank	Organization	Flows	% Flows	Bytes (GB)	% Bytes	Outbound	Inbound
1	Apple	11,172,676	6.15	5,417	12.91	791	4,627
2	Netflix	519,633	0.29	5,094	12.14	82	5,012
3	Akamai	16,907,100	9.30	4,815	11.48	131	4,683
4	Google	33,788,336	18.59	3,536	8.43	470	3,066
5	CANARIE	500,082	0.28	3,238	7.72	38	3,200
6	Facebook	7,505,585	4.13	2,891	6.89	130	2,761
7	Microsoft	37,201,566	20.46	2,034	4.85	935	1,098
8	Amazon	25,083,071	13.80	1,941	4.63	210	1,731
9	Fastly	2,934,594	1.61	1,386	3.30	45	1,341
10	UChicago	3,400	0.00	1,185	2.82	16	1,169

Table 2. Top 10 External Organizations by Traffic Volume on Day2020 (2020-09-23)

Rank	Organization	Flows	% Flows	Bytes (GB)	% Bytes	Outbound	Inbound
1	Amazon	12,936,245	14.82	3,259	11.70	928	2,331
2	Akamai	6,225,932	7.13	3,140	11.27	79	3,061
3	Apple	3,950,781	4.53	2,545	9.14	392	2,154
4	Netflix	421,738	0.48	2,393	8.59	89	2,304
5	Microsoft	20,200,909	23.15	2,286	8.20	1,027	1,259
6	Google	15,818,810	18.13	2,268	8.14	744	1,524
7	CANARIE	328,570	0.38	1,551	5.57	21	1,531
8	Facebook	1,548,066	1.77	1,094	3.93	56	1,038
9	Shaw	145,454	0.17	924	3.32	585	339
10	Oracle	37,193	0.04	853	3.06	241	612

Second, our findings show that video streaming and entertainment services are significant contributors to campus traffic, even in 2020. However, the increase in outbound traffic volume after the lockdown has reduced the degree of asymmetry.

Figure 1 also shows several distinct outages in monitor data collection (e.g., a week in late March, plus a few days in May, late July, and mid-October). These outages were due to intensive scanning attacks on the university network that crashed our monitor. These attacks were more frequent and more extreme during the pandemic than in the previous year. The main takeaway here is that *campus network traffic has changed in both expected (e.g., decline in inbound, increase in outbound) and unexpected ways (e.g., intensive scanning attacks).*

4.2 Structural Analysis

To better understand the changes in network traffic, we first examined the traffic volumes for hypergiants, such as Google and Microsoft. Table 1 shows the Top 10 external organizations based on total byte traffic volume (in Gigabytes) on a weekday in Fall 2019 (**Day2019**: 2019-09-24). As a representative day, this table illustrates the pre-pandemic traffic pattern for hypergiants. Apple tops the list at 5.4 TB/day, due to the multiple services it offers, such as iCloud and Apple TV.

Table 3. Top 10 Internal Subnets by Traffic Volume on Day2019 (2019-09-24)

Rank	Subnet	Flows	% Flows	Bytes (GB)	% Bytes	Outbound	Inbound
1	NAT 1	96,802,932	53.25	26,547	63.27	2,434	24,113
2	NAT 2	30,148,603	16.58	4,780	11.39	884	3,896
3	Guest WiFi	7,292,797	4.01	2,050	4.89	170	1,880
4	Other (4)	104,936	0.06	1,210	2.88	36	1,174
5	WLAN	388,178	0.21	399	0.95	28	371
6	Other (6)	4,107,762	2.26	381	0.91	62	320
7	RezNet 1	385,428	0.21	326	0.78	15	311
8	RezNet 2	417,543	0.23	319	0.76	18	300
9	RezNet 3	417,322	0.23	315	0.75	32	283
10	Other (10)	380,008	0.21	312	0.74	17	295

Table 4. Top 10 Internal Subnets by Traffic Volume on Day2020 (2020-09-23)

Rank	Subnet	Flows	% Flows	Bytes (GB)	% Bytes	Outbound	Inbound
1	NAT 1	23,247,481	24.38	12,824	46.03	2,105	10,719
2	NAT 2	27,246,242	28.58	6,285	22.56	1,280	5,005
3	Guest WiFi	2,627,619	2.76	1,208	4.34	97	1,111
4	VPN (217)	18,315	0.02	900	3.23	698	202
5	Other (6)	5,244,110	5.50	705	2.53	181	523
6	Admin (83)	133,204	0.14	366	1.31	8	357
7	Other (84)	392,222	0.41	297	1.07	273	24
8	Other (33)	178,259	0.19	284	1.02	15	269
9	Other (19)	218,287	0.23	264	0.95	15	249
10	Other (14)	250,297	0.26	232	0.83	8	224

Netflix (5.1 TB/day) is second with a large number of subscribers and high popularity. The other organizations are popular hypergiants, with some primarily offering their own services, such as Facebook, Google, and Microsoft, and others providing network infrastructure and CDNs (Content Delivery Networks), such as Akamai, Amazon, and Fastly. CANARIE is Canada's national research and education backbone network, connecting Canadian universities, educational institutions, and research organizations to each other and to the Internet.

Table 2 shows the results for the corresponding day in 2020 (**Day2020**: 2020-09-23) to illustrate hypergiant traffic during the lockdown. Significant changes in usage patterns are evident in this table, with Amazon and Akamai now at the top, and significant declines for Apple and Netflix. The latter declines are attributable to fewer users on campus. Table 2 also shows Shaw (a major ISP in western Canada) that was not even in the top 20 on Day2019.

One interesting observation when comparing Tables 1 and 2 is that while traffic for most organizations declined significantly from Day2019 to Day2020, Amazon's traffic increased substantially from 1.9 TB/day to 3.3 TB/day. One contribution to this growth is Zoom, since its services are mainly deployed on AWS, and expanded during the pandemic [30]. In particular, Zoom's traffic on

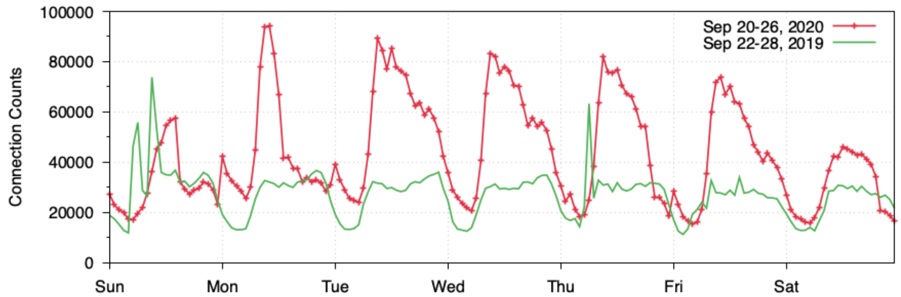

Fig. 2. Hourly connections initiated to authentication servers.

our campus rose from 34 GB on Day2019 to 1,358 GB on Day2020, and represents about 4% of total campus traffic. The key takeaway is that *Zoom traffic, at over 1.3 TB/day, is now comparable to the traffic of other hypergiants.*

Insights can also be gleaned by looking at the internal breakdown of campus network traffic. Table 3 shows pre-pandemic traffic on Day2019, broken down by subnets within our campus network. The top subnets on this list include a BYOD subnet for unmanaged devices, campus WiFi subnets, student residences, and some popular locations with NAT access to the Internet.

Table 4 shows the corresponding traffic breakdown for Day2020 to represent internal subnet traffic patterns during the pandemic. While the top three subnets remain the same, their traffic volumes are much lower, since fewer users are on campus. Several new subnets appear in Table 4, including a subnet for VPN traffic, and a subnet used by UCalgary's administration to update the campus community about the pandemic situation. The labels (numbers in parentheses) for these subnets show their relative pre-pandemic rankings on Day2019. Also of note, the traffic volumes from several student residences[2] decreased, since occupancy was limited; these subnets no longer appear in the Top 10. The main insight from our analysis is that *there were significant structural changes in network usage, both internally and externally. For example, VPN and Admin usage rose, while RezNet and WiFi decreased. The latter contribute to the concomitant decreases in Apple and Netflix traffic.*

4.3 Authentication

Our next analysis focuses on the authentication-related traffic, as we study the network usage patterns of our campus community during the pandemic. All faculty, staff, and students must authenticate themselves with their credentials when using enterprise services, such as email, LMS, VPN, and so on.

[2] We have not analyzed the residence traffic in detail, since the number of users seems low. Ulkani *et al.* [33] studied pandemic effects on student residence traffic at UCSD, finding changes (for example) in Zoom and OSN usage.

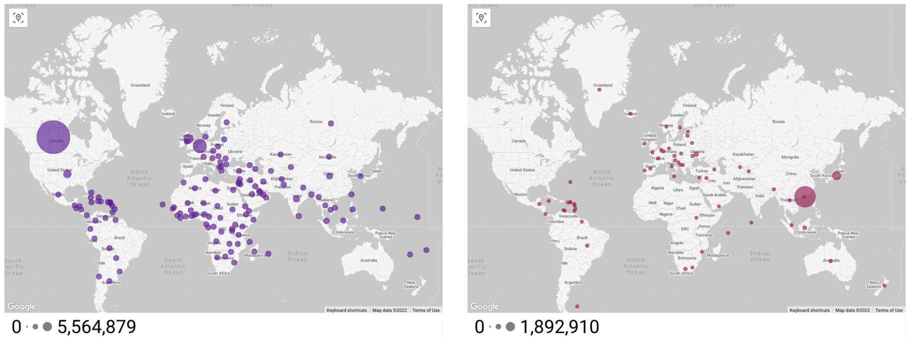

(a) Countries with increase in connections (b) Countries with decrease in connections

Fig. 3. Changes in authentication connections from Sept 2019 to Sept 2020 based on the countries of origin. The maximum numbers in the legends (under the maps) demonstrate the maximum change that a country experienced, i.e., increase (in Canada) for (a) and decrease (in Hong Kong) for (b).

Figure 2 shows the authentication-related traffic for two selected weeks from our empirical dataset: one in September 2019, and the corresponding one in September 2020. The graph shows the number of connections initiated to the authentication servers in each one-hour interval during the week.

There are two main observations evident from Figure 2. First, prior to the pandemic (September 2019), authentication traffic tended to have two peaks per day on weekdays, with one peak in the morning, and one in the evening. This pattern reflects users logging in from home (e.g., checking email, accessing course Web pages) as part of their daily routine both before and after their time on campus[3] for the workday. Second, there is a substantial increase in authentication traffic in September 2020, as a new cohort of students joins the campus community, and many people are working and learning from home. There is a single peak to the traffic each day, with most authentications happening in the morning, and some sessions lasting several hours[4], if not all day, since there was no need to logout and commute to/from campus anymore. Weekend traffic is substantially lower than weekday traffic, as expected. The weekend peaks are also time-shifted to slightly later in the morning. The main insight is that *working from home shifts the usage patterns and leads to prolonged sessions with campus servers.*

Figure 3 demonstrates the changes to the number of connections to the authentication servers, comparing data from September 2019 to September 2020. Each bubble represents a country and shows the absolute amount of change in the authentication connection counts. The size of the bubbles is relative to the max-

[3] Recall that any additional authentication sessions initiated while on campus would not be observable from our monitor.

[4] The mid-July configuration change to the monitor restart interval (now 6 hours) contributes to the observed increase in connections as well.

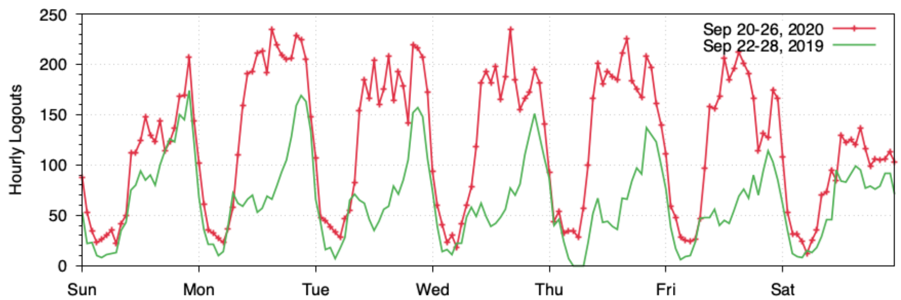

Fig. 4. Hourly LMS Connections during two weeks in Sept 2019 and Sept 2020.

imum change observed (i.e., increase in connections from Canada). Figure 3(a) shows the countries whose number of authentication connections increased, while Figure 3(b) shows the countries with a decrease in their counts. These maps show that connection count increases were most prominent from locations in Canada, followed by the Netherlands, UK, and the US. For connection count decreases, Hong Kong had the largest change, with Japan, Ukraine, and Indonesia next.

4.4 Learning Management System (LMS)

Figure 4 shows hourly connection counts to our LMS during a week in September 2020 and the corresponding week in September 2019. UCalgary uses D2L (Desire2Learn) for LMS, and it is hosted in Quebec, Canada. Despite being in the cloud (AWS), users are redirected to the campus authentication servers at both the start and the end of LMS sessions, enabling counting of this traffic.

Figure 4 shows significant changes in LMS traffic patterns, similar to those observed for the authentication traffic. In September 2019 (before the pandemic), students were regularly on campus, so their LMS authentication traffic was not always visible from our monitor. A peak in the evening when most people were back home is most evident on the green line. When working from home, however, this traffic is more observable throughout the day, as reflected in the higher activity levels in September 2020, with significant changes in its pattern. Diurnal patterns are still evident, with a decline on weekends.

4.5 Remote Access

This subsection discusses the usage of three popular remote access protocols, namely SSH, VPN, and RDP.

Figure 5(a) illustrates the daily connection counts for these protocols for the entire calendar year of 2020, while Figure 5(b) shows the daily byte traffic volume associated with these connections. Note that the vertical axes are logscale (base 10) for better visibility of the entire data. Overall, these results show the dominance of SSH (purple line) in terms of the number of connections (some

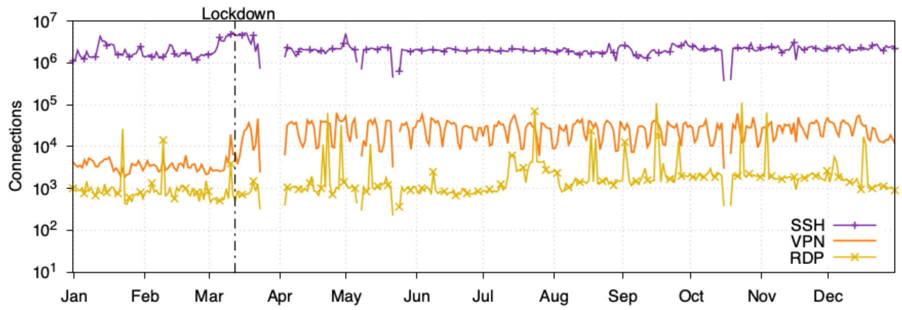

(a) Daily connection counts

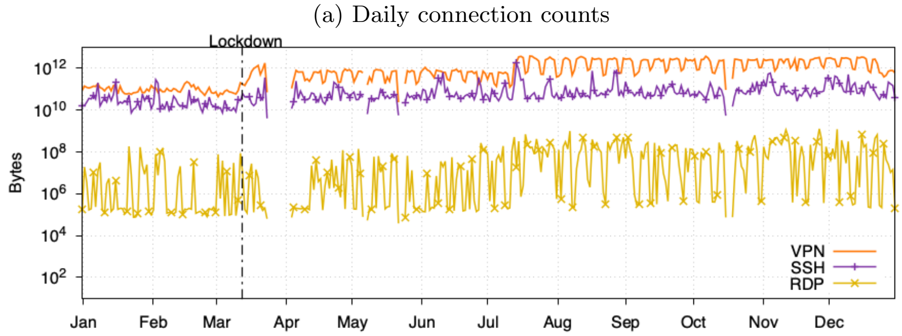

(b) Daily byte traffic volume

Fig. 5. Comparison of SSH, VPN, and RDP usage in 2020.

of which may be generated by scripts or automated processes), while VPN connections (orange line) account for the most data bytes. RDP (gold line) has the lowest activity for both connections and data volumes since it is only applicable for Windows users, and requires a registered system on campus in order to establish an RDP connection. Therefore, lower usage for RDP is unsurprising.

Our further investigation revealed that the increase in the number of SSH connections right before the lockdown is attributable to the increase in inbound scanning activity. Interestingly, the SSH connection count remains pretty steady throughout the year and does not exhibit the typical human-driven weekly patterns evident in the VPN traffic. However, the SSH data volume did increase 2-5x compared to the pre-pandemic baseline in February, reflecting changes to the monitor's visibility of this traffic after the July configuration change. The 6-hour restart interval improves visibility into long-duration TCP connections (refer to Appendix). This is important for applications like VPN and SSH, which often last several hours, if not all day, and it explains the larger proportionate increase in byte traffic volume than in connection count. On our network, SSH usage seems more research-driven rather than education-driven.

Daily VPN connections and data volume both increased after the lockdown by a factor of about 10x. This increase occurred almost immediately following the work-at-home mandate in mid-March 2020. A later increase is also evident

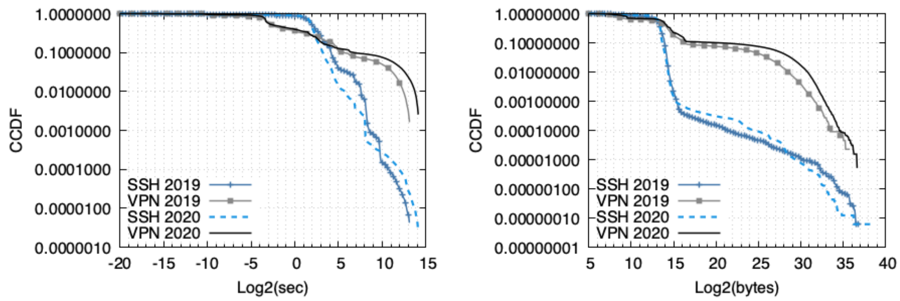

(a) Connection durations

(b) Traffic bytes

Fig. 6. LLCD of connection durations and traffic bytes for SSH and VPN connections during two separate weeks (September 22-28, 2019 and September 20-26, 2020).

in mid-July 2020, when the change in the monitor configuration enhanced observability of longer-duration connections. Although there are fewer VPN connections than SSH, VPN connections tend to have longer durations and transfer more data bytes than SSH connections. Figure 6 illustrates these effects.

The growth in VPN traffic is consistent with observations made by others [14]. For our campus network, the VPN has a broader set of users than the other remote access protocols, since it is available to the entire community. The VPN has greater flexibility as well, since traffic from multiple network applications can be transferred via its connection. Therefore, such an increase in VPN usage is not surprising. In fact, after the lockdown, the primary option to access the campus network was to use the VPN. Many students returned to their home cities or countries, and a lot of newly admitted international students had to commence their programs from abroad. Using the VPN has been the primary means to facilitate this access.

A separate analysis (not shown here) of the origin cities of VPN connections confirms that the increased connection count comes from a larger set of external IPs accessing the network from all over the world. The main insight is that *VPN usage increased dramatically, in terms of connections, data volumes, session duration, IP addresses, and geographical distribution.* Our further investigations did not find any evidence of VPN-related performance degradation on the campus network or repercussions for the nearby clients.

5 Zoom Measurement Results

This section provides an in-depth look at Zoom network traffic, motivated by the growth and volume of this traffic as identified in the previous section. We begin with a look at videoconferencing applications to provide a comparison point for the Zoom traffic.

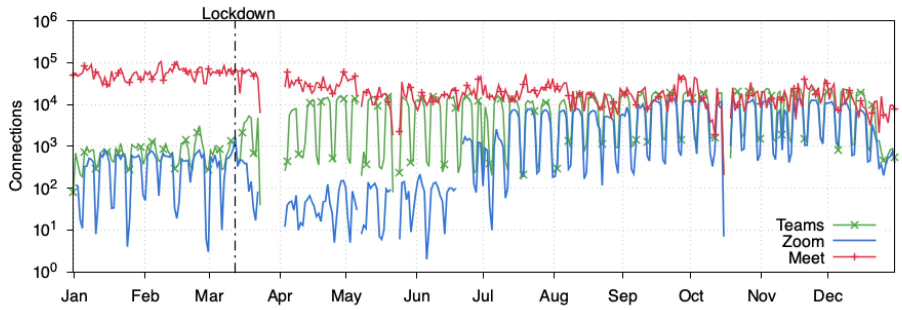

(a) Daily connection counts

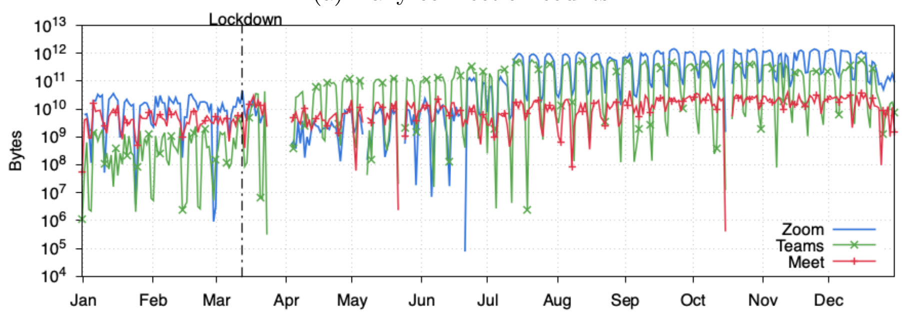

(b) Daily byte traffic volume

Fig. 7. Comparison of Zoom, Teams, and Meet usage in 2020.

5.1 Videoconferencing Apps

This subsection discusses the network traffic measurement results for three online meeting applications (i.e., Zoom, Microsoft Teams, and Google Meet). Google Meet is a relatively new video conferencing app launched by Google in 2017. Prior work indicates that Meet usage increased during the COVID-19 pandemic, especially after Google relaxed its meeting size constraints for unpaid users [13]. We identify Zoom traffic based on the ports and IP ranges[5] provided in their Web site documentation. Similar principles apply to our identification of traffic for Teams and Meet. We show graphs of daily connection counts and traffic volume for each of these applications and compare them accordingly.

Figure 7(a) illustrates the daily connection counts for the three applications, while Figure 7(b) shows the corresponding daily byte traffic volumes for each. Both plots show the entirety of calendar year 2020, illustrating the traffic generated by on-campus users when accessing these externally-supported applications. The gaps in the plots are due to the monitor outages mentioned earlier.

Figure 7 shows the emergence of Zoom in our post-secondary learning environment in 2020. Prior to the work-at-home order in March 2020, Google Meet

[5] https://support.zoom.us/hc/en-us/articles/201362683-Network-firewall-or-proxy-server-settings-for-Zoom

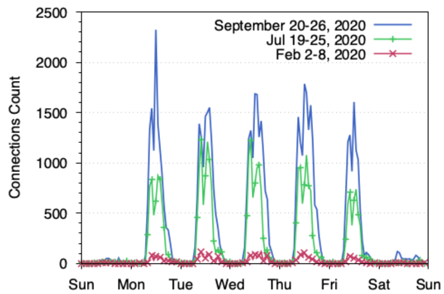

(a) Connections

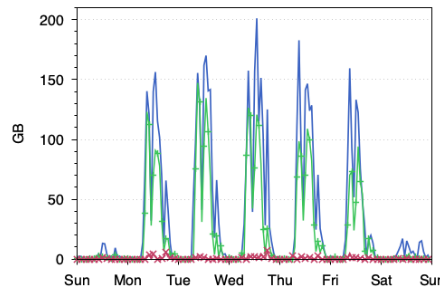

(b) Traffic

Fig. 8. Hourly connections and byte traffic volume for Zoom on three selected weeks in February, July, and September, 2020.

(red line) was the most popular conferencing app, with almost 100K connections per day. At this time, Zoom (blue line) had only 1K connections per day. By September 2020, however, Zoom had reached a level of connection activity comparable to Meet, while far exceeding Meet in data traffic volume. Similarly, Microsoft Teams traffic (green line) grew significantly for remote work and learning after March 2020, and has usage patterns very similar to Zoom.

The number of connections to Google Meet has actually decreased after the lockdown. One reason is fewer people on campus, and another is that Teams and Zoom were adopted as the official online meeting and conferencing app for our university. In particular, the total number of Meet connections from July onwards has decreased by 60-70% with respect to that number in February. However, the byte traffic volume for Google Meet did not decline much at all, suggesting more prolonged usage. These observations also suggest that a significant portion of connections to Meet are system-generated probes by the Meet app on the phones or when users access Gmail.

The daily connections to Teams, and its data traffic volume, increased tenfold right after the lockdown. This surge reflects the shift of administrative meetings (for faculty and staff still present on campus) to the remote format. On the other hand, daily Zoom connections and traffic declined after March 2020 since few students remained on-campus. It was not until the Summer 2020 semester that Zoom usage grew, since more classes were offered then.

In mid-July 2020, we made a configuration change in the monitor (as described earlier), which enabled better tracking of long-duration connections. Consequently, there are increases observed in connections and data volumes for both Zoom and Teams since then. With this change, we have a more complete view to compare with the baseline before the lockdown. For example, comparing the measurements in July, August, and September with February (baseline) shows that the total number of connections to Zoom in July is about 9.5x that in February, and this ratio for August and September is 11x and 20x, respectively. The corresponding ratios for the aggregated byte traffic volume are 27x, 36x,

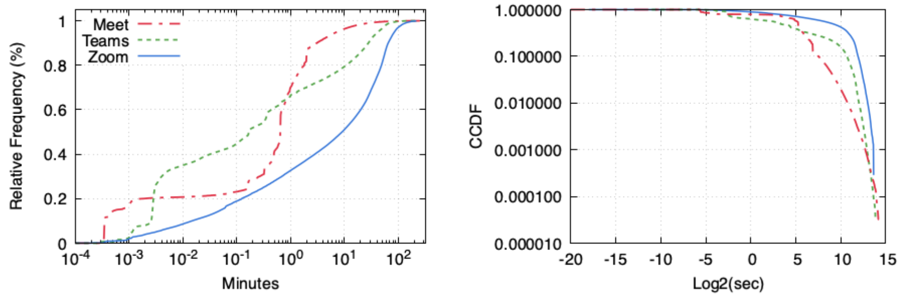

(a) CDF (b) LLCD

Fig. 9. Distribution of connections durations for three applications (Sept 21-25, 2020).

and 48x. These numbers illustrate the pronounced effects of Zoom following the lockdown.

The growth in traffic for Teams is even more dramatic. The total number of connections in July, August, and September are (respectively) 9x, 11x, and 14x that for February, while total byte traffic volumes are 424x, 447x, and 448x that in February. This increase in Teams traffic after the lockdown shows the prevalence of the application among staff and faculty who are still on campus after the lockdown. In February, Zoom traffic volume was 27x larger than the Teams traffic. However, this dominance decreases to 1.6x, 2x, and 2.6x in July, August, and September, respectively. These trends may reflect different bit rate, video resolution, or compression settings in the two applications [21]. Even with fewer people on campus, there has been *a significant increase in Zoom and Teams traffic on our campus network*.

Figure 7 shows a strong weekly usage pattern for Zoom and Teams, both in the connection counts and the data traffic volume. Every hump represents five consecutive working days of network activity, while the valleys show the weekends where those activities are reduced. However, this weekly pattern is less prominent in the Google Meet traffic, especially for connection counts, which implies the system-generated nature of many of these connections.

Figure 8 takes a closer look at diurnal usage patterns in Zoom traffic at a finer-grain time scale. Figure 8(a) illustrates the hourly counts for Zoom connections in three separate weeks from February, July, and September. Figure 8(b) shows the hourly byte traffic volume (in GB) for the same weeks, with inbound and outbound traffic combined. In both plots, there is a clear diurnal pattern, with increases in connections and byte traffic volume during normal working hours, and a decline overnight. Recall that the week in February was pre-pandemic, and the Zoom traffic was negligible. Nonetheless, the connections were established during working hours. The week in July represents the lockdown period. Although many restrictions were lifted by that time, it was after the monitor's configuration change, and the data is more complete. The week

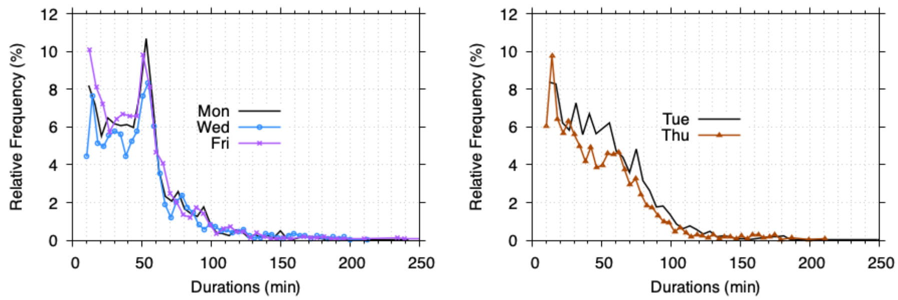

(a) Monday, Wednesday, and Friday (b) Tuesday and Thursday

Fig. 10. Distribution of Zoom connections durations (Sept 21-25, 2020).

in September is after Fall 2020 classes began, and we see increased traffic, as expected.

Figure 8 shows two notable peaks per weekday in the selected weeks from July and September. The first one is in the morning and the second is in the afternoon, both during working hours. On some days, there is a third peak in the late evening, especially in the traffic volume graph. All these peaks in the network traffic represent diurnal patterns from human-driven behavior.

Figure 9(a) illustrates the Cumulative Distribution Function (CDF) of the connection durations for each of the three meeting applications under study during five working days of September 2020. For Zoom, 80% of the connections are less than 50 minutes. For Teams, 90% of the connections last less than one hour. For Google Meet, the vast majority of connections have very short durations, often less than a minute, once again suggesting the machine-generated[6] nature of them rather than human-generated. However, the tail of the distribution for all three applications extends well to the right, with some connections lasting up to 5.5 hours, as can be seen in Figure 9(b). The main takeaway from Figure 9 is that *Zoom connections tend to have longer durations than Teams and Meet*, reflecting usage of Zoom for classes and workshops with prolonged durations.

We next analyze connection durations for Zoom traffic in particular. At UCalgary, most courses are offered on either the Monday/Wednesday/Friday (MWF) schedule with 50-minute lectures, or the Tuesday/Thursday (TuTh) schedule with 75-minute lectures. Since most classes were delivered via Zoom after the lockdown, we expect to see some evidence of that in the distribution of Zoom connection durations.

Figure 10 shows the empirical distribution of Zoom connection durations for five consecutive working days from Fall 2020. Figure 10(a) for MWF confirms the expected peak around 50 minutes. Figure 10(b) for TuTh shows a small peak near 75 minutes duration, with a wide range of other values observed. Note that for courses with labs or tutorials, the timings may be different. For example,

[6] A more detailed analysis shows that some of these are for STUN (Session Traversal Utilities for NAT) protocol traffic on UDP port 19302.

Table 5. Breakdown of transport protocols of Zoom connections on 2020-09-23.

Protocol	Connections	Outbound	Inbound
TCP	308,688	6.16 GB	9.18 GB
UDP	20,461	361 GB	981 GB

tutorial slots are usually 50 minutes, regardless of which day of the week they occur. These observations are congruent with our expectations, and *confirm the widespread usage of Zoom for class delivery at UCalgary.*

5.2 Detailed Traffic Analysis

Via active measurement experiments, we have gained further insights [6] into the structure of Zoom sessions. Note that there are several different ways to set up Zoom, depending on the client application, server deployment, or cloud solution in use. Our university uses the default approach with remote Zoom servers, and no Zoom Meeting Zones located inside the campus.

The connection process to initiate or join a meeting depends on the client's application (e.g., desktop app, mobile app, or Web browser). When using Zoom apps for one-on-one meetings between two parties, direct peer-to-peer connections are often used to carry the media packets. For meetings with more than two participants, a client-server architecture is used, with a cloud-hosted media server as the central point for collecting and distributing media packets for all participants in the Zoom session. Furthermore, such a typical Zoom session involves four logical connections: one TCP connection for control and management of the session (including chat interactions), and three UDP connections, one for audio, one for video, and one for screen sharing (if used). If the Zoom client is unable to connect via the usual procedure, they are directed to use the Web client, which uses TCP only.

To measure Zoom connections and client application usage on our campus, we picked a representative day (Day2020) and examined transport protocol usage based on the number of connections, as well as inbound and outbound data traffic volumes. Table 5 shows this information. TCP accounts for only 1.67% of outbound traffic, and 0.93% of inbound traffic, indicating that few people use Web access to join meetings. Rather, they use the standard procedure of a Zoom meeting connection using client applications.

For a typical Zoom meeting, a client should have three UDP connections for every TCP connection. However, the results in Table 5 show that the number of TCP connections is 15x larger than the number of UDP connections. There are two possible explanations for this discrepancy. First, there might be network connectivity or performance issues when users connect to Zoom from certain subnets (e.g., WiFi), causing TCP problems. Second, there might be many short-lived TCP connections for administrative management of Zoom sessions. We explore both of these possibilities next.

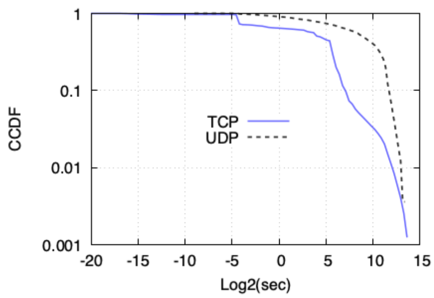

Fig. 11. LLCD of connections durations to Zoom on 2020-09-23.

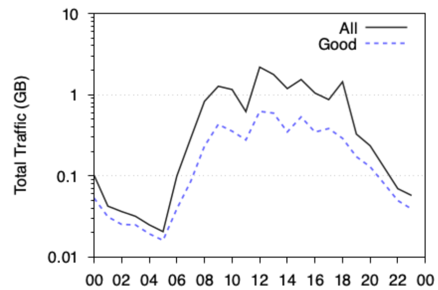

Fig. 12. Hourly TCP traffic of Zoom connections on 2020-09-23.

Figure 11 provides evidence to support these hypotheses. The plot illustrates the Log-Log Complementary Distribution (LLCD) function for Zoom connection durations. Considering the logarithmic scale of the y-axis, we can see a rather significant portion of TCP connections have small durations (under 30 seconds) that cannot be attributed to typical meetings. UDP sessions tend to have longer durations that reflect actual meetings, although the non-negligible portion of relatively short-lived UDP sessions partly indicates the performance issues resulting in Zoom connection interruptions. Waiting rooms, a feature utilized in some Zoom meetings, can also be another cause for the short-lived UDP sessions.

Figure 12 provides another perspective on Zoom session issues on our campus network. It shows the hourly TCP traffic of Zoom connections on Day2020. Connections with typical SYN-FIN handshakes as seen by our monitor are deemed "Good". During the peak hours of the day, only about half of the byte traffic (note the log scale) is exchanged on Good TCP connections, implying that *many connections suffer when too many users on the same network connect to busy Zoom servers*. Note that the administrative machine-generated TCP connections are short-lived with only a few kilobytes of traffic, which do not contribute significantly to overall traffic volume. Although this issue gives us insight into some implications of online learning on our campus network, identifying its root cause requires further investigation, which we leave as future work.

5.3 Zoom Session Management

To better understand Zoom sessions on our campus network, we have analyzed Zoom server usage, as well as the administrative traffic generated between our campus VPN server and the Zoom servers. These results are described next.

Figure 13 is a profile-rank plot to show how connections and byte traffic volumes are distributed across clients and Zoom servers. Figure 13(a) shows the IP frequency-rank profile for servers and clients on Day2020, as well as a year later (**Day2021**: 2021-09-22). Figure 13(b) plots the corresponding IP volume-rank profile for those two days.

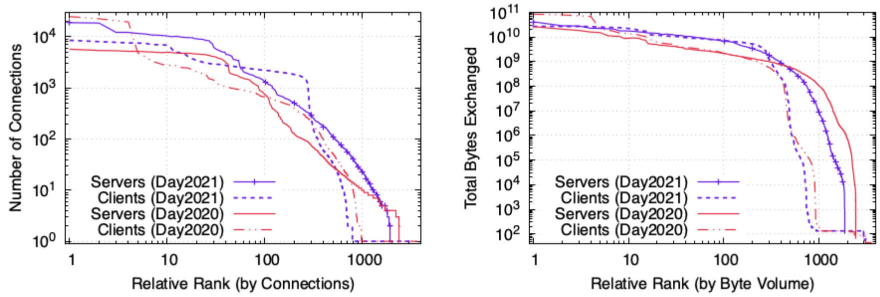

(a) Connections (b) Byte Traffic Volume

Fig. 13. IP profile-rank plots for Zoom traffic on Day2020 and Day2021.

Several key insights emerge from these two graphs. First, in 2020, four client IPs dominated the Zoom traffic, while connections are more widely distributed among a larger set of IPs in 2021. This change reflects the presence of more people on campus, with many using BYOD wireless devices. Second, the load increase on Zoom's servers is also evident in these graphs, both in connection counts (2-3x) and traffic volume (1.5x). Third, although about 2,000 server IPs are seen on a daily basis, most of the traffic is handled by only a couple hundred servers. Furthermore, two stand out in the frequency-rank, reflecting roles in Zoom session management for the campus network. Last but not least, the top 20 server IPs in frequency-rank do not contribute much traffic volume, implying the role of zone controllers, directing clients to selected Zoom Multimedia Routers (MMRs). The main takeway here is that *most of our campus Zoom traffic is handled by a relatively small set of Zoom servers*, leading to possible load issues on those servers. Zoom and other vendors need to provide more detailed information in their client-side dashboard to assist with customer support[7].

Within our own campus network, we have identified one specific server that is directly involved in Zoom session management. This server has at least two different roles. First, it communicates with Zoom servers at the start of each new hosted Zoom meeting to exchange a fixed-size payload, which might be a certificate or authentication credential for licensed users. Second, it generates ICMP "port unreachable" messages to Zoom servers when a Zoom session is aborted, or when an authenticated participant departs prematurely from a meeting. The takeway message from these observations is that *Zoom sessions are complex from the network point of view, and induce extra administrative overhead.*

[7] For example, a light (green, yellow, or red) on the client's view to indicate the performance of the Zoom server from the server's perspective, and possibly tracking over time to summarize the percentage of total meeting time where server performance was green, yellow, or red.

6 Conclusion

The COVID-19 pandemic has had a profound impact on many aspects of people's lives over the past year. In this paper, we provide a detailed look at the network-level effects on inbound and outbound Internet usage on a large campus edge network with over 30,000 users.

The main highlights from our study include the changes in the volume, timing, and directionality of traffic. With fewer users on campus, we observed dramatic changes in the inbound and outbound traffic volumes, as well as a reduction in the degree of asymmetry in the traffic. That is, inbound traffic still dominates outbound, but not by as much as it did prior to the pandemic. There are some perceptible differences in the daily timing of network usage, since commuting to campus is no longer the norm. Furthermore, a geographic analysis of the authenticated users for our campus network shows an increasingly international spread.

Pronounced shifts are also evident in network application usage. The increased traffic volume for Zoom and Teams is dramatic (e.g., 20x-450x), and VPN usage is also much higher (20x) than ever before. Most applications show strong daily and weekly patterns, consistent with the normal workday schedule, even when working from home. Research traffic seems less affected by the pandemic than teaching and learning traffic.

Finally, the results of our analysis reveal that there are issues with Zoom TCP connections and session management on our campus network when many people on campus connect to a limited set of regional Zoom servers during peak hours [6]. These problems manifest themselves with a plethora of short-lived TCP connections, and compromise the user-perceived quality of Zoom sessions. Furthermore, these problems are likely to grow as UCalgary adopts a blended learning model (i.e., a mix of in-person and online learning) in the upcoming academic year.

Acknowledgements

The authors thank the PAM 2022 reviewers and shepherd Tobias Fiebig for their constructive suggestions that helped to improve our paper. The authors are also grateful to UCIT for facilitating our collection of campus-level network traffic data, and to the team at ARC for technical support in the storage and management of our data. Financial support for this research was provided by Canada's Natural Sciences and Engineering Research Council (NSERC).

Appendix

The monitor reconfiguration mentioned earlier happened in the course of a week. On July 6, 2020, we changed the reset interval from one hour to every three hours to test the robustness of the monitor against the large volume of scanning activity and how disabling the scanning module is effective. The experiment was

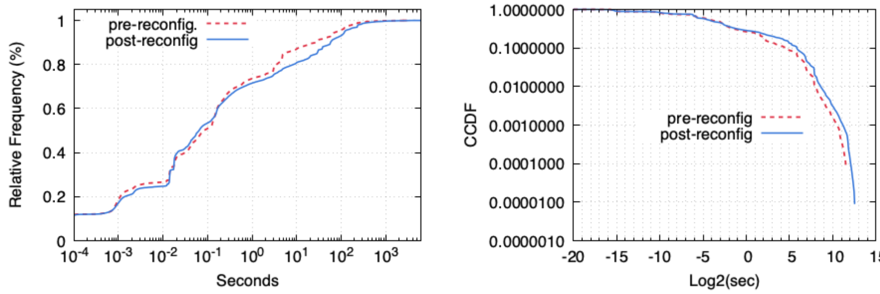

(a) CDF (b) LLCD

Fig. 14. Distributions of connection durations during five working days of June 29, 2020 to July 3, 2020 (before monitor reconfiguration) and five working days of July 13, 2020 to July 17, 2020 (after monitor reconfiguration).

successful, and on July 13, 2020, we again changed the reset interval to every six hours. We then settled with that interval as our subsequent resource monitoring suggested that a longer interval may cause problems.

Figure 14 shows the distribution of connection durations for five working days from June 29 to July 3, 2020 (representing before reconfiguration) and another five working days from July 13 to July 17, 2020 (representing after reconfiguration). Both distributions follow a very similar pattern, with the post-reconfiguration graph stretching slightly to the right and longer tail on the LLCD plot, showing a heavier tail for the distribution that attributes to the connections lasting between 1 to 6 hours (note the log2-based x-axis). However, the most significant difference between these distributions (not evident in this figure) is that more than 615 million connections were captured during these post-reconfiguration days, while this number for the pre-reconfiguration days was more than 570 million. There is about 45 million difference between the number of connections in these distributions, out of which only about 574 thousand lasted between 1 to 6 hours. It shows that the reconfiguration not only helped in capturing longer connections (which is very impactful for some applications, such as Zoom and VPN) but also more connections in general, due to fewer restarts per day.

References

[1] V. Adhikari et al. "Measurement Study of Netflix, Hulu, and a Tale of Three CDNs". In: *IEEE/ACM Transactions on Networking* 23.6 (2014), pp. 1984–1997.

[2] M. Arlitt and C. Williamson. "Internet Web Servers: Workload Characterization and Performance Implications". In: *IEEE/ACM Transactions on Networking* 5.5 (1997), pp. 631–645.

[3] T. Balliester and A. Elsheikhi. "The Future of Work: A Literature Review". In: *ILO Research Department Working Paper* 29 (Mar. 2018).

[4] L. Breslau et al. "Web Caching and Zipf-like Distributions: Evidence and Implications". In: *Proceedings of IEEE INFOCOM Conference*. Vol. 1. New York, NY, USA, Mar. 1999, pp. 126–134.

[5] H. Chang et al. "Can You See Me Now? A Measurement Study of Zoom, Webex, and Meet". In: *Proceedings of ACM IMC*. Nov. 2021, pp. 216–228.

[6] A. Choi et al. "Zoom Session Quality: A Network-Level View". In: *Proceedings of PAM Conference*. Mar. 2022, online.

[7] Zoom Video Communications. *Configure Meeting Connector Controller Port Forwarding*. `http://aiweb.techfak.uni-bielefeld.de/content/bworld-robot-control-software/`. 2020.

[8] M. Crovella and B. Krishnamurthy. *Internet Measurement: Infrastructure, Traffic and Applications*. John Wiley & Sons, Inc., 2006.

[9] T. Favale et al. "Campus Traffic and E-Learning during COVID-19 Pandemic". In: *Computer Networks* 176 (July 2020). Article 107290, pp. 1–9.

[10] A. Feldmann et al. "The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic". In: *Proceedings of ACM Internet Measurement Conference (IMC)*. ACM. Oct. 2020, pp. 1–18.

[11] T. Fiebig et al. "Heads in the Clouds: Measuring the Implications of Universities Migrating to Public Clouds". In: (2021).

[12] P. Gill et al. "Youtube Traffic Characterization: A View from the Edge". In: *Proceedings of the 7th ACM Internet Measurement Conference (IMC)*. San Diego, CA, USA, Oct. 2007, pp. 15–28.

[13] Google. *COVID-19 Community Mobility Reports*. `https://www.google.com/covid19/mobility/`. 2021.

[14] BITAG Technical Working Group. *2020 Pandemic Network Performance*. `https://www.bitag.org/documents/bitag_report.pdf`. Apr. 2021.

[15] S. Klenow et al. "Campus-Level Instagram Traffic: A Case Study". In: *Proceedings of IEEE MASCOTS*. Rennes, France, Oct. 2019, pp. 228–234.

[16] C. Labovitz. *Early Effects of COVID-19 Lockdowns on Service Provider Networks: The Network Soldiers On!* `https://www.nokia.com/blog/early-effects-covid-19-lockdowns-service-provider-networks-networks-soldier/`. Mar. 2020.

[17] C. Labovitz. *Network Traffic Insights in the Time of COVID-19: April 9 Update*. `https://www.nokia.com/blog/network-traffic-insights-time-covid-19-april-9-update/`. Apr. 2020.

[18] A. Lamb et al. "The Vertica Analytic Database: C-Store 7 Years Later". In: *Proc. VLDB Endow.* 5.12 (Aug. 2012), pp. 1790–1801. ISSN: 2150-8097. DOI: `10.14778/2367502.2367518`. URL: `https://doi.org/10.14778/2367502.2367518`.

[19] S. Liu et al. "Characterizing Service Provider Response to the COVID-19 Pandemic in the United States". In: *Proceedings of PAM Conference*. Germany, Mar. 2021, pp. 20–38.

[20] A. Lutu et al. "A Characterization of the COVID-19 Pandemic Impact on a Mobile Network Operator Traffic". In: *Proceedings of the ACM Internet

Measurement Conference (IMC). Pittsburgh, PA, USA, Oct. 2020, pp. 19–33.

[21] K. MacMillan et al. "Measuring the Performance and Network Utilization of Popular Video Conferencing Applications". In: *Proceedings of ACM IMC*. Nov. 2021, pp. 229–244.

[22] Microsoft. *Prepare your organization's network for Microsoft Teams.* `https://docs.microsoft.com/en-us/microsoftteams/prepare-network/`. 2021.

[23] A. Mislove et al. "Measurement and Analysis of Online Social Networks". In: *Proceedings of the 7th ACM Internet Measurement Conference (IMC)*. San Diego, CA, USA, Oct. 2007, pp. 29–42.

[24] L. Nagel. "The Influence of the COVID-19 Pandemic on the Digital Transformation of Work". In: *International Journal of Sociology and Social Policy* 40.9/10 (Dec. 2020), pp. 861–875.

[25] V. Paxson. "Bro: A System for Detecting Network Intruders in Real-time". In: *Computer Networks* 31.23-24 (1999), pp. 2435–2463.

[26] V. Paxson. "Empirically Derived Analytic Models of Wide-Area TCP Connections". In: *IEEE/ACM Transactions on Networking* 2.4 (Aug. 1994), pp. 316–336.

[27] F. Porpiglia et al. "Traditional and Virtual Congress Meetings during the COVID-19 Pandemic and the Post-COVID-19 Era: Is it time to change the paradigm?" In: *European Urology* 78.3 (2020), p. 301.

[28] H. Pratama et al. "The Trend in Using Online Meeting Applications for Learning During the Period of Pandemic COVID-19: A Literature Review". In: *Journal of Innovation in Educational and Cultural Research* 1.2 (2020), pp. 58–68.

[29] A. Ramachandran and N. Feamster. "Understanding the Network-level Behavior of Spammers". In: *Proceedings of ACM SIGCOMM Conference*. Pisa, Italy, Sept. 2006, pp. 291–302.

[30] Amazon Inc. - Press Room. *AWS and Zoom Extend Strategic Relationship.* `https://press.aboutamazon.com/news-releases/news-release-details/aws-and-zoom-extend-strategic-relationship`. Nov. 2020.

[31] M. Santana and M. Cobo. "What is the Future of Work? A Science Mapping Analysis". In: *European Management Journal* 38.6 (Dec. 2020), pp. 846–862.

[32] F. Schneider et al. "Understanding Online Social Network Usage from a Network Perspective". In: *Proceedings of the 9th ACM Internet Measurement Conference (IMC)*. Chicago, IL, USA, Nov. 2009, pp. 35–48.

[33] Alisha Ukani, Ariana Mirian, and Alex C Snoeren. "Locked-in during lockdown: undergraduate life on the internet in a pandemic". In: *Proceedings of the 21st ACM Internet Measurement Conference*. 2021, pp. 480–486.