Locked-In during Lock-Down: Undergraduate Life on the Internet in a Pandemic

Alisha Ukani
UC San Diego
San Diego, California, USA
aukani@ucsd.edu

Ariana Mirian
UC San Diego
San Diego, California, USA
amirian@eng.ucsd.edu

Alex C. Snoeren
UC San Diego
San Diego, California, USA
snoeren@cs.ucsd.edu

ABSTRACT

Governments around the world enacted stay-at-home orders in response to the COVID-19 pandemic, which changed many aspects of life, including how people interacted with the Internet. These draconian restrictions on in-person social interactions were perhaps most acutely felt by people living alone. We study the changes in network traffic of one such population, students remaining in the (single-occupancy) on-campus dormitories at a large residential educational institution during the onset and initial few months of the lock-down. Specifically, we analyze how students shifted their online work and leisure behaviors at an application level. Further, we segment the population into domestic and international students, and find that even within these two broad sub-populations, there are significant differences in Internet-based behavior. Our work provides a focused lens on pandemic Internet usage, examining both 1) a concentrated user population and 2) the differing impacts of a global pandemic on disparate sub-populations.

CCS CONCEPTS

• **Networks** → **Network measurement**; **Public Internet**; • **Information systems** → **Web applications**.

ACM Reference Format:

Alisha Ukani, Ariana Mirian, and Alex C. Snoeren. 2021. Locked-In during Lock-Down: Undergraduate Life on the Internet in a Pandemic. In *ACM Internet Measurement Conference (IMC '21), November 2–4, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3487552.3487828

1 INTRODUCTION

The COVID-19 pandemic disrupted many aspects of life across the globe, but perhaps the most significant impact was felt by those under government-mandated stay-at-home orders. While enforcement varied widely across jurisdictions, many impacted populations were forced to spend almost all of their time in their residence, with little-to-no in-person social interactions. While such circumstances have occurred in the past, rarely have they applied to such privileged populations, with so little notice, and for such extended periods of time. Moreover, for those individuals living alone in Internet-connected homes, the pandemic presents a unique opportunity to study the (changing) role that the Internet—and various application classes in particular—plays when it provides the main form of outside interaction in their lives.

We present the first measurement study of a pandemic-induced natural experiment: "How does university students' Internet usage change when they are forced to remain in their individual residences with almost no in-person contact?" In particular, we report on the network usage of several thousand mostly undergraduate students at UC San Diego, a major residential research university over a four-month period before, during, and after the onset of COVID-19. These students were in a very real sense "trapped" in their dorm room for the duration of the academic term and ostensibly had nowhere to go, as the university asked all residents to return home shortly after the World Health Organization (WHO) declared a pandemic (and international air travel ground to a halt), but before the issuance of a region-wide stay-at-home-order.

The Internet was likely the dominant—if not exclusive—means of interaction for the vast majority of these students. While campus policies regarding student movement varied during the four months of our study, congregating in groups larger than three was forbidden once the lock-down began, as were on-campus visitors. Though the university provides cable TV service, and many if not all students have cell phones (there are no land-line phones in the dormitories), TV is a one-way medium, and most modern cell phones preferentially operate on WiFi when available. We assume that students' phone-based Internet activity is captured by our network tap, although we cannot exclude that some students may have used their own, self-paid cellular data plans (or other, more fringe technologies) for portions of their connectivity.

Within this captive audience, we analyze Internet usage both in aggregate and across application classes. In particular, we are able to understand how work and leisure changed for many of these students at an application level, and provide a finer granularity of analysis than previous work. Classes at the university transitioned to an online modality shortly after the beginning of the pandemic, so the students' in-class engagement is well captured by their use of UC San Diego's favored tool, Zoom. Recreationally, we are able to segregate traffic by device type, and identify a large number of IoT devices and game consoles. Moreover, by geolocating the sites they visit, we are able to segregate the student population into (presumed) domestic and international students, and report upon differences in their behaviors.

Overall, we find several distinct trends. While aggregate network traffic initially spikes in April near the beginning of the pandemic, it returns to near pre-pandemic levels by May, although not for all applications. Moreover, we find that international students spend less time on US-based social media applications than their domestic

IMC '21, November 2–4, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-9129-0/21/11.
https://doi.org/10.1145/3487552.3487828

counterparts, but spend more time on Steam, which highlights the importance of separating sub-populations within analyses. In sum, our study adds a unique additional perspective to the growing corpus of COVID-19 Internet measurements.

2 RELATED WORK

A great deal has been published about Internet traffic during the COVID-19 pandemic, but generally from the perspective of large network or service operators. Their network-wide, or, conversely, service-only vantage points prevent them from painting a coherent picture regarding the network activities of any particular user group. Perhaps the most comprehensive published report uses ISP and IXP data to study how the volume of different traffic classes changed throughout the pandemic [7]. Some of their overall findings—such as the convergence of diurnal patterns to that of pre-pandemic weekends—are not apparent in our population, while others—like the increase in gaming and video conferencing—are clearly visible. We frequently present our results in a similar style, and, where relevant and possible, report over the same time periods (e.g., we report traffic volume for the same four weeks, substituting the week of May 14, 2020 for June 18, 2020 to stay within our academic term).

Others have also reported on network traffic at educational institutions, notably Favale et al. [6], who focus on the prevalence of remote-work tools such as VPNs, Microsoft Teams, and their university's custom teaching software. However, their study reports on the traffic of the entire university population across multiple campuses, not just residents; staff and faculty presumably were not living in their offices during the pandemic. While we find a similar dramatic uptick in educational technology tools, we are able to report on the recreational aspect of students' activities as well.

Network and service operators have published both the impacts of the pandemic on their services and their reactions to them [3, 10, 12]. Indeed, most reports—both in the peer-reviewed literature and the blogosphere—suggest the network itself held up well under the changing demand. For example, a study of Italian network operators [4] found that they were proactive in expanding network capacity, primarily through an increase in hardware resources, in order to match increased demand. A study examining the response of ISP providers in the United States affirms this increase in traffic, but notes that network utilization returns to pre-pandemic levels in part due to ISP capacity augmentation [11]. While we similarly find a marked increased in per-capita demand among students after the campus shutdown, the sharp decline in residential user population led to a decrease in overall network traffic.

The use of social media during COVID-19 has also been studied in an analysis of YouTube video sharing on Twitter [13]. This analysis finds a high negative correlation between media sharing and mobility; we similarly find that our "trapped" users greatly increase their social media and overall Internet usage.

3 DATA SOURCES

Our study focuses on the devices used by the residential campus population (primarily undergraduates) at UC San Diego between February 1, 2020 to May 31, 2020. A range of circumstances, notably the quick onset of pandemic travel restrictions and lock-downs, left many undergraduates trapped on campus where they were required to isolate in their on-campus residences. This situation provides a unique context for analyzing Internet-traffic usage which we were able to investigate using a passive monitoring infrastructure already in place at UC San Diego [5].

There are three primary components that serve as inputs for our data set: 1) raw bidirectional network traffic from the campus network, 2) DHCP logs, and 3) DNS resolutions. The raw network traffic is mirrored from a switch connecting the on-campus residences to the campus backbone, but specifically excludes traffic from certain networks due to high traffic volumes; excluded networks include parts of UC San Diego, Google Cloud, Amazon, Microsoft Azure, Riot Games, Twitch, Qualys, and Apple.

The system uses Zeek [17] to extract flows from the set of connections between each device and remote server. Devices in the network are assigned dynamic, temporary IP addresses by DHCP, which we normalize using contemporaneous DHCP logs to convert these dynamic IP address to per-device MAC addresses. To protect user privacy, the IP and MAC addresses for the devices we study are anonymized, and the raw data is discarded after being processed. As well, to avoid analyzing traffic from campus visitors we discard information for devices that appear on the network for fewer than 14 days. In addition, we use contemporaneous DNS logs to convert remote IP addresses (i.e., the servers communicating with the devices we study) to domain names (hence, allowing us to distinguish between different services in use). Finally, we classify individual on-campus MAC devices as being desktop, mobile or IoT devices using multiple heuristics, including analysis of User-Agent strings and organizationally unique identifiers (OUIs) extracted from traffic data. For IoT devices specifically, we employ the methods devised by Saidi et al. with a threshold of 0.5 [15]. Such heuristics are inherently imperfect, so to estimate the error in our approach we manually reviewed 100 random devices in our dataset and verified that 84 were correctly classified.[1] Further details about this collection pipeline, including the device classification approach, are provided by Dekoven et al. [5]. Our data collection and privacy controls have been reviewed and approved by a range of campus oversight entities including UC San Diego's institutional review board (IRB), our campus-wide cybersecurity governance committee, and our campus network operations and cybersecurity groups. In particular, because we do not collect identifiable information about individual users and report results in aggregate, our study has been declared exempt by our IRB.

4 AGGREGATE ANALYSIS

First, we analyze population-level changes and explain how that affects our subsequent analysis. Many students left campus during March 2020, as seen in Figure 1. This graph plots the number of devices active for each day; devices are more likely to have network activity on weekdays than weekends, creating regular dips and spikes. However, this graph shows that students started leaving campus even before classes became fully remote. Before the shutdown, there was a peak of 32,019 active devices; this dipped to a low of 4,973 active devices during the shutdown.

[1]Only two devices in this sample were affirmatively misclassified (e.g. labeling a device as laptop when it was actually a desktop) and the dominant source of error (14 devices) was omission (i.e., devices conservatively classified as"unknown").

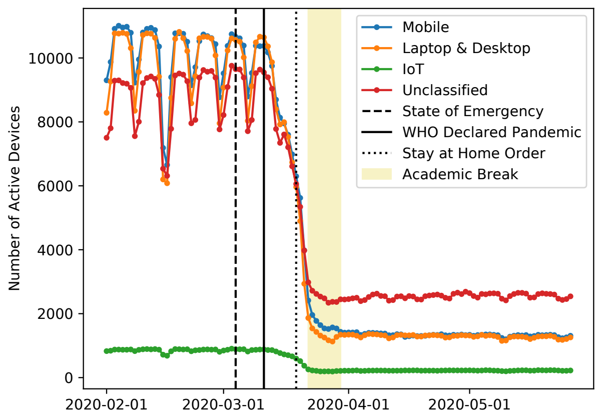

Figure 1: The number of active devices per day, broken down by device type.

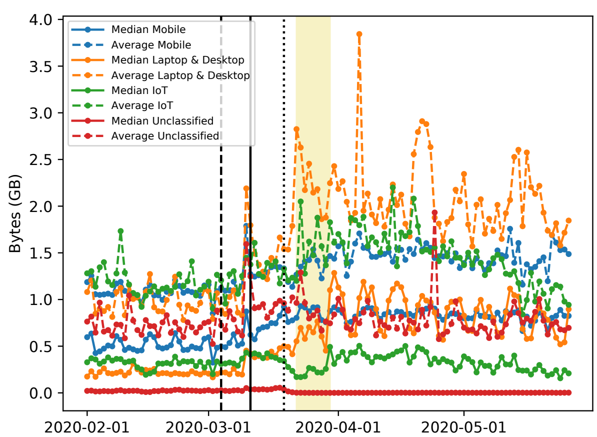

Figure 2: The average and median bytes of active devices per day by device type.

The important dates, indicated by dark vertical lines in our graphs, are as follows:

- 3/4/20: regional authorities issue a state of emergency
- 3/11/20: WHO declares COVID-19 to be a pandemic
- 3/19/20: regional authorities issue a stay-at-home order
- 3/22/20: academic break starts
- 3/30/20: academic break ends; classes resume online

We also examine the distribution of devices by device type in this graph. The number of desktop/laptop and mobile devices on campus seem to follow a 1:1 ratio. After the campus shutdown, the number of unclassified devices dominates the number of IoT, mobile, and desktop/laptop devices[2].

To understand how much traffic each type of device generates, we plot the mean and median bytes per device in Figure 2. We note that some high-volume traffic devices skew the means to be much greater than the medians; this is especially noticeable for IoT and unclassified devices, where the difference spans several orders of magnitude. As a result, the rest of the analysis in this work will rely on median values. The median traffic in this graph shows that while unclassified devices dominate in number, they do not dominate in volume of traffic. Instead, pre-shutdown mobile devices dominate, and post-shutdown mobile and desktop/laptop devices exhibit roughly equal volumes of traffic.

[2] Unclassified devices are those that we cannot identify using any of the methods describe in the previous section; while further analysis is required, we suspect that they are actually mobile and desktop devices with large outliers in device behavior.

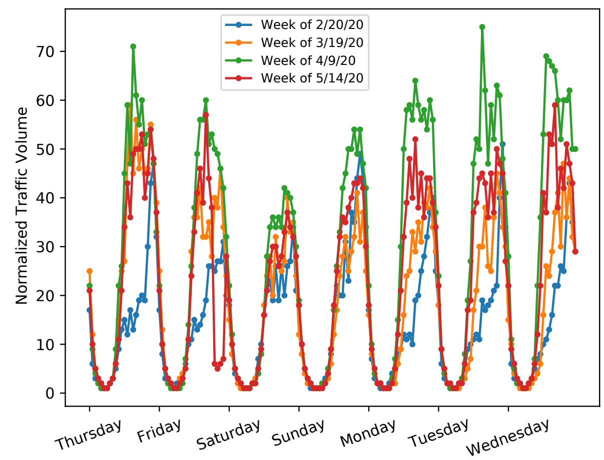

Figure 3: Normalized median volume of traffic per device per hour of week for four weeks of the measurement period.

We find 6,522 devices in total remained on campus after the shutdown; we refer to this group of devices as *post-shutdown users*. Our subsequent analyses focus on these devices to ensure that changes we see in the data reflect changes due to the COVID-19 shutdown, and not changes due to a shift in demographics.

4.1 Overall Internet Usage

As students spent more time indoors and less time socializing, time spent online drastically increased. This change is reflected in the total volume of traffic, which increases by 58% from February to April and May 2020. Traffic in April and May 2020 was 53% higher than in 2019, which indicates that the pandemic was a driving factor behind these changes. Furthermore, on average, users visited 34% more distinct sites in April and May 2020 than in February 2020, which shows users expanded the range of sites they visited.

Notably, however, the weekend dips in traffic persist—a trend not found in other measurement studies of this time period, such as by Feldmann et al. [7]. Figure 3 shows how the hourly volume of traffic changed over the course of the early pandemic months; the data is normalized by the minimum volume of traffic across all weeks. This graph shows that on weekdays during the shutdown, traffic spikes earlier in the day and peaks at higher volumes than in February. In contrast, weekends are relatively unchanged.

4.2 Student population

We hypothesize that foreign students are over-represented in the set of post-shutdown users because it would have been more difficult for these students to find flights to return home at the start of the

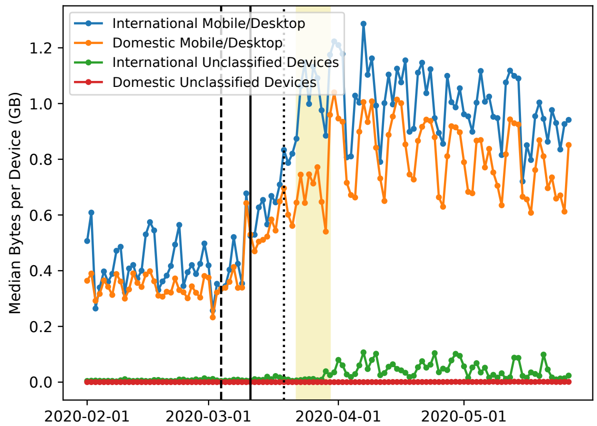

Figure 4: Median bytes per device, excluding Zoom traffic, for international and domestic post-shutdown users. We consider mobile and desktop devices separately from unclassified devices, and exclude IoT devices here.

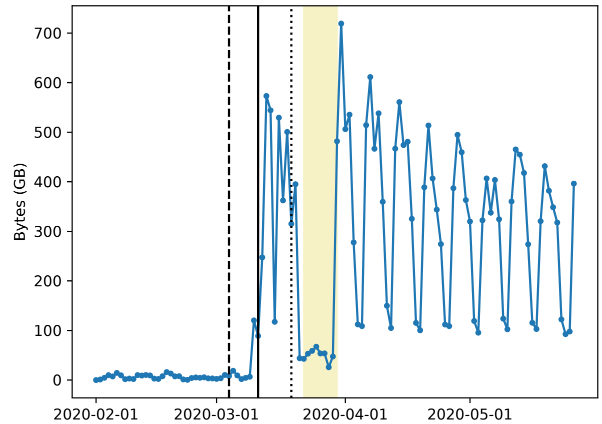

Figure 5: Daily aggregate Zoom traffic for post-shutdown users from February through May 2020.

pandemic. Moreover, we suspect these students to have different patterns of Internet usage compared to their domestic peers, so we seek to separate the two classes of users.

Due to privacy concerns, we do not have ground truth regarding which devices in our dataset are owned by international students. Instead, we identify (likely) international students based on their network traffic. First we collect the geolocation data for every IP address that was visited by a post-shutdown user during the month of February, excluding CDNs (Akamai, AWS, Cloudfront, and Optimizely). We exclude these CDNs because they give information about the user's *device* location, but not the location of the sites the user is visiting. Then, for each device, we calculate the geographic midpoint of the destination of each of that device's connections during the month of February. We weight each connection by its number of bytes and then translate this weighted midpoint into geographic coordinates; if a user's midpoint falls outside the borders of the United States, we classify them as an international student.

Our method of classifying international students is conservative because an international student's weighted midpoint may fall in a different part of the country if they visit a mix of domestic and foreign servers. However, this approach still allows us to identify students who were primarily visiting foreign websites. This approach identifies 1,022 devices presumed to be used by international students, which constitutes 18% of all identified post-shutdown users. In the Fall before the pandemic, reports indicated that about 25% of the entire student body population at UC San Diego was comprised of International students [14], however there are no reports of how many students who remained on campus during the lock-down were domestic or international students.

Given this labeling of foreign and domestic students, we compare their Internet usage by plotting the median volume of traffic per device in Figure 4. In this graph, we exclude Zoom (the tool used for online classes) traffic as it is both large (c.f. Figure 5) and not significantly different between populations. The biggest difference between these two user groups is during academic break, when the volume of traffic increases for international students but remains stable for domestic students. The volume of traffic also stays elevated for international students for the duration of the term relative to their domestic counterparts.

5 APPLICATION USAGE

Here we consider usage of both work and leisure applications.

5.1 Online classes

When classes at the university transitioned fully online, they were primarily hosted on the video conference platform Zoom. Other universities, companies, and social events also moved to Zoom, resulting in over 300 million daily meeting participants as of April 2020 [19]. To analyze Zoom traffic in our dataset, we identify all connections that resolve to a `zoom.us` domain. We also analyze connections where an IP address matches a list of IP addresses from Zoom support [18], and use the Internet Archive Wayback Machine to find any IP addresses that were previously listed on this page, but were subsequently removed.

We plot the aggregate Zoom traffic in Figure 5. While the overall Zoom traffic increased, there are periodic dips that occur during the weekends. The weekday traffic is most active from 8am to 6pm on weekdays, which corresponds to online classes. On weekends, there is a small spike in traffic in the afternoon (not shown), which may indicate that people are using Zoom for entertainment (e.g. calls with family and friends) or for extracurricular activities (e.g. club meetings).

5.2 Social media

As people had to stay indoors during the pandemic, the popularity of some social media sites has increased dramatically. Notably, TikTok's popularity increased by 75% from January to September 2020 [9]. In this section, we study the changes in social media usage throughout the COVID-19 shutdown. We focus on Facebook, Instagram, and TikTok in particular.

For each platform, we manually analyzed traffic from a laptop and mobile device to create signatures of the application's behavior

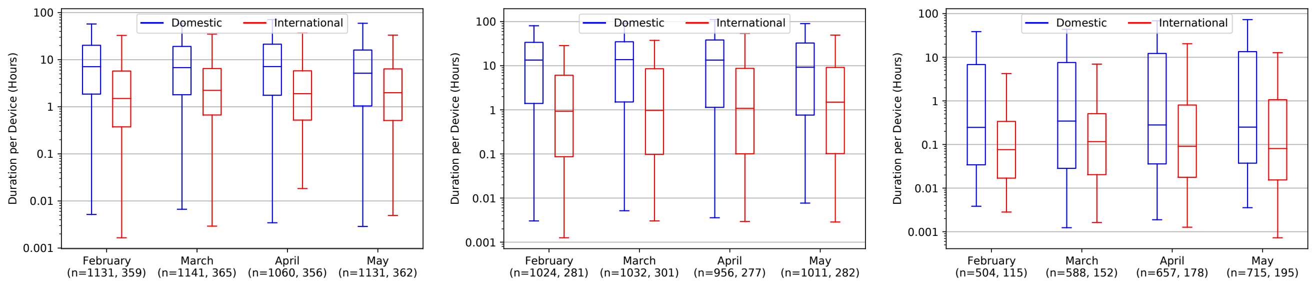

(a) (b) (c)

Figure 6: Box-and-whiskers plot of mobile duration for domestic and international users of (a) Facebook, (b) Instagram, and (c) TikTok. The whiskers extend from the 1st to the 95th percentile. (Note the logarithmic y axes.)

(i.e., the set of sites visited by each). We then compute the duration of each user's sessions. While Zeek provides the duration of each network flow, the social media sites often use multiple domains to serve content to users. For example, in a single Facebook session, a client would receive traffic from `facebook.com`, `facebook.net`, and `fbcdn.net`. So, to compute the duration of an entire user session, we find the bounds of overlapping flows from different domains belonging to the same site.

Another complication in computing duration is that the aforementioned Facebook domains serve content for both Facebook and Instagram services. We use a simple heuristic to differentiate Facebook and Instagram sessions: if any of the domains in a set of overlapping flows delivers Instagram-only content (e.g. traffic from `instagram.com`), then we mark the entire session as an Instagram session. Otherwise, we mark the session as Facebook. This heuristic is reasonable according to our manual traffic analysis, but may overstate Facebook usage and under-represent Instagram usage.

After computing duration for each site, we generate graphs of each user's aggregate duration per month for February through May. We analyze only mobile traffic because there is not significant desktop/laptop traffic even before the pandemic. Figure 6a shows these durations for Facebook mobile users. For domestic users, Facebook usage was relatively unchanged from February through March, but decreased in May. However, the median duration for international students increased during the campus shutdown. While Facebook was more popular for domestic users than international users in February, in May we see that the decrease in domestic usage dampened the distinction between user groups.

We next plot Instagram duration in Figure 6b. For domestic students, the median is relatively unchanged from February through April, but decreases in May. The first quartile also decreases in April and May, indicating that some users were not as active on Instagram later in the campus shutdown. In contrast, the median for international students increases in May. The first and third quartiles also increase from February to March and stay steady for the following months. So international students increased their usage while domestic students maintained or decreased their usage.

Finally we show TikTok duration in Figure 6c. The median duration for domestic users increases from February to March, then decreases in April, and returns all the way back to February's level in May. However, the third quartile and 99th percentile both increase steadily over the months. This indicates that a portion of domestic users kept increasing their TikTok usage, while some users went back to pre-pandemic levels of usage in May. International users were much less active on TikTok than domestic users, but their median usage also increased in March and April compared to February. Like domestic students, the 99th percentile for international students also continuously increased over the months, but the 1st percentile also continuously decreased. This suggests a lot more variance in TikTok usage for this user group.

5.3 Gaming

An important part of online entertainment is gaming. Prior work [7] found that overall gaming traffic increased, and that people started playing games all the time, instead of just in the evenings or on weekends. We also find that gaming traffic increased in our target population. In this section we analyze two important platforms: Steam and Nintendo Switches.

5.3.1 Steam. We developed a signature for Steam, an online platform for PC games, from the set of domains that their customer support recommends whitelisting [16]. We used this signature to plot the total bytes of Steam traffic per user, which is shown in Figure 7a. For each month, we show the traffic for any device that visits Steam, though the set of devices are different for each month. We see that domestic students increase their Steam usage in March, but this usage falls in April and May. International students increase their usage even more during March and April, but again this usage falls in May. The graph of the number of Steam connections, shown in Figure 7b, presents a slightly different picture. Domestic students' median drops over time, while international students' median increases in March and then drops again. We suspect these graphs are different due to game releases or due to the way each game operates. In any event, it seems students initially turned to Steam as source of entertainment during the early periods of the campus shutdown, but found other diversions as time passed.

5.3.2 Nintendo Switches. Despite the Nintendo Switch being released in 2017, sales soared during the COVID-19 pandemic. This

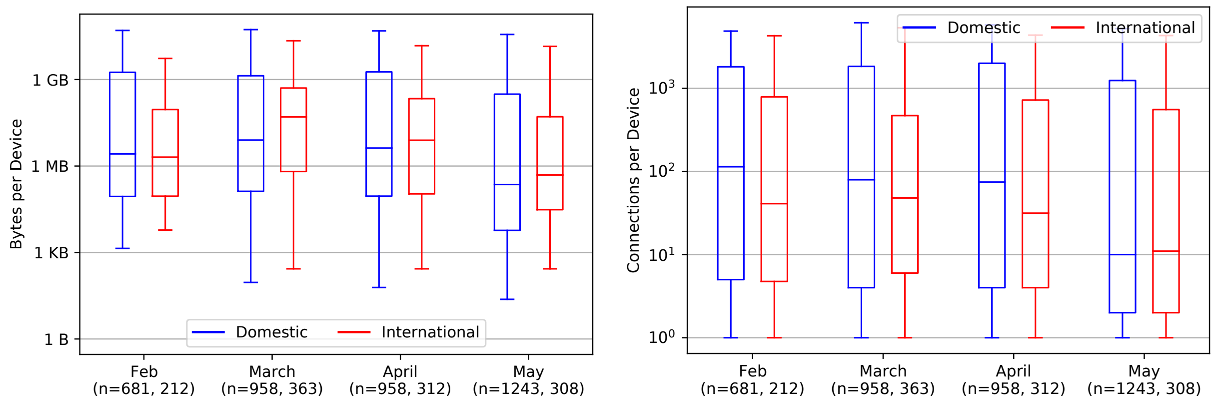

(a) (b)

Figure 7: Box-and-whiskers plot of (a) total bytes and (b) number of connections of Steam traffic per domestic and international post-shutdown user. Whiskers extend from the 1st to 95th percentile.

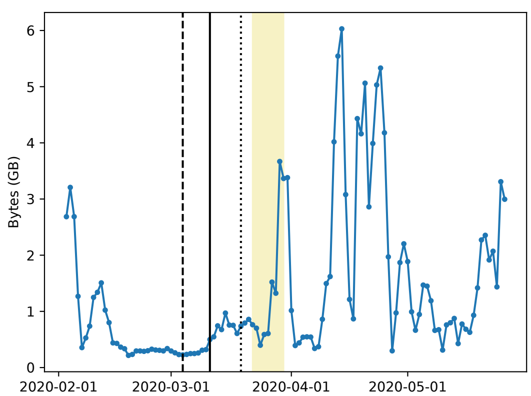

Figure 8: Moving average of gameplay traffic from Nintendo Switch devices per day.

surge has been attributed to an increased demand for indoor entertainment, as well as the release of the *Animal Crossing: New Horizons* game on March 20, 2020 [8].

To detect Nintendo Switch devices from our network logs, we measured the network traffic of a Switch to create a list of domains that a Switch contacts regularly. We cross-checked this list with 90DNS [1], a tool that lets users block contact between a jailbroken Switch and Nintendo servers. Given this signature, we classify devices in our dataset as Switches if at least 50% of their traffic is to the identified Nintendo servers. As users left campus and took their devices home, the number of Nintendo switches fell markedly from 1,097 to 267. We also identified 40 new Switches that first appeared in April and May, indicating that people were seeking out gaming for entertainment during lock-down.

We measured the Nintendo server domains that are used for system updates, game updates and downloads, and other non-gameplay traffic, and confirmed these domains with another script for jailbroken Switches to block Nintendo servers [2]. We filtered these domains out of our network traffic logs to determine actual gameplay and plot a 3-day moving average of this traffic for switches active in both February and May in Figure 8. While there are heavy spikes of usage during academic break and the early part of the Spring academic term, usage returned to almost pre-pandemic levels in late April and early May before increasing again. We suspect that Switch traffic falls immediately after spring break due to classes resuming, but rises as boredom kicks in during the middle of the term. Given their fixed use, we did not attempt to segregate Switch device users.

6 CONCLUSION

In this paper we provide another viewpoint to understanding how Internet-based behaviors changed with the COVID-19 pandemic. By leveraging the existing passive network monitoring pipeline at UC San Diego, we are able to identify changes in application behaviors (for work and leisure) as well as sub-population differences in the residence halls. We find that while per-device traffic increased dramatically in April of 2020, traffic volumes returned to pre-pandemic levels in May. Moreover, we find that entertainment usage increased, classroom-related (i.e., Zoom) platforms were not as utilized on the weekends as the weekdays, and international students displayed different changes in recreational behavior than their domestic counterparts. Our study not only provides a unique viewpoint on the Internet usage of isolated individuals, but also shows that sub-populations exhibited markedly different behaviors even at the same university.

ACKNOWLEDGEMENTS

This work was supported in part by the National Science Foundation through grants CNS-1629973 and CNS-1705050. We are indebted to Geoff Voelker, Stefan Savage, Aaron Schulman, Gautam Akiwake, and Alex Liu for their suggestions and feedback on earlier versions of our analyses.

REFERENCES

[1] 90dns. http://gitlab.com/a/90dns.
[2] Switchblockerforpihole. http://github.com/buggerman/SwitchBlockerForPiHole.
[3] Timm Böttger, Ghida Ibrahim, and Ben Vallis. How the Internet reacted to Covid-19: A perspective from Facebook's edge network. In *Proceedings of the ACM Internet Measurement Conference*, IMC '20, page 34–41, New York, NY, USA, 2020. Association for Computing Machinery.
[4] Massimo Candela and Antonio Prado. Italian operators' response to the COVID-19 pandemic. *SIGCOMM Comput. Commun. Rev.*, 51(1):26–31, March 2021.
[5] Louis F. DeKoven, Audrey Randall, Ariana Mirian, Gautam Akiwate, Ansel Blume, Lawrence K. Saul, Aaron Schulman, Geoffrey M. Voelker, and Stefan Savage. Measuring security practices and how they impact security. In *Proceedings of the Internet Measurement Conference*, IMC '19, page 36–49, New York, NY, USA, 2019. Association for Computing Machinery.
[6] Thomas Favale, Francesca Soro, Martino Trevisan, Idilio Drago, and Marco Mellia. Campus traffic and e-learning during COVID-19 pandemic. *Computer Networks*, 176:107290, 2020.
[7] Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis. The lockdown effect: Implications of the COVID-19 pandemic on internet traffic. In *Proceedings of the ACM Internet Measurement Conference*, IMC '20, page 1–18, New York, NY, USA, 2020. Association for Computing Machinery.
[8] Tom Huddleston. How 'Animal Crossing' and the coronavirus pandemic made the Nintendo Switch fly off shelves, Jun 2020.
[9] John Koetsier. Massive TikTok Growth: Up 75% This Year, Now 33X More Users Than Nearest Direct Competitor, September 2020.
[10] Craig Labovitz. Effects of COVID-19 lockdowns on service provider networks. https://youtu.be/lg-YHkWFjIE, 2020.
[11] Shinan Liu, Paul Schmitt, Francesco Bronzino, and Nick Feamster. Characterizing service provider response to the covid-19 pandemic in the united states. In *Passive and Active Measurement*, pages 20–38, Cham, 2021. Springer International Publishing.
[12] Andra Lutu, Diego Perino, Marcelo Bagnulo, Enrique Frias-Martinez, and Javad Khangosstar. A characterization of the COVID-19 pandemic impact on a mobile network operator traffic. In *Proceedings of the ACM Internet Measurement Conference*, IMC '20, page 19–33, New York, NY, USA, 2020. Association for Computing Machinery.
[13] Yelena Mejova and Nicolas Kourtellis. Youtubing at home: Media sharing behavior change as proxy for mobility around COVID-19 lockdowns. In *13th ACM Conference on Web Science*, WebSci '21. Association for Computing Machinery, 2021.
[14] Gary Robbins. Uc san diego's foreign enrollment hits record. the pandemic may stop it from going higher. https://www.sandiegouniontribune.com/news/education/story/2020-11-16/uc-san-diego-foreign-student-enrollment, 2020.
[15] Said Jawad Saidi, Anna Maria Mandalari, Roman Kolcun, Hamed Haddadi, Daniel J. Dubois, David Choffnes, Georgios Smaragdakis, and Anja Feldmann. A haystack full of needles: Scalable detection of iot devices in the wild. In *Proceedings of the ACM Internet Measurement Conference*, IMC '20, New York, NY, USA, 2020. Association for Computing Machinery.
[16] Steam Support. Required ports for steam. https://support.steampowered.com/kb_article.php?ref=8571-GLVN-8711.
[17] Zeek. Zeek protocol analyzers website. https://docs.zeek.org/en/current/script-reference/proto-analyzers.html.
[18] Zoom. Network firewall or proxy server settings for Zoom.
[19] Zoom. 90-day security plan progress report: April 22, April 2020.