Zoom Session Quality: A Network-Level View

Albert Choi, Mehdi Karamollahi, Carey Williamson, and Martin Arlitt

University of Calgary, Calgary, Alberta, Canada
{albert.choi1,mehdi.karamollahi,cwill,marlitt}@ucalgary.ca

Abstract. Zoom is a popular videoconferencing application for remote work and learning. In 2020, our university adopted Zoom for delivering online lectures during work-from-home restrictions. Starting in September 2021, however, our university offered both in-person and online classes, leading to increased Zoom usage on our campus network. In this paper, we study this Zoom network traffic in two different ways. First, we perform small-scale active measurements on individual Zoom test sessions to understand communication patterns and traffic structure. Second, we use large-scale passive measurement of campus-level Zoom traffic to understand usage patterns and performance problems. Our results identify 4x growth in Zoom traffic on our campus network since 2020, as well as network-related issues that affect Zoom session quality.

Keywords: Network measurement · Zoom video conferencing · QoE

1 Introduction

Since the onset of the COVID-19 pandemic in early 2020, Zoom has become the primary platform for online learning at many universities worldwide, including the University of Calgary. Zoom grew significantly as a video conferencing platform during the pandemic, due to its ease of use, and the ability to host meetings with a large number of participants [2,17].

In September 2021, our university began to transition back to in-person learning for 50% of its courses. Many students who enrolled for in-person classes also had Zoom-based online lectures for other courses. Due to course scheduling, this situation often required students to be on campus for their online classes, which significantly increased Zoom traffic on our campus network, and led to several anecdotal reports of Zoom performance problems.

In this paper, we describe our approach for studying Zoom network traffic on campus, and our results from analyzing this network traffic. Our work is motivated by fairly broad questions such as how large the Zoom user community is on campus and the number of classes being joined by students while on campus. The data that we collect and analyze also gives insights into the behaviours of students when joining Zoom sessions from on campus, such as when and for how long they join Zoom sessions, and whether they use their camera or microphone. We are also interested in the user-level Quality of Experience (QoE) for Zoom sessions. Our results provide a better understanding of how Zoom is used on our campus network, and identify several network performance issues.

Our main contributions include tools for analyzing empirical Zoom network traffic data, as well as insights regarding Zoom session quality from a network-level perspective. Our results quantify the performance of Zoom with respect to bandwidth usage and QoE. Table 1 summarizes our main observations, which are discussed more fully in the rest of the paper.

Table 1: Empirical Observations about Zoom Network Traffic

Name	Observation	Section
Control	Zoom uses a TCP connection to manage each session, including chat.	4.1
Channels	Zoom uses separate UDP ports for audio, video, and screen sharing.	4.1
Adaptivity	Zoom uses bandwidth probing to dynamically adapt video bit rates.	4.1
Disruptions	Many Zoom sessions experience disruptions on TCP or UDP connections.	4.3
Usage	Zoom has diurnal usage patterns that are driven by class schedules.	5.1
Growth	Zoom traffic on our campus network has grown 4x over the past year.	5.1
Patterns	Session structure and camera usage are discernible from traffic analysis.	5.2
Robustness	Zoom sessions are resilient, and can re-establish connections as needed.	5.3

The remainder of this paper is organized as follows. Section 2 provides a brief overview of prior related work. Section 3 describes our active and passive approaches for collecting and analyzing Zoom network traffic. Section 4 presents results from our small-scale experiments with Zoom, while Section 5 discusses results from our campus-wide look at Zoom traffic. Section 6 discusses the performance implications of our results. Finally, Section 7 concludes the paper.

2 Related Work

Several prior works have studied pandemic effects on Internet traffic, including enterprise networks [1,5], ISP networks [7], mobile networks [9], and academic environments [4,15]. Here we highlight selected papers from the literature.

Feldmann *et al.* [5] studied pandemic effects on Internet traffic, as viewed by ISPs, IXPs, and mobile network operators. They identified 15-20% growth in Internet traffic volumes within a week of the lockdowns taking place in Europe. Their paper identified new network applications being used in work-from-home environments, including Zoom and Teams, that contributed to the traffic growth. Our study focuses specifically on Zoom in a campus setting, with an emphasis on network-level and application-level performance.

Favale et al. [4] studied changes to network traffic induced by the COVID-19 pandemic in an academic setting. The traffic patterns were observed from Politecnico di Torino, a medium-sized university located in Italy. Their study differs from ours in that the online e-learning platform studied was not Zoom, but rather a custom in-house solution developed based on the BigBlueButton framework. The authors found that incoming traffic was greatly reduced across three university campuses that were studied, since few students were on campus to generate inbound traffic. For the campus that hosted the e-learning platform,

however, the outbound network traffic more than doubled. While their university did not adopt Zoom at all, many universities did. Thus our results should be broadly applicable for many Zoom users in academic environments.

Two recent papers [3,10] studied videoconferencing applications, including Zoom. Chang *et al.* [3] developed a cloud-based framework for testing and comparing the QoE for Meet, WebEx, and Zoom. Their work used emulated clients in the cloud to identify QoE differences for these three applications, based on architecture, infrastructure, geographic location, lag, video bit rate, and network bandwidth constraints. They used free-tier versions of these applications, and real mobile devices, in meetings with up to 11 emulated participants. MacMillan *et al.* [10] compared Meet, Teams, and Zoom in an experimental testbed. Their work focused on network utilization, robustness, and fairness [14] between and among these videoconferencing applications under different emulated network conditions. They also studied the scaling properties of these applications with up to 8 participants. Our work differs from both of these papers in that we focus on developing tools to analyze empirically-observed Zoom sessions generated by thousands of users on our campus network. Our companion paper [6] provides a longitudinal (macroscopic) view of Zoom, Teams, Meet, VPN, and other applications on our campus network during the pandemic; here, we provide a detailed (microscopic) look at Zoom.

There are also several Zoom-specific papers that focus on privacy and security issues [11,12], rather than network performance or QoE. Marczak and Scott-Railton [12] studied Zoom network traffic to investigate potential security vulnerabilities with the platform. Mahr et al. [11] analyzed the types of data that Zoom sends in either encrypted or unencrypted form. In our paper, we focus on the network traffic patterns and application-level performance of Zoom.

3 Methodology

Our work uses a combination of active and passive approaches to network traffic measurement. We use the resulting data to study on-campus Zoom sessions from both network-level and application-level perspectives.

We first analyze Zoom traffic at a small-scale to extract insights regarding communication patterns. We do so using Wireshark [16] for a packet-level view of Zoom test sessions under our control. Although all Zoom traffic is encrypted, we can still obtain useful information regarding the Zoom application, such as IP addresses and port numbers, session structure, and bandwidth consumption.

Data collection is done using Windows and Linux laptops. Wireshark is used to capture packets locally during Zoom test sessions. Data collection is done in a controlled environment where we are the only users in a Zoom session. Our active measurements only involve users who opted in for data collection.

Using knowledge from the small-scale experiments, we subsequently analyze campus-wide Zoom traffic at our university, which has approximately 35,000 students, faculty, and staff, and a 10 Gbps external connection to the Internet. The campus-level traffic data consists of summarized traffic logs from our campus

network monitor (Endace DAG; Dell PowerEdge running Zeek). The connection logs provide detailed information about each observed connection, such as start time, protocol endpoints, bytes/packets transferred, and connection duration (see Appendix 1 for an example). These logs are collected in cooperation with our campus IT staff, and with permission obtained via our campus research ethics review process. We conduct TCP/IP traffic analysis using packet headers only (no payloads), focusing on aggregate traffic characteristics from thousands of users, most of whom use transient IP addresses from DHCP/NAT.

We collect our passive measurement data from a mirrored stream of all traffic that passes through the campus edge router. This means that we can observe all traffic that has at least one endpoint on the campus network and at least one endpoint on the Internet. In particular, Zoom traffic is observable when students or instructors on campus connect to a Zoom server. We identify Zoom traffic based on the 118 IPv4 network prefixes[1] indicated on the Zoom Web site. We also use information about the Zoom connection process[2] and local DNS logs to distinguish Zoom MultiMedia Routers (MMR) from other Zoom server types (e.g., API, CDN, logging, Web, XMPP, Zone Controller).

4 Small-scale Measurements

This section presents results from our small-scale measurements of Zoom test sessions using Wireshark. The purpose is to understand the structure of Zoom traffic, and enable Zoom analysis using our campus-level connection logs.

4.1 Zoom Session Structure

Our first experiment[3] focused on the structure of a Zoom test session that lasted about 45 minutes. The meeting host was on campus, using the university-licensed version of Zoom on a wireless laptop. The second participant was on their own laptop at home. A third participant joined the meeting late, and then left early.

Figure 1(a) shows a time-series plot of the packet traffic during this Zoom test session, based on the Wireshark trace. After launching the meeting, the host waited for others to join. During this time, communication took place between the host's laptop and the Zoom MMR server, using one TCP control connection (port 443), and three separate UDP connections[4] for audio, video, and screen-sharing (data), respectively. All media traffic was sent to the MMR server on UDP port 8801, with periodic keep-alive packets when idle. TCP is used to manage the session, and handle any chat messages[5] sent during the meeting.

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

[2] https://zoom.us/docs/doc/Zoom Connection Process Whitepaper.pdf

[3] Other experiments tested different features (e.g., camera, microphone, chat, screen-sharing, waiting room) to learn more about Zoom, similar to the approach in [8]. See Appendix 2 for two additional examples of such sessions.

[4] Though UDP is connectionless, we refer to these as UDP connections or *channels*.

[5] Documents sent via chat use a separate TCP connection to an XMPP server.

When the second participant joined 6 minutes later, Zoom switched into peer-to-peer (P2P) mode (shown in purple). In this mode (P), all media traffic is delivered directly between the two participants using a single UDP connection with ephemeral ports at each end [3]. However, the TCP control (C) connection to the MMR server remained in place. The total bit rate in P2P mode fluctuated between 3 Mbps and 6 Mbps, with an average of 5 Mbps.

When the third participant joined 11 minutes later, Zoom switched back into server mode, with all traffic from each laptop being sent to the MMR server on UDP port 8801. In this mode, Zoom uses three UDP channels (with ephemeral ports at each client) to send audio, video, and screen-sharing data to each participant. The packet rates and sizes on each channel are quite distinctive (see the CDF plots in Figure 1(b) and (c)), enabling simple heuristics[6] to identify the channels. These channels are colour-coded in the graph with blue for video (V), orange for audio (A), green for screen-sharing data (D), and black for total UDP traffic volume (T). The total bit rate in this mode was about 3.5 Mbps, substantially higher than the 1.0-1.5 Mbps rate reported for free-tier Zoom [3,10]. As is evident from the graph, video accounts for most of the UDP traffic observed.

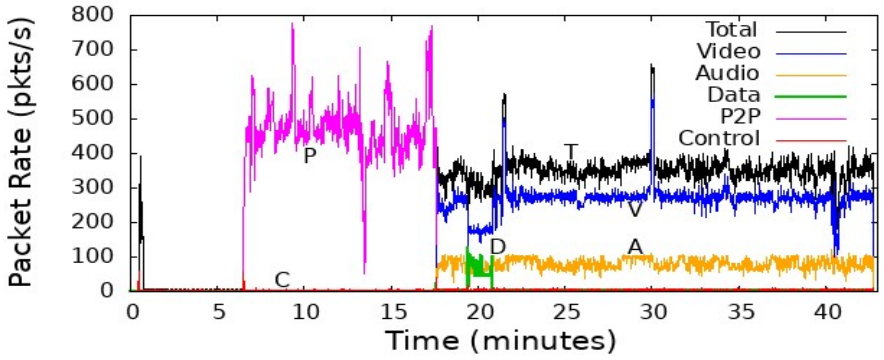

(a) Packet Rate Time Series

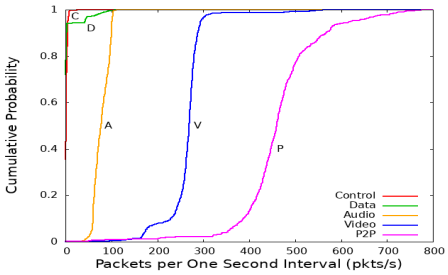

(b) Packet Rate CDF

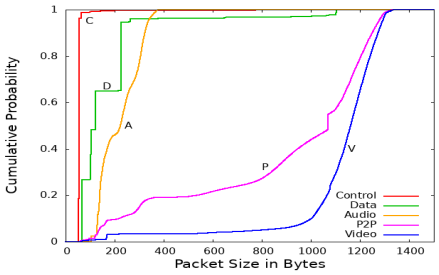

(c) Packet Size CDF

Fig. 1: An Example of a Zoom Test Session

[6] We use threshold-based strategies, with average packet size as the primary feature, and average packet rate as a secondary feature. Directionality is also important.

Several other observations are evident from Figure 1(a). First, the overall packet rate decreased when switching from P2P mode to Zoom server mode, likely reflecting bandwidth management techniques used by the server. Second, when the third participant leaves, Zoom does *not* switch back into P2P mode; rather, it stays in central server mode for the remainder of the session. Third, when screen-sharing happens (at 20-minute mark), the video bit rate is dynamically reduced until screen-sharing is complete. Fourth, the two upward spikes suggest that Zoom uses dynamic bandwidth probing to adapt the video bit rate during the session. Each spike represents a higher packet rate from the MMR server for 10 seconds. Finally, the TCP control connection (red) lasts throughout the Zoom session, regardless of the number of participants.

4.2 Zoom Session Profiles

In addition to time-series plots of packet traffic, we construct *session profiles* to visualize Zoom session structure and identify performance anomalies.

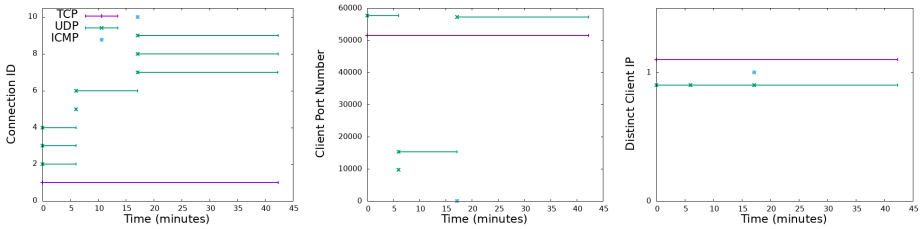

(a) Connection Profile (b) Client Port Profile (c) Client IP Profile

Fig. 2: Session Profile for a Single Client in the Zoom Test Session

Figure 2 shows the session profile for the Zoom test session from Figure 1. Figure 2(a) shows the *connection profile*, which illustrates the TCP and UDP connections used during a Zoom session. The vertical axis shows a monotonically increasing ID number for each connection (TCP, UDP, or ICMP) in the order of their creation, and the horizontal axis shows the elapsed time duration for each connection. A normal Zoom session has four horizontal lines for each participant, representing one TCP connection and three UDP connections launched almost simultaneously, and with very similar durations. An exception is when switching between P2P mode (with a single UDP connection) and Zoom server mode, since three new UDP connections are created then, while the TCP connection remains the same. Figure 2(b) is a *port profile*, which shows the port numbers used by each client connection, which sometimes reveals OS-specific behaviours. Notably, the UDP port numbers here are consecutive (suggestive of Microsoft Windows), and the P2P session occurred on a distinctly different port than the others. Because of the vertical scale of the graph, some of the concurrent UDP connections blend together, though they are separate connections as seen in

Figure 2(a). Figure 2(c) shows the *IP profile*, which stitches together disrupted connections (if any), and groups them based on the client's IP address. We vertically offset TCP from UDP so that they do not obfuscate each other, and bundle the UDP connections together to reduce visual clutter (e.g., this helps a lot when many clients are being shown).

4.3 Anomalous Zoom Sessions

The session profile plot provides a useful visualization tool for assessing Zoom session quality, and identifying anomalous Zoom sessions.

Figure 3 shows the session profile from a Zoom meeting that had very poor quality (i.e., two Zoom restarts, and two additional audio outages for the user). Figure 3(a) shows that this particular session had numerous disruptions to the TCP control channel, which had trouble reconnecting, and finally stabilized[7] after about 10 minutes. There were about 40 TCP connection attempts, though only four Zoom impairments were perceptible at the user level. The two Zoom restarts resulted in new UDP connections for all three channels (labelled 'AVD' in Figure 3(c)). There were also two other disruptions to the audio channel (labeled 'A'), which resulted in the creation of a new UDP connection. The dynamic port selection for the disrupted audio channel is evident in Figure 3(b), while that for the TCP control channel is less evident, since these port numbers are often contiguous. Figure 3(c) shows the IP profile for this session. The overall structure of the Zoom session is clearer on this plot. In particular, the TCP channel had many disruptions, which affected the control channel and Zoom connectivity. The UDP channels had fewer disruptions, but still degraded session quality.

Surprisingly, even small Zoom meetings in P2P mode can experience disruptions. As an example, the second row of Figure 3 shows the session profile for a two-person meeting, with one on campus, and the other at home. The on-campus laptop for which we have a Wireshark trace has TCP (red) and UDP (orange) connections that start near time 50 on the graph. The UDP connection was normal, but the TCP connection was disrupted repeatedly for over half an hour. Analysis of the Wireshark trace and the campus-level data shows numerous connection resets, often in alternating fashion with another campus host (green for TCP) that was trying to set up a different meeting with the same Zoom server. The graph also shows a third host (blue for TCP, and magenta for UDP) that had started another meeting with the same Zoom server about an hour earlier. Its session was fine until near the end, when the other sessions started.

Figure 3 leads to several insights. First, it is quite common for multiple meetings on our campus network to share the same Zoom server at the same time. This behaviour differs from that reported in prior work, where every new meeting used a different Zoom server [3]. We attribute this difference to the large number of Zoom meetings on campus, and the limited set of Zoom servers available in our region. Second, things can go wrong when too many Zoom meetings concurrently

[7] The client contacted a Zone Controller (ZC), and then switched the TCP control channel (but not the UDP channels) to a different Zoom MMR server.

use the same busy Zoom server. Our Wireshark traces show lots of TCP duplicate ACKs, spurious retransmissions, and the Zoom server sending an "Encrypted Alert" to terminate the TLS session with the client, which then has trouble reconnecting to the server. Third, the network-level effects are most evident on TCP connections. Detailed analysis shows a mix of unsuccessful connection attempts, resets by the server, and resets by the client. Several exhibit sluggish Server Hello messages in the TLS handshake (e.g., taking 2-8 seconds instead of 25-30 milliseconds), causing the client to timeout and reset the connection.

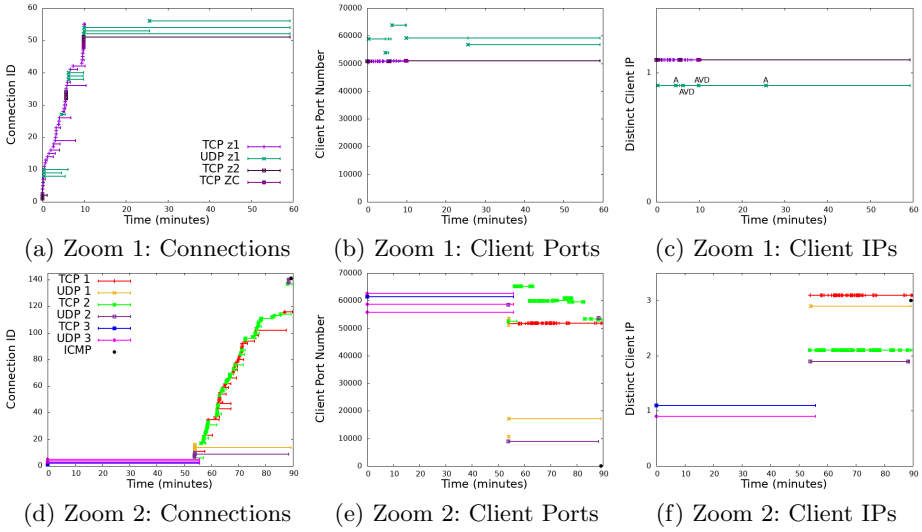

(a) Zoom 1: Connections (b) Zoom 1: Client Ports (c) Zoom 1: Client IPs

(d) Zoom 2: Connections (e) Zoom 2: Client Ports (f) Zoom 2: Client IPs

Fig. 3: Session Profile Examples for Two Anomalous Zoom Meetings

5 Large-Scale Measurements

Our next set of measurements focus on larger examples of Zoom sessions extracted from our campus network traffic logs. For these sessions, we have no packet-level Wireshark traces, since we do not have permission to collect such traces during live classes or meetings. Nonetheless, we use the empirically-observed traffic data from connection logs to infer information about the Zoom sessions.

5.1 Zoom Usage Patterns

Figure 4(a) shows a time series plot of the aggregate Zoom traffic observed on our campus network for a representative day (Wed Sept 22, 2021) from our Fall 2021 semester. The lines on the graph represent the total number of concurrent UDP 8801 connections (purple) to Zoom MMR servers, as well as the corresponding

number of Zoom sessions (green) and Zoom meetings (blue) determined using our traffic analysis tools. On this day, the peak load reached about 3,500 concurrent UDP connections, representing about 1,100 Zoom sessions (i.e., users) in about 250 different concurrent Zoom meetings. There were a total of 3,679 different Zoom meetings on this day. Each meeting had about 4 on-campus users, on average, and each user session had about 3 UDP connections.

Figure 4(a) shows that Zoom usage exhibits a diurnal pattern that corresponds to when people are on campus. The busy period starts in the morning with usage peaking mid-day[8] and declining towards the evening. Class schedules are also evident in the overall Zoom traffic, with classes starting every hour between 8:00am and 5:00pm on this specific day.

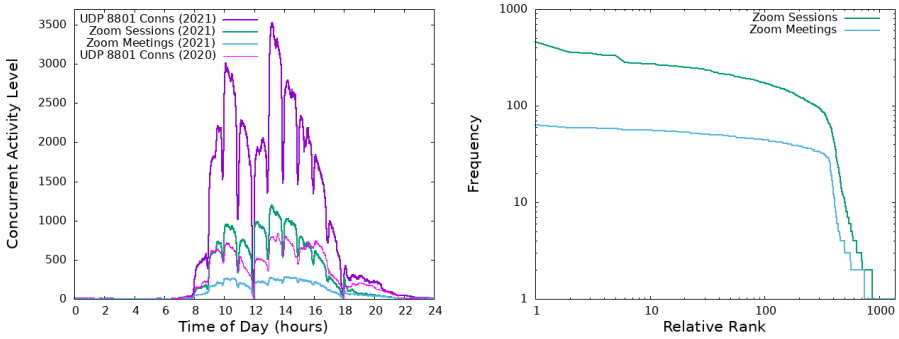

(a) Zoom Traffic Patterns (b) Zoom Server Usage

Fig. 4: Aggregate Zoom Traffic Characteristics: Demand vs. Supply

Figure 4(a) also illustrates the growth in Zoom traffic on our campus network between Fall 2021 and the corresponding day (Wed Sept 23, 2020) from the Fall 2020 semester, when only 20% of courses were offered in person. We use the count of concurrent UDP 8801 connections for this comparison (i.e., purple line versus thin magenta line). The results show that Zoom traffic has increased 4x from Fall 2020 to Fall 2021; in fact, it now exceeds 1 TB per day [6]. This growth in demand is substantial, and reflects the impacts of mixed modalities for learning, with 30-50% of students back on campus.

Figure 4(b) takes a different perspective, by looking at the supply of Zoom servers. This graph is an IP frequency-rank plot (on a log-log scale), showing the relative frequency with which different Zoom MMR server IPs are used for meetings. The data is from a one-week period in September 2021, during which there were 17,120 Zoom meetings, involving a total of 1,374 Zoom MMR

[8] The dip at 12 noon is an artifact of our campus network monitor, which is restarted every 6 hours to avoid possible crashes during high-volume scans [6]. This restart unfortunately loses information about connections in progress at 6:00am, noon, 6:00pm, and midnight. This artifact is most evident at 12 noon, when load is higher.

servers. However, Figure 4(b) shows that the usage of Zoom servers is highly non-uniform. In our dataset, the top three /24 Zoom network prefixes accounted for 42%, 35%, and 13% of the meetings, using 183, 150, and 76 MMR servers, respectively. Each of these MMR servers was used, on average, about 40 times during the week, or about 6 times each day.

These results show that the number of Zoom servers available to our campus users is limited, with about 400 servers handling 90% of the meetings. The reason for this is our campus Zoom configuration, which directs work-related meetings to regionally-hosted Zoom servers within Canada (primarily Vancouver and Toronto, as well as Etobicoke, Ontario). The other Zoom server IPs in our dataset appear only a few times, possibly from personal (free-tier) Zoom usage, or on-campus people attending meetings hosted by colleagues elsewhere.

We believe that this mismatch between supply and demand is the root cause of the Zoom performance anomalies that we have observed on our campus network. At peak times of the day, some MMR servers are managing multiple large meetings, and are unable to cope with the corresponding load. This phenomenon might be specific to our university's Zoom configuration, but could occur elsewhere if the supply of regional Zoom servers is quite limited.

5.2 Session Characteristics

By focusing on a single Zoom server IP address at a time, we can identify specific Zoom sessions for classes or meetings. As a sanity check for these sessions, we also check the consistency of the end time, and the data volumes exchanged on each connection. For such an event, we can then determine the number of participants, and assess the arrival patterns for the attendees.

Based on the relative data volumes sent in each direction on the UDP channels, we can also estimate the proportion of participants that are using their camera, microphone, or screen-sharing during a Zoom call. Anecdotally, it has been observed that most students turn their cameras off when attending a Zoom lecture. Our measurements suggest that about 10% of participants are using their camera, and the others have the camera off for most or all of the session.

Figure 5(a) shows an example from a one-hour Zoom session with about 40 on-campus participants. The horizontal axis shows the average packet size sent by each connection, and the vertical axis shows the average packet size received. The points are colour-coded based on our (heuristic) classification of the channels. The points cluster quite tightly into logical groupings, with the video category varying the most. The graph shows that one of the Zoom participants was using screen-sharing, with an average packet size of 600 bytes, while the others were not. Similarly, one participant sent more audio/video than the others, and received more control information. These observations imply that there was a single presenter during the session, with several questions asked via audio or chat. These patterns are consistent with a lecture-based class.

Figure 5(b) shows an example of a large meeting with 120 Zoom participants, about 30 of whom were on campus (including one of the co-authors). The patterns for the UDP channels resemble those in Figure 5(a), though more cameras

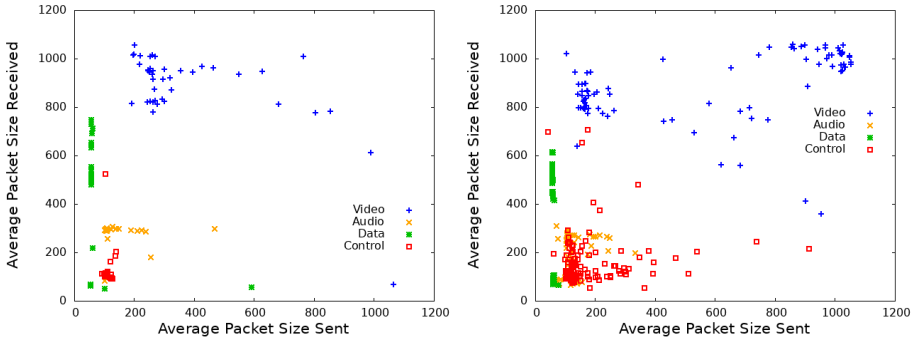

(a) Zoom 3: Lecture Class (b) Zoom 4: Large Meeting

Fig. 5: Scatterplot View of Zoom Channel Usage

(about 40%) were on during the meeting. In addition, the TCP control channel indicates lots of interaction via the chat interface for questions and answers.

5.3 Session Quality

Figure 6 shows the session profile[9] for the large Zoom meeting. There was a steep arrival pattern for the connections, because of the Zoom waiting room used to admit attendees. The consistency in the end times for the connections suggests that all participants were (likely) in the same meeting.

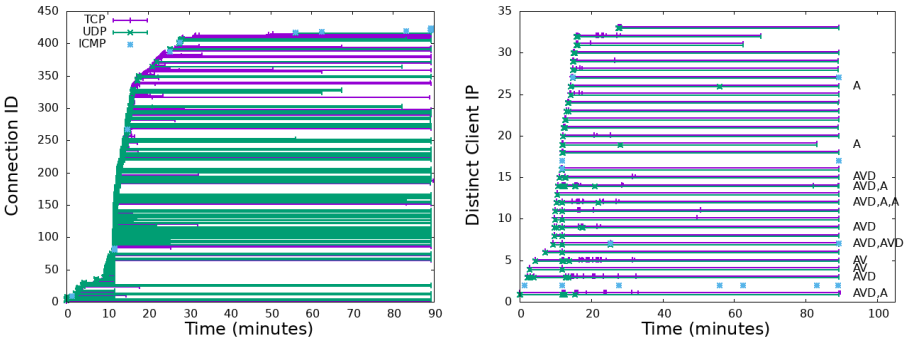

(a) Client Connections (b) Client IP Addresses

Fig. 6: Zoom 4: Session Profile for Participants in Large Zoom Meeting

Figure 6 shows that about a dozen (40%) of the 30 participants had disruptions to their TCP connections during the Zoom session. We have also annotated

[9] We exclude the port profile, which is too cluttered to be useful.

the graph with A, V, or D to indicate disruptions to the UDP-based audio, video, or data channels, respectively, that occurred during the meeting. About 10 of the participants (30%) had disruptions, with many of these participants on NAT addresses. Even users on static IP addresses (e.g., IPs on rows 7, 25, and 27) had some disruptions on their UDP channels.

5.4 Anomalous Zoom Behaviour

Figure 7 shows one of the most interesting Zoom sessions that we found in our data. This graph shows the IP profile for a session with about 20 on-campus participants. This session seems to be a lecture-oriented class that lasted about an hour. Several disruptions occurred during this session, at almost periodic intervals that are about 11 minutes apart. Based on the session profile, it appears that the session host had their network connectivity disrupted (note the ICMP messages), which affected every other meeting participant.

The timing gaps between UDP connections are about 2-10 seconds each, consistent with Zoom restarts. Even more interesting, Zoom seems to enter a "failover" mode with *four* concurrent UDP connections: three[10] new ones on a randomly chosen[11] Zoom server, and one on the original server. On the graph, this mode manifests in alternating fashion upon each new disruption. This might reflect Forward Error Correction (FEC) to preserve audio quality [10], since the extra UDP channel has packet rates and sizes consistent with audio.

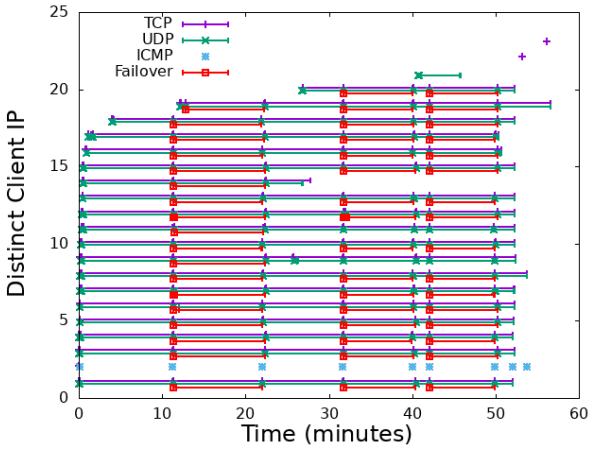

Fig. 7: Zoom 5: Example of Zoom Failover Mode

[10] There is also a TCP control connection to this new server (not shown on graph).

[11] Curiously, two clients (IPs 9 and 11) choose the *original* Zoom server to handle failover as well, inducing extra connection overhead at an inopportune time.

We have not yet identified the root cause of this anomaly, but have found several other examples like it in our data. Using scatterplot analysis on this example, we have determined that all participants have their cameras on, and screen-sharing is in use, though the presenter does not seem to be on campus. This might suggest that the problem is a home networking issue for the presenter, such as buffer overflow at a WiFi router, leading to repeated Zoom failures that affect all participants in a synchronized fashion. We are trying to recreate this phenomenon experimentally in our ongoing work.

6 Discussion

In this section, we highlight the main performance implications from our study, and offer several suggestions for network administrators and Zoom software engineers to improve Zoom performance in the future. We also discuss the technical limitations of our work.

6.1 Performance Implications and Recommendations

The main problem identified in our study is one of supply versus demand. The use of mixed learning modalities on our campus during the Fall 2021 semester quadrupled Zoom traffic demand compared to Fall 2020, while directing this traffic to regional Zoom servers in Canada has inadvertently constrained the supply of MMR servers available. As a result, these servers often host multiple Zoom meetings at the same time, some of which are large and long duration meetings. At busy times of the day, some MMR servers seem to become sluggish (e.g., TLS), compromising session quality and the user-level Zoom experience.

For campus network administrators, our main recommendation is to ensure that a sufficiently large pool of regional servers is available. If this is not possible, then a second option is to reduce the default video bit rate on Zoom sessions, so that traffic loads are lower. A third option is to host a Zoom Meeting Zone within the campus network, though this might be expensive, and would no longer be needed when the pandemic is over. Finally, it is important to consider the load that Zoom places on the campus WiFi network and the NAT infrastructure, to ensure that these are not performance bottlenecks for Zoom.

For Zoom network engineers, one recommendation is to ensure that their regional data centers are adequately resourced to handle peak traffic loads. Better monitoring of these facilities could also identify performance anomalies sooner. Another recommendation is to improve Zoom's load balancing, which currently seems to be random, rather than load-based. These policies lead to non-uniform loads on MMR servers (Figure 5(b)), and poor selections of failover servers when needed (Figure 8). A third recommendation is to consider the use of QUIC, rather than UDP, for media streaming, if they are not already doing so. QUIC could provide a simpler solution for robust media streaming, without the need for elaborate FEC and failover strategies. In the longer term, a final recommendation is to consider the use of network-layer or application-layer *multicast*, which seems like a potential solution to reduce network traffic and server load.

6.2 Limitations

There are several limitations to our study. First, our monitoring infrastructure only sees traffic that traverses the edge of the campus network, so we inherently underestimate Zoom traffic when some meeting participants are off-campus. Second, our heuristics for UDP channel classification (i.e., audio, video, data) are simple threshold-based policies, and may not be robust to the many variations possible in empirical Zoom traffic. Third, a similar caveat applies to our heuristics for Zoom meeting classification: overlapping meetings on the same Zoom server make meeting identification quite a challenge, and campus-level NAT means that the same client IP can be in multiple Zoom meetings (same or different) at the same time. Fourth, the restarts of our monitor, plus any packet losses there, result in underestimation of Zoom traffic volume. Finally, our traffic analysis approach currently misses Zoom traffic exchanged in P2P mode, since neither endpoint is a Zoom address. Identifying this traffic (at least when it crosses the campus edge network) is part of our ongoing work.

7 Conclusion

This paper has presented a detailed analysis of Zoom network traffic on a university campus network. Through small-scale experiments, we identified the structural properties of Zoom sessions. We then used the knowledge and insights from the small-scale study to analyze large-scale Zoom traffic at the campus level.

The main take-home messages from our paper are as follows. First, Zoom usage on our campus has increased 4x with the transition to a mix of in-person and online course delivery. Second, this traffic can stress campus network infrastructure, including WiFi and NAT, due to many concurrent meetings, temporally-correlated arrivals, high video bit rates, and long-lasting sessions. Third, many Zoom sessions experience disruptions, seemingly triggered by high demand on a limited supply of regional Zoom servers. These disruptions can degrade the QoE for Zoom users. Finally, our simple analysis and visualization tools provide an effective way to identify and quantify such problems.

Our ongoing work seeks to corroborate the results from our network-level view with the application-level view provided by the Zoom console (dashboard) used by our campus network administrators.

Acknowledgements

The authors thank the PAM 2022 reviewers and shepherd Matteo Varvello for their constructive suggestions that helped to improve our paper. Summer student Kiana Gardner helped with our active measurements, including the collection of Wireshark traces from Zoom test sessions. The authors are grateful to University of Calgary Information Technologies (UCIT) and the Conjoint Faculties Research Ethics Board (CFREB) for enabling the collection of our passive network traffic measurement data. Financial support for this work was provided by Canada's Natural Sciences and Engineering Research Council (NSERC).

Appendix 1: Data Format and Analysis Tools

Figure 8 shows an example of the connection log data from the Zoom test session in Figure 1. This format uses selected columns from the Zeek connection log [13]. Each line summarizes the network traffic on one connection (TCP, UDP, or ICMP). In this example, A.B.C.D is a laptop on the campus WiFi network, K.L.M.N is a laptop on a home network, and W.X.Y.Z is a Zoom MMR server. The number of users varied between 1 and 3, but the third user was off campus, and thus does not appear in the log. In this example, there were two UDP connection attempts before P2P mode was fully established. Also, an ICMP "port unreachable" message was sent when switching back to server mode.

Timestamp	Src_IP	SPort	Dest_IP	DPort	Prot	Duration	State	PSent	BytesSent	PRecd	BytesRecd
3371.758208	A.B.C.D	51525	W.X.Y.Z	443	tcp	2539.271654	RSTR	12214	1692606	21103	1646631
3372.166462	A.B.C.D	57643	W.X.Y.Z	8801	udp	361.544867	SF	3209	3096498	3008	2793039
3372.391270	A.B.C.D	57644	W.X.Y.Z	8801	udp	361.320242	SF	760	73731	194	20933
3372.515465	A.B.C.D	57645	W.X.Y.Z	8801	udp	361.196108	SF	311	33321	396	41924
3733.570248	K.L.M.N	38099	A.B.C.D	9756	udp	0.157993	SF	23	2622	21	2373
3733.592228	K.L.M.N	45276	A.B.C.D	15326	udp	666.015837	SF	211955	179935922	283443	252424232
4399.609065	A.B.C.D	57193	W.X.Y.Z	8801	udp	1511.283689	SF	7054	1394954	4778	344920
4399.609081	A.B.C.D	57192	W.X.Y.Z	8801	udp	1511.408204	SF	114638	25945097	95585	17170572
4399.609405	A.B.C.D	57194	W.X.Y.Z	8801	udp	1511.407976	SF	354039	388913210	289850	318527540
4399.612464	K.L.M.N	3	A.B.C.D	3	icmp	0.002069	OTH	8	4086	0	0

Fig. 8: Zeek Connection Log Entries for Zoom Test Session (anonymized)

We have written C and Python programs to parse such log entries and produce graphical visualizations of Zoom sessions using gnuplot. Our C programs (called `zoomparse.c`, `zoomplot.c`, and `zoomcount.c`) produce a textual summary, intermediate data for graph plotting, and a statistical summary of Zoom sessions, respectively. We also have a Python program that parses full Zeek connection log entries, and produces a summary of Zoom sessions and Zoom meetings. The latter program relies on a database of Zoom server IP addresses and server types. Our software tools and graph plotting scripts are available from `http://www.cpsc.ucalgary.ca/~carey/software.htm`

Appendix 2: Additional Zoom Test Sessions

We collected Wireshark traces of several other Zoom sessions in order to identify typical and atypical behaviours. Figure 9 shows two unusual examples.

Figure 9(a) shows the packet traffic for a small meeting with three participants (all on their home networks), during which the presenter used the screen-sharing function to scroll through a large PDF document. In this example, the screen-sharing data volume (green) fluctuated dramatically, and actually exceeded the video traffic volume for most of the session.

Figure 9(b) shows the video and audio packet traffic for an on-campus participant during a seven-person Zoom meeting. (We exclude the data and control

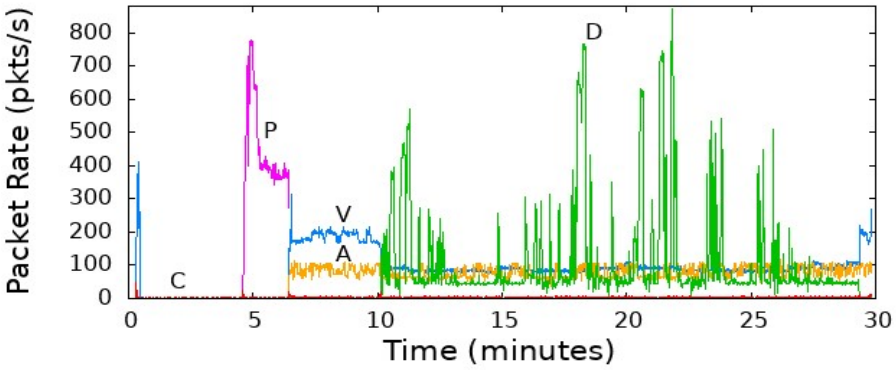

(a) Intensive Screen-Sharing Activity

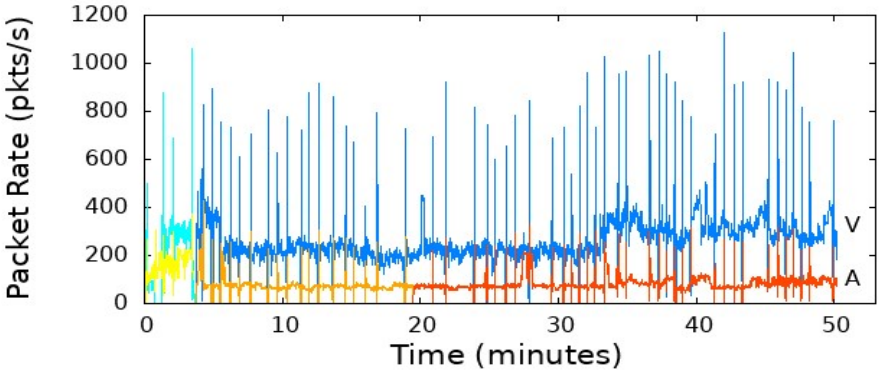

(b) Extreme Spikes in Zoom Traffic

Fig. 9: Additional Examples of Zoom Test Sessions in Wireshark

traffic from the graph, since it is negligible.) There are extreme spikes in the traffic during this Zoom session, which had very poor QoE (the different colors in the graph show the audio and video disruptions). One of the spikes, near the 20-minute mark, reflects Zoom's bandwidth probing, which lasts for 10 seconds. The other spikes, however, are more extreme, and seem almost periodic. Each spike in the Wireshark trace lasts for only a second or two, and is preceded by a 2-3 second interval with no packets at all. Furthermore, the same pattern occurs in both the audio and video traffic (as well as non-Zoom traffic in the trace).

We do not believe that the traffic spikes in Figure 9(b) are attributable to Zoom servers. Rather, this phenomenon could reflect congestion on the campus WiFi network (e.g., a large backlog at an AP), or could be a measurement artifact from running Wireshark on the same laptop as the Zoom session. We have observed this pattern in at least three different Wireshark traces, but have not yet been able to recreate it experimentally.

These examples help illustrate the variety of traffic patterns observed during our Zoom test sessions.

References

1. Böttger, T., Ibrahim, G., and Vallis, B.: How the Internet Reacted to Covid-19: A Perspective from Facebook's Edge Network, Proceedings of the ACM Internet Measurement Conference (IMC), pp. 34-41, Pittsburgh, PA, October 2020. https://doi.org/10.1145/3419394.3423621
2. Carman, A.: Why Zoom became so popular, The Verge. https://www.theverge.com/2020/4/3/21207053/ zoom-video-conferencing-security-privacy-risk-popularity (accessed September 19, 2021).
3. Chang, H., Varvello, M., Hao, F., and Mukherjee, S.: Can You See Me Now? A Measurement Study of Zoom, Webex, and Meet, Proceedings of ACM IMC, pp. 216-228, November 2021. https://doi.org/10.1145/3487552.3487847
4. Favale, T., Soro, F., Trevisan, M., Drago, I., and Mellia, M.: Campus Traffic and e-Learning during COVID-19 Pandemic, Computer Networks, Vol. 176, Article 107290, pp. 1-9, July 2020. https://doi.org/10.1016/j.comnet.2020.107290
5. Feldmann, A., Gasser, O., Lichtblau, F., Pujol, E., Poese, I., Dietzel, C. et al.: The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic, Proceedings of ACM IMC, pp. 1-18, Pittsburgh, PA, October 2020. https://doi.org/10.1145/3419394.3423658
6. Karamollahi, M., Williamson, C., and Arlitt, M.: Zoomiversity: A Case Study of Pandemic Effects on Post-Secondary Teaching and Learning, to appear, Proceedings of Passive and Active Measurement (PAM) Conference, online, March 2022.
7. Liu, S., Schmitt, P., Bronzino, F., and Feamster, N.: Characterizing Service Provider Response to the COVID-19 Pandemic in the United States, Proceedings of PAM Conference, pp. 20-38, Germany, March 2021. https://doi.org/10.1007/978-3-030-72582-2_2
8. Lu, Y., Zhao, Y., Kuipers, F., and Mieghem, P.: Measurement Study of Multi-Party Video Conferencing, Proceedings of IFIP Networking Conference, pp. 96-108, Chennai, India, May 2010. https://doi.org/10.1007/978-3-642-12963-6_8
9. Lutu, A., Perino, D., Bagnulo, M., Frias-Martinez, E., and Khangosstar, J.: A Characterization of the COVID-19 Pandemic Impact on a Mobile Network Operator Traffic, Proceedings of ACM IMC, pp. 19-33, Pittsburgh, PA, October 2020. https://doi.org/10.1145/3419394.3423655
10. MacMillan, K., Mangla, T., Saxon, J., and Feamster, N.: Measuring the Performance and Network Utilization of Popular Video Conferencing Applications, Proceedings of ACM IMC, pp. 229-244, November 2021. https://doi.org/10.1145/3487552.3487842
11. Mahr, A., Cichon, M., Mateo, S., Grajeda, C., and Baggili, I.: Zooming into the Pandemic! A Forensic Analysis of the Zoom Application, Forensic Science International: Digital Investigation, Vol. 36, Article 301107, p. 7, March 2021. https://doi.org/10.1016/j.fsidi.2021.301107
12. Marczak, B., and Scott-Railton, J.: Move Fast and Roll Your Own Crypto: A Quick Look at the Confidentiality of Zoom Meetings, Citizen Lab Research Report No. 126, University of Toronto, April 2020 (accessed September 23, 2021).
13. Paxson, V.: Bro: A System for Detecting Network Intruders in Real-time, Computer Networks, Vol. 31, No. 23, pp. 2435-2463, December 1999. https://doi.org/10.1016/S1389-1286(99)00112-7
14. Sander, C., Kunze, I., Wehrle, K., and Rüth, J.: Video Conferencing and Flow-Rate Fairness: A First Look at Zoom and the Impact of Flow-Queuing AQM, Proceedings of PAM Conference, pp. 3-19, Germany, March 2021. https://doi.org/10.1007/978-3-030-72582-2_1

15. Ukani, A., Mirian, A., and Snoeren, A.: Locked-In during Lock-Down: Undergraduate Life on the Internet in a Pandemic, Proceedings of ACM IMC, pp. 480-486, November 2021. https://doi.org/10.1145/3487552.3487828
16. Wireshark.org, Wireshark Frequently Asked Questions, https://www.wireshark.org/faq.html (accessed September 22, 2021).
17. Zoom, Zoom Video Conferencing Plans and Pricing, https://zoom.us/pricing (accessed September 19, 2021).