Audio Streaming Application Performance: A Comparative Study of Spotify and YouTube Music

Lachlan Moore
Supervisor: Carey Williamson
lachlan.moore@ucalgary.ca

Abstract

Music streaming is an increasingly popular Internet service. The music streaming market is dominated by two main companies: Spotify, and YouTube Music. In this study, we characterise Spotify and YouTube Music's network performance at both a micro and macroscale. Our microscale observations capture specific packet level data, while the macroscale data is taken from our campus network traffic. The results of this study show heavy traffic usage of both music streaming services, which show diurnal patterns and heavy tailed behaviour. YouTube Music and Spotify differ in their network infrastructures, protocol, usage, and popularity.

1 Introduction

Every year there is an increasingly large market share for music streaming services. With the ongoing COVID-19 pandemic, the demand for services to be online has only further increased this popularity. These streaming services account for 83% of all revenue generated by the music industry [5]. Due to the massive market share for music streaming, there is also an increasing number of companies providing this service, including Spotify and YouTube Music.

Spotify is the world's largest and most popular music streaming software with 365 million active users [9]. Spotify offers free access to 70 million unique tracks and 2.9 million podcasts with advertisements, while some additional features such as commercial-free listening are offered via subscription. They also offer features such as public/collaborative playlists, music stories, and friend activity lists. According to a paper in 2018, [3] an average listener spends about 25 plus hours of music listening per month.

YouTube Music, a relatively new service, is the fastest growing music streaming service [1]. YouTube Music is the successor of Google Play Music and has 60 million tracks. YouTube Music offers similar services to Spotify, including free accounts with advertisements, with the main differences being access to watching music videos and listening to radio stations.

The main focus of this study is the comparison of Spotify and YouTube Music's usage at a network level. We specifically look deeper into Spotifys' overall network performance. Both services are compared in the context of our campus of over 33,000 students, faculty, and staff.

CPSC 503.08, December 2021, Calgary, AB, Canada

Our study involves the collection and analysis of data at both a micro and macroscale. In the microscale, data is collected at the packet level from both music streaming services, Spotify and YouTube Music, in a controlled environment. By analysing the micro data, we then narrow down the overall network data stream, at the macroscale, to just that of Spotify. This in turn allows us to provide a full campus overview of Spotify's network performance, as well as comparison of the two at a packet level.

This study seeks to answer the following questions:

- How much campus network traffic is generated by these two music streaming services?
- How are these two services similar or different in their architecture, traffic patterns, and protocol usage?
- Are there differences in the network-level or user-level performance for these two applications?

The rest of this paper is organised into five additional sections. Section 2 provides background information and prior related work on music streaming apps. Section 3 describes the passive and active methodologies for collection and analysis of data. Section 4 presents the microscale measurement results, while Section 5 presents the macroscale measurement results. Finally, Section 6 concludes the paper.

2 Previous Work

There are previous studies on video based media streaming, although few study audio based media streaming. While Spotify has been extensively researched, most contemporary research fails to compare it with the fastest growing music streaming site, YouTube Music. Because of this, a comparison of the two services, from a network standpoint is novel.

Sackl *et al.* [10] investigated the impact of temporal impairments between video and audio streaming services. By conducting a subjective user study they found that music streaming users are less tolerant to initial delays and stalling.

Schwind *et al.* [11] conducted an QoE analysis (active measurement) on Spotify's mobile application in order to find network related properties. They found that Spotify not only buffers an entire song, but also prefetches subsequent songs in a playlist to avoid possible network delays.

Basher *et al.* [2] compares peer-to-peer (P2P) and Web traffic from a large scale network to discuss the implication of increased P2P traffic. They suggest that new models are necessary for Internet traffic and provide flow-level distributional models for Web and P2P traffic.

Laterman *et al.* [7] characterised two video streaming services, Netflix and Twitch, at the campus level. By studying traffic on a campus edge router they were able to see the similar traffic characteristics of the two video services and how their network performance might be improved.

Kreitz *et al.* [6] gives an overview of the protocol and peer-to-peer (P2P) architecture used in Spotify and its performance. They found that 8.8% of music data came from Spotify's servers, another 35.8% of data coming from the P2P network, and the remaining percentage from cached data. Because Spotify no longer uses a P2P architecture [4], our study deviates from Kreitz's study.

3 Methodology

This section describes the active and passive measurement techniques used to collect and analyse music streaming services. In this study, we collected and analysed network traffic from the University of Calgary campus network. The collected data was analysed in controlled environments for the microscale, while the macroscale was from overall campus traffic. At the microscale, we used Wireshark to measure and analyse the packet data. Utilizing this gathered packet data, we were able to greatly narrow down the campus network traffic to specific servers to analyse.

3.1 Active Measurement

Active measurement involves the injection of network traffic for the purpose of measurement, and is used to understand performance. Small-scale active measurement was used to reveal important information about how the streaming service operates over the Internet. This process identifies specific packet level information such as domain names, IP addresses, protocols, port numbers, and TCP/UDP connections.

To capture our packet data, we used Wireshark. By capturing packet level information, we are able to validate the IP addresses. This in turn allows us to confirm that these IP addresses are consistent between each of our tests with Spotify and YouTube Music. Our methodology was exclusively tested in Spotify and YouTube Music's Google Chrome Web applications.

Prior to running our tests, many other extraneous applications were disabled or their Internet access was terminated. Each experiment was performed in various locations on campus. To maintain consistency between experiments, Web applications were used and loaded from a fresh browser session, and the same song was played during each test. From Spotify or YouTube Music's main page, the song 4:00 A.M. [8] by Taeko Onuki was searched and played for its full duration. These tests each spanned around around six minutes and captured data for the entirety of that duration.

Each test was run from the campus wireless network at various times throughout the afternoon, between 11:00 - 15:30. The campus AirUC-guest network, eduroam, as well as the edu secure networks were all used in testing. Locations of the tests varied throughout campus and were conducted in locations such as MacEwan Hall, Math Science Laboratories, Craigie Hall, and Science B. All active measurements were done on the same single laptop device, an HP Spectre x360 13-aw2020ca (Quad-Core 4.7 GHz CPU, 16 GB RAM), running a Windows 10 operating system.

Table 1 describes the location, time, date and network used for each test. Results from each location and time were consistent among tests, but with more packets being lost at busier locations as well as at busier times of the day.

Table 1. Metadata for Active Measurement Experiments

Streaming Service	Location	Time	Date	Network
YouTube Music	Math Science Laboratories	15:03	2021-09-24	Edu secure
	MacEwan Hall	11:31	2021-10-05	Eduroam
	Science B	13:15	2021-10-13	AirUC
	Craigie Hall	14:15	2021-10-15	Edu secure
Spotify	Math Science Laboratories	15:24	2021-09-24	Edu secure
	MacEwan Hall	11:40	2021-10-05	Eduroam
	Science B	13:22	2021-10-13	AirUC
	Craigie Hall	14:24	2021-10-15	Edu secure

3.2 Passive Measurement

Passive measurement involves observing the overall traffic flow from one or more network vantage points, without injecting any additional traffic. We passively collected network data from a campus edge router connected directly to the Internet. Our campus network is used by 33,000 students, faculty, and staff.

To capture network traffic, we used a capture card on the campus's edge router. Our data collection mechanism records connection-level information and not packet-level payloads. The captured traffic is stored as Bro connection logs which contain timestamps, IP addresses and ports of the source and destination, connection duration, connection state, as well as the number of packets and bytes sent and received.

To analyse the captured data, we used various Bash and Python scripts. These pulled specific data from the Bro logs for analysis and observation.

4 Microscale Measurement Results

In this section, we present our comparison between Spotify and YouTube Music at the packet level.

4.1 Service Infrastructure

The first result that our tests with Wireshark were able to provide was the server layout of both Spotify and YouTube Music. Each of the services is using various addresses, with Spotify having around 15 servers and YouTube Music having

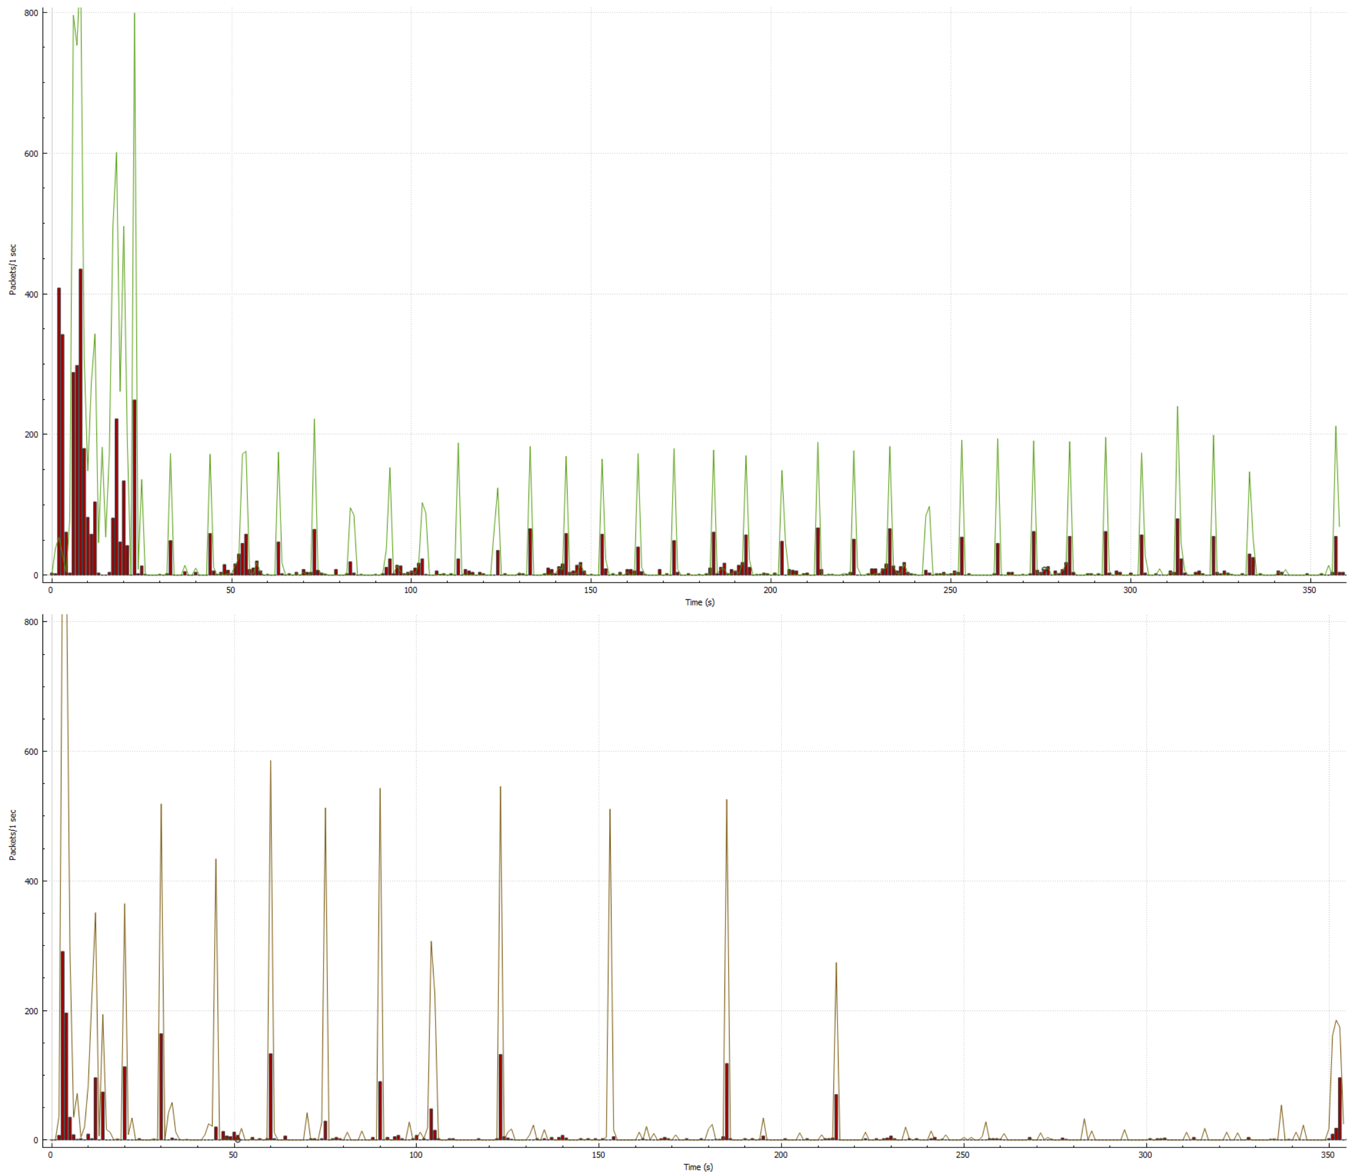

Figure 1. Time series graphs of Wireshark test sessions for Spotify (Top) and YouTube Music (Bottom) (Oct 13, 2021)

around 30. Spotify was found to be using content delivery networks (CDN) belonging to Akamai and Fastly. The Spotify addresses belonged to the 35.184.0.0/13, and 34.64.0.0/10 network, Akamai belonged to the 23.192.0.0/11 network, and Fastly belonged to the 151.101.0.0/16 networks. YouTube Music shared various hosts all belonging to Google's infrastructure. These belonged to various hosts on the 142.250.0.0/15, 172.217.0.0/16, and 199.212.24.0/24 networks.

Spotify and YouTube Music both primarily use secure TCP connections for data transfer. Spotify produced HTTP range requests during Wireshark testing while YouTube Music used standard HTTP. On the other hand, YouTube Music also used QUIC, presumably Google's version, in addition to TCP connections.

Observations of Spotify and YouTube Music were made at a subjective level while using each service in testing. When listening to the same song in the same location with the same listening devices, YouTube Music seemed to have a higher audio quality than Spotify. Although while moving or in locations where Internet connection is variable, Spotify provides less pauses in music.

Table 3 shows the distribution of packet sizes for Spotify and YouTube Music tests on October 13, 2021. We can see that Spotify is sending twice as many packets than YouTube Music. The packet size distribution is bimodal, with many large data packets, and many small control packets (ACKs).

Figure 1 shows the time series graphs of our Wireshark tests from October 13th, 2021. The vertical axis is the number

Table 2. Main similarities and differences observed between Spotify and YouTube Music

Characteristic	Similarities	Differences
Service	Both are music streaming services.	YouTube Music also offers access to music videos.
Buffering	Both services use a burst buffering at the same rate.	YouTube Music uses fewer (but larger) bursts than Spotify.
Protocols	Both use secure HTTP periodic connections and TCP.	YouTube Music also uses QUIC, and Spotify uses HTTP range requests.
Infrastructure	Both services make use of multiple servers.	Spotify uses CDNs while YouTube Music uses its own servers.
Servers	A small set of servers handle most of the workload.	YouTube Music has more overall servers than Spotify.
Traffic Volume	Both services have high volume traffic and continue to grow.	Spotify send twice as many packets as YouTube Music.
Users	Both services provide music to large number of users.	Spotify has significantly more active users than YouTube Music.

Table 3. Packet size distributions of Spotify and YouTube Music (Oct 13, 2021)

Packet Lengths	0-39	40-79	80-159	160-319	320-639	640-1279	1280-2559	Total	Average
YouTube Music	0	4190	720	173	178	181	5984	11426	806.07
Spotify	0	8006	1488	353	441	364	9798	20450	751.12

of packets sent, the horizontal axis is the time of the Wireshark trace in seconds, the red boxes represent the number of unacknowledged packets. The top graph shows the Spotify session, and the bottom graph shows the YouTube Music session.

We can see that both Spotify and YouTube Music are sending packets in bursts. These bursts are spaced quite similarly, about ten seconds between each burst, for each streaming service. Spotify has consistently spaced bursts of about 200 packets each throughout the whole session, while YouTube Music's bursts are about double in size and are done less than half as often as Spotify. These different rates of data can be explained for various reasons such as, different bit rates, media quality levels, song buffering, etc.

Finally, Figure 1 shows high initial rates of buffering in both graphs. Spotify shows significantly more initial buffering compared to YouTube Music, with 3 large spikes of packets to YouTube Music's single packet spike. This is consistent with the overall higher rate of buffering that Spotify is doing in comparison to YouTube Music.

4.2 Comparison

Using the packet-level information collected from Wireshark, we can characterise the similarities and differences between Spotify and YouTube Music. Table 2 displays the main characteristics observed, with those being service type, buffering approach, protocols used, infrastructure servers, traffic volume, and type of users. For example, in the buffering approach we see a very similar approach in timing, but different amounts of packets being sent.

To map the IP addresses of Spotify and YouTube Music, we used DNS logs to find the IP addresses associated with the host names. Additionally, by comparing multiple experiments we also observed common high traffic IP addresses being used. Utilizing both of these methods we were able to obtain a distinct list of the main IP addresses used by each music streaming service. When conducting the same active measurement methods on other applications from the same service provider, for example, YouTube Music and YouTube are both owned by Google, we found overlapping IP addresses. In our case, we found that both YouTube and YouTube Music used heavy connections from Google Video IP addresses (199.212.24.0/24). These findings created a problem in distinguishing traffic between YouTube Music and YouTube at a network level.

For the remainder of this study, we will focus on characterising Spotify's network performance at the campus level.

5 Macroscale Measurement Results

In this section, we present our results of Spotify's network performance at the campus level.

5.1 Traffic Profile

Figure 2 shows the number of connections initiated to Spotify in each hour of the day on our campus over a 24 hour period, on Wednesday September 22nd, 2021. The purple line refers to the amount of connections that were made to Spotify servers. The vertical axis shows the number of connections. The horizontal axis shows the time of day in hours.

We can observe a clear daytime active usage pattern for Spotify, as expected. Network traffic changes related to the time of day, as a result of common human usage of a network. This graph shows us a steep growth in connections starting in the morning, a peak around the early afternoon, and a slightly slower decline in the late evening.

Table 4 shows how many connections, bytes sent and received, clients, and servers there were throughout the day of September 22nd. We are observing just under one million connections, 300 GBs of data, 16000 clients and fifteen different servers.

Another observation is that there is an increasing amount of connections throughout the early afternoon of the day, despite fewer clients compared to other parts of the day. This suggests that due to the time of day many students are moving between lectures, lunch and other meetings and hence dropping and re-establishing connections to the campus network. Overall this behaviour would result in more connections to Spotify servers, explaining this anomaly. In our Bro logs the median length of connections shortens from 184.4 seconds to 120.7 seconds between 10:00 - 11:00.

Figure 3 shows the frequency of connections and byte volume for servers and clients of Spotify on September 22nd,

Table 4. Connections, bytes, clients and servers for Spotify (Sept 22, 2021)

	Connections	Bytes Sent	Bytes Received	Clients	Servers
Spotify	935 920	81.0 GB	221.4 GB	16 784	15

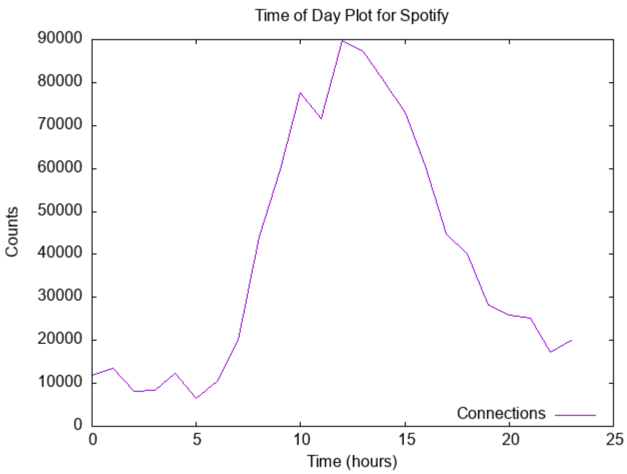

Figure 2. Traffic profile for Spotify (Sept 22, 2021)

2021. The top graph shows the frequency rank of connections, while the bottom graph shows the frequency rank of byte volume. Both graphs use a logarithmic scale on the vertical and horizontal axis.

Both graphs in Figure 3 display the non-uniformity of traffic across the IP addresses. We see in the upper graph that about five servers, three Spotify and two Akamai, account for 95% of the connections made throughout the day. Additionally three of these five servers, one Spotify and the two Akamai, also make up for 97.4% of bytes sent. The two byte heavy servers belong to the Akamai CDN and make up for 48.1% of the bytes being sent, while only making up for 14.2% of connections. The remaining 49.3% of bytes being sent belong to a single Spotify server. We can observe from this that Spotify contains servers to handle data transfer, mainly being the Akamai CDN, and servers to handle connection management.

With the client-side of the graph we also see highly non-uniform traffic in terms of connections and bytes received. The top few client IP addresses each account for over 10,000 TCP connections per day, and over 1 GB of data traffic volume per day. At around the 300th ranked client IP address, we see a significant drop in the connection count and the number of bytes, although the drop in connections is steeper in comparison to the byte count. This is mainly due to how our campus network is set up using Network Address Translation (NAT). On the AirUC 136.159.213.0/24 subnet, there are 255 IP addresses given to multiple devices. This means that one IP address has multiple users; students, staff and/or faculty, on a single IP address. As there are more clients on one IP address, we also see a larger use of Spotify among those multiple clients. The drop off at around the 300th user is due to static IP addresses. These static addresses belong to various public and staff desktop computers around campus.

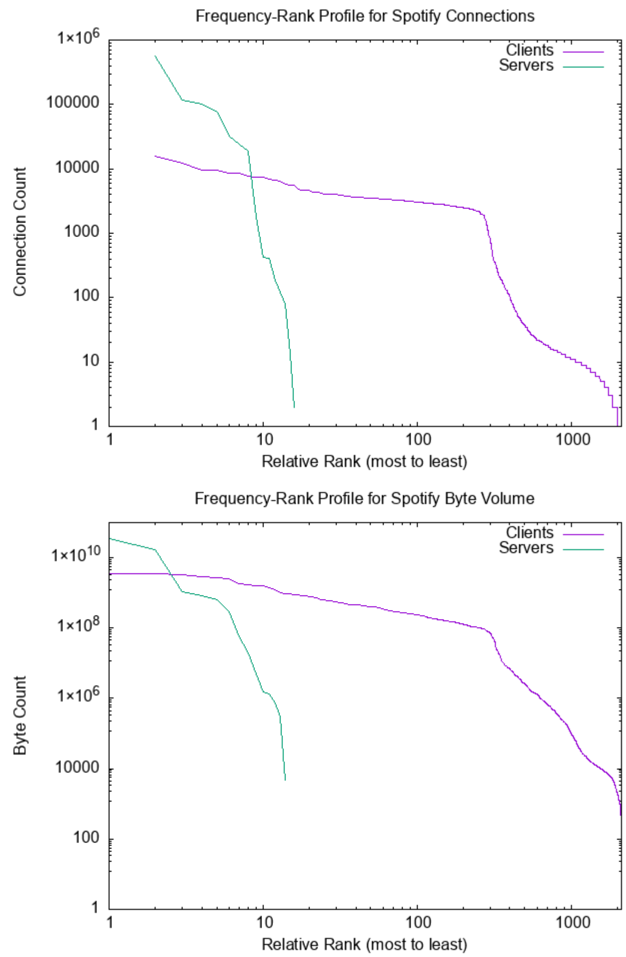

Figure 3. Frequency-rank profiles for Spotify connections and byte volumes (Sept 22, 2021)

We also see a large difference in the scale between the number of connections and bytes. Some users and servers create a large number of connections and bytes, while some have very few connections and byte transfers. The mean transfer size between Spotify and a user was 118295 bytes, while the median was 3034 bytes. While most connections are small, the large connections are way bigger and generate the majority of the data. We may infer from this that there is a distinct difference in the types of users, power users and regular users, using Spotify and the types of servers within Spotify's infrastructure.

5.2 Durations and Transfer Sizes

Next, we analyse the durations and transfer sizes of the TCP connections created by Spotify sessions. We are specifically focusing on the tails of these distributions, looking for evidence of heavy-tailed behaviours.

Figure 4 provides a look into the tail behaviours of observed TCP connections, using log-log complementary distribution (LLCD) plots. The leftmost graph is for the connection sizes in bytes, while the right graph is for the connection lengths in seconds. Both graphs show graphical evidence of heavy-tailed distribution. This is an unexpected behaviour as a music streaming service, Spotify is not expected to exhibit traits of applications that transfer large amounts of data.

A final observation that can be made from Figure 4 is the difference between the sent and received data on the leftmost graph. The received data plot is consistently above the amount of sent data. This observation reinforces previously seen data in Table 4, that more information is being sent to our campus than is being received. This is expected as the population of Spotify users on campus consume Spotify content rather than create content.

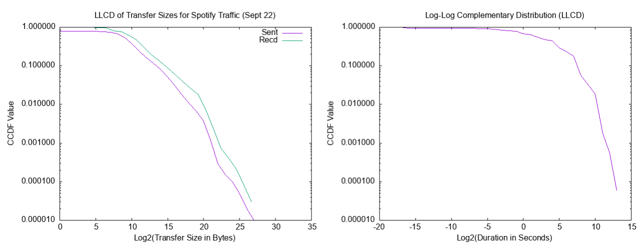

Figure 4. LLCD Plots of Empirical Distributions: Transfer Sizes (left) Connection Duration (right)

5.3 Servers

Next, we analyse the specific transfer sizes of the TCP connections sent by the five most connection-heavy Spotify servers. Once again we specifically focus on the tails of these distributions, looking for evidence of heavy-tailed behaviour.

Figure 5 shows the tail behaviour for the Spotify servers with the five highest connections, using LLCD plots. IP addresses 23.220.167.17 and 23.220.167.66 both belong to the Akamai CDN, while the other three are Spotify addresses. All five servers show evidence of heavy-tailed behaviours, with both Akamai addresses exhibiting these traits more than the Spotify addresses. This is consistent with the overview of Spotify's behaviour. From this we can infer how Spotify uses the Akamai CDNs. With the larger transfer sizes coming from Akamai's servers, large, popular, static media would mainly be stored in the Akamai CDN. Because smaller connections come from Spotify, most of the smaller Web interface connections are handled within Spotify's own servers.

5.4 Performance Implications

Observing about 300GB of daily traffic volume for one music streaming service was higher than expected. Considering

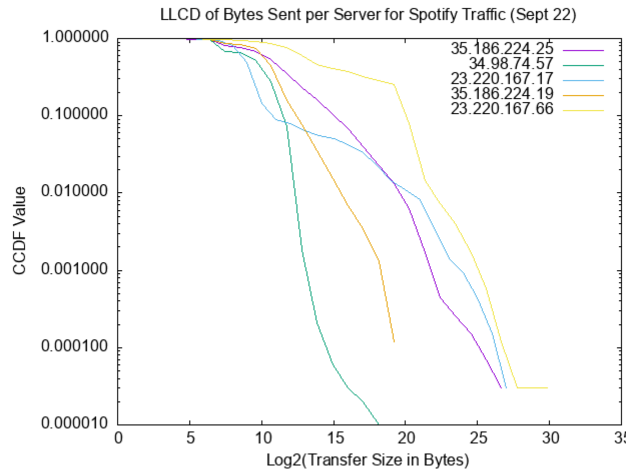

Figure 5. LLCD Plot of Five Spotify Addresses with Most Connections

this traffic and how it may grow in the future once students fully return to campus, is important to understanding and improving our campus network. We observed that one of the many music streaming services produces a significant amount of traffic on the campus network. This traffic will undoubtedly grow as more music streaming services start to provide other features such as video sharing, podcasts, audiobooks and more.

These results highlight the importance of CDNs when it comes to the delivery of network data. From our study, we see the use of the Akamai CDNs by Spotify to deliver their content. CDNs allow for the quick delivery of content to nearby users, which is important given the rise of popularity of new streaming services.

6 Conclusion

In this study, we used passive and active measurement to analyse one entire day of music streaming traffic on our campus network. This data was used to identify the similarities and differences between Spotify and YouTube Music at a packet level as well as characterise Spotify at a campus level.

Our study has presented several interesting findings. The number of connections created by Spotify is large and has about 300 GB of data sent and received between users and servers per day. We see a diurnal usage pattern for Spotify as usage of the network changes throughout the day. Spotify also had a heavy tail distribution for connection duration and transfer sizes.

Music streaming services are still gaining traction, as more variety is created and more features are being added for users. The traffic characterisation of Spotify is bound to change as music streaming services evolve. This study can help future studies of music streaming services to characterise their network performance. As the COVID-19 pandemic begins to end, it will be interesting to see how network traffic changes and adapts to the influx of users once again.

References

[1] MIDiA Research 2021. 2021. Global music subscriber market shares Q1 2021. Retrieved September 13, 2021 from https://www.midiaresearch.com/blog/global-music-subscriber-market-shares-q1-2021

[2] N. Basher, Aniket Mahanti, Anirban Mahanti, C. Williamson, and M. Arlitt. 2008. A Comparative Analysis of Web and Peer-to-Peer Traffic. *Proceedings of WWW* (April 2008), 287–296.

[3] Goodwater Capital. 2018. Understanding Spotify: Making Music Through Innovation. (March 2018). https://www.goodwatercap.com/thesis/understanding-spotify#important-disclosures

[4] Ernesto Van der Sar. 2014. Spotify Starts Shutting Down Its Massive P2P Network. Retrieved September 16, 2021 from https://torrentfreak.com/spotify-starts-shutting-down-its-massive-p2p-network-140416/

[5] J. Friedlander. 2021. YEAR-END 2020 RIAA REVENUE STATISTICS. Retrieved September 15, 2021 from https://www.riaa.com/reports/2020-year-end-music-industry-revenue-report/

[6] G. Kreitz and F. Niemela. 2010. Spotify – Large Scale, Low Latency, P2P Music-on-Demand Streaming. *IEEE 10th Intl. Conference on Peer-toPeer Computing* (Aug. 2010), 1–10.

[7] M. Laterman, M. Arlitt, and C. Williamson. 2017. A Campus-Level View of Netflix and Twitch: Characterization and Performance Implications. *Proceedings of SCS SPECTS* (2017).

[8] Taeko Onuki. 2021. *4:00 A.M.* Retrieved November 22, 2021 from https://www.youtube.com/watch?v=4tEXaW8tPso

[9] Spotify Technology S.A. 2021. Spotify Company Info. Retrieved September 13, 2021 from https://newsroom.spotify.com/company-info/

[10] A. Sackl, S. Egger, and R. Schatz. 2013. Where's the Music? Comparing the QoE Impact of Temporal Impairments between Music and Video Streaming. *Proceedings of Fifth International Workshop on Quality of Multimedia Experience (QoMEX)* (2013), 64–69.

[11] A. Schwind, L. Haberzettl, F. Wamser, and T. Hoßfeld. 2019. QoE Analysis of Spotify Audio Streaming and App Browsing. *Proceedings of 4th Workshop on QoE-based Analysis and Management of Data Communication Networks* (Oct. 2019), 25–30.