



**HAL**  
open science

# Characterizing Service Provider Response to the COVID-19 Pandemic in the United States

Shinan Liu, Paul Schmitt, Francesco Bronzino, Nick Feamster

► **To cite this version:**

Shinan Liu, Paul Schmitt, Francesco Bronzino, Nick Feamster. Characterizing Service Provider Response to the COVID-19 Pandemic in the United States. Passive and Active Measurement Conference 2021, Mar 2021, Virtual, France. hal-03135284

**HAL Id: hal-03135284**

**<https://hal.science/hal-03135284>**

Submitted on 24 Feb 2021

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Characterizing Service Provider Response to the COVID-19 Pandemic in the United States

Shinan Liu<sup>1</sup>, Paul Schmitt<sup>2</sup>, Francesco Bronzino<sup>3</sup>, Nick Feamster<sup>1</sup>

<sup>1</sup> University of Chicago  
`{shinanliu,feamster}@uchicago.edu`

<sup>2</sup> Princeton University  
`pschmitt@cs.princeton.edu`

<sup>3</sup> Université Savoie Mont Blanc  
`francesco.bronzino@univ-smb.fr`

**Abstract.** The COVID-19 pandemic has resulted in dramatic changes to the daily habits of billions of people. Users increasingly have to rely on home broadband Internet access for work, education, and other activities. These changes have resulted in corresponding changes to Internet traffic patterns. This paper aims to characterize the effects of these changes with respect to Internet service providers in the United States. We study three questions: (1) How did traffic demands change in the United States as a result of the COVID-19 pandemic?; (2) What effects have these changes had on Internet performance?; (3) How did service providers respond to these changes? We study these questions using data from a diverse collection of sources. Our analysis of interconnection data for two large ISPs in the United States shows a 30–60% increase in peak traffic rates in the first quarter of 2020. In particular, we observe traffic downstream peak volumes for a major ISP increase of 13–20% while upstream peaks increased by more than 30%. Further, we observe significant variation in performance across ISPs in conjunction with the traffic volume shifts, with evident latency increases after stay-at-home orders were issued, followed by a stabilization of traffic after April. Finally, we observe that in response to changes in usage, ISPs have aggressively augmented capacity at interconnects, at more than twice the rate of normal capacity augmentation. Similarly, video conferencing applications have increased their network footprint, more than doubling their advertised IP address space.

## 1 Introduction

The COVID-19 pandemic has resulted in dramatic shifts in the behavioral patterns of billions of people. These shifts have resulted in corresponding changes in how people use the Internet. Notably, people are increasingly reliant on home broadband Internet access for work, education, and other activities. The changes in usage patterns have resulted in corresponding changes in network traffic demands observed by Internet service providers. Many reports have noted some of the effects of these changes from service provider networks [1, 5], application providers [19, 23], and Internet exchange points [20]. Generally, previous findings

and conventional wisdom suggest that while overall traffic demands increased, the Internet responded well in response to these changing demands.

Previous work has shed light on the nature of the resulting changes in traffic patterns. In Europe, Internet exchange points saw a 15–20% increase in overall traffic volumes [3], in some cases resulting in peaks in round trip latency in some countries (e.g., Italy) that were approximately 30% higher than normal [12]. For cellular networks in the UK [16], because users were less mobile, downlink traffic volume decreased by up to 25%. While some of the characteristics of shifting traffic demands are known, and certain aspects of the Internet’s resilience in the face of the traffic shifts are undoubtedly a result of robust design of the network and protocols, some aspects of the Internet’s resilience are a direct result of providers’ swift responses to these changing traffic patterns. This paper explores these traffic effects from a longitudinal perspective—exploring traffic characteristics during the first half of 2020 to previous years—and also explores how service providers *responded* to the changes in traffic patterns.

Service providers and regulatory agencies implemented various responses to the traffic shifts resulting from COVID-19. AT&T and Comcast have made public announcements about capacity increases in response to increases in network load [1, 5]. The Federal Communications Commission (FCC) also announced the “Keep Americans Connected” initiative to grant providers (such as AT&T, Sprint, T-Mobile, U.S. Cellular, Verizon, and others) additional spectrum to support increased broadband usage [9]. Web conferencing applications Zoom and WebEx were also granted temporary relief from regulatory actions [9]. These public documents provide some perspectives on responses, but to date, there are few independent reports and studies of provider responses. This paper provides an initial view into how some providers responded in the United States.

We study the effects of the shifts in Internet traffic resulting from the COVID-19 pandemic response on Internet infrastructure. We study three questions:

- *How did traffic patterns change as a result of COVID-19?* Traffic volumes and network utilization are changing as a reaction to changes in user behaviors. It is critical to measure the exact alterations in a long time span.
- *What were the resulting effects on performance?* Considering an expected surge around the dates when states issued stay-at-home orders or declared states of emergency, we seek to observe possible changes in the latency and throughput of network traffic across locations. Further, different ISPs also have different capacity and provisioning strategies, which provides us a finer granularity based on these differences.
- *How did ISPs and service providers respond?* Finally, to deal with the usage boosts and performance degradations during the COVID-19 response, operations and reactions of ISPs and service providers were taken which may explain the changes in network performance. The answer to this question informs us of the networks robustness and their effective disaster provisioning strategies. These questions have become increasingly critical during the COVID-19 pandemic, as large fractions of the population have come to depend on reliable Internet access that performs well for a variety of applications, from video conferencing to remote learning and healthcare.

To answer these questions, we study a diverse collection of datasets about network traffic load, through granular measurements, proprietary data sharing agreements, and user experiences, as well as extensive baseline data spanning over two years.

**Summary of findings.** First, we study the traffic pattern changes in the United States (§4) and find that, similar to the changes previously explored for European networks, our analysis reveals a 30–60% increase in peak traffic volumes. In the Comcast network in particular, we find that downstream peak traffic volume increased 13–20%, while upstream peak traffic volume increases by more than 30%. Certain interconnect peers exhibit significant changes in the magnitude of traffic during the lockdown. Second, we observe a temporary, statistically significant increase in latency lasting approximately two months (§5). We observe a temporary increase of about 10% in average latency around the time that stay-at-home orders were issued. Typical latency values returned to normal a few months after these orders were put in place. We also find heterogeneity between different ISPs. Finally, we explore how service providers responded to this increase in traffic demands by adding capacity (§6). ISPs aggressively added capacity at interconnects, more than 2x the usual rates. On a similar note, application service providers (e.g., video conferencing apps) increased the advertised IP address space by 2.5–5x to cope with the corresponding 2–3x increase in traffic demand.

## 2 Related Work

The pandemic response has modified people’s habits, causing them to rely heavily on the Internet for remote work, e-learning, video streaming, etc. In this section, we present some previous efforts in measuring the effects of COVID-19 and past disaster responses on networks and applications.

**Network Measurements during COVID-19.** Previous work has largely focused on aggregate traffic statistics surrounding the initial COVID-19 lockdowns. Traffic surged about 20% in Europe for broadband networks [12]. In the United States, a blog post [18] reveals that the national downstream peak traffic has recently stabilized, but in the early weeks of the pandemic, it showed a growth of 20.1%. For wireless networks in the US, volume increases of up to 12.2% for voice and 28.4% for data by the top four providers were shown in an industry report [6]. Mobile networks in the UK reported roughly 25% drops in downlink data traffic volume [16]. Industry operators have self-reported on their network responses largely through blog posts [1, 5, 14, 17].

For traffic performance changes, different patterns appear in different regions. Facebook shows that less-developed regions exhibited larger performance degradations through their analysis of edge networks [2]. Network latencies were approximately 30% higher during the lockdown in Italy [12]. According to an NCTA report, networks in the United States saw less congestion [18]. Due to decreased user mobility, cellular network patterns have shifted [16]: The authors found a decrease in the average user throughput as well as decreased handoffs. Feldmann et al. [12] observed that the fixed-line Internet infrastructure was able

to sustain the 15–20% increase in traffic that happened rapidly during a short window of one week.

Our work differs from and builds on these previous studies in several ways: First, this study extends over a longer time frame, and it also uses longitudinal data to compare traffic patterns *during* the past six months to traffic patterns in previous years. Due to the nascent and evolving nature of COVID-19 and corresponding ISP responses, previous studies have been limited to relatively short time frames, and have mainly focused on Europe. Second, this work explores the ISP *response* to the shifting demands and traffic patterns; to our knowledge, this work is the first to begin to explore ISP and service provider responses.

**Application Measurements during COVID-19.** Previous work has also studied application usage and performance, such as increases in web conferencing traffic, VPN, gaming, and messaging [12]. Favale et al. studied ingress and egress traffic from the perspective of a university network and found that the Internet proved capable of coping with the sudden spike in demand in Italy [8]. Another paper used network traffic to determine campus occupancy at the effect of COVID-19 related policies on three campus populations across Singapore and the United States [25]. The cybercrime market was also statistically modeled during the COVID-19 era to characterize its economic and social changes [24].

**Network measurements of other disasters.** While COVID-19 responses are ongoing and evolving, making measurement efforts incomplete, network responses under other disastrous events can be informative. In 2011, the Japan earthquake of Magnitude 9.0 caused circuit failures and subsequent repairs within a major ISP. Nationwide, traffic fell by roughly 20% immediately after the earthquake. However, surprisingly little disruption was observed from outside [4]. In 2012, Hurricane Sandy hit the Eastern seaboard of the United States and caused regional outages and variances over the network [15]. For human-caused disasters such as the September 11th attacks, routing, and protocol data were analyzed to demonstrate the resilience of the Internet under stress. Their findings showed that although unexpected blackouts did happen, they only had a local effect [21]. Oppressive regimes have also caused Internet outages, such as a complete Internet shutdown due to censorship actions during the Egypt and Libya revolts [7], where packet drops and BGP route withdrawals were triggered intentionally.

Although there have been several preliminary measurements of the effects of the COVID-19 response, none have holistically studied traffic data, performance analysis, routing data, and ISP capacity information together, as we do in this paper. It is crucial to collect and correlate such information to better understand the nature of both traffic demands, the effects of these changes on performance, and the corresponding responses. This paper does so, illuminating the collaborative view of responses of service providers in the United States.

### 3 Data

We leverage multiple network traffic datasets to facilitate our study:

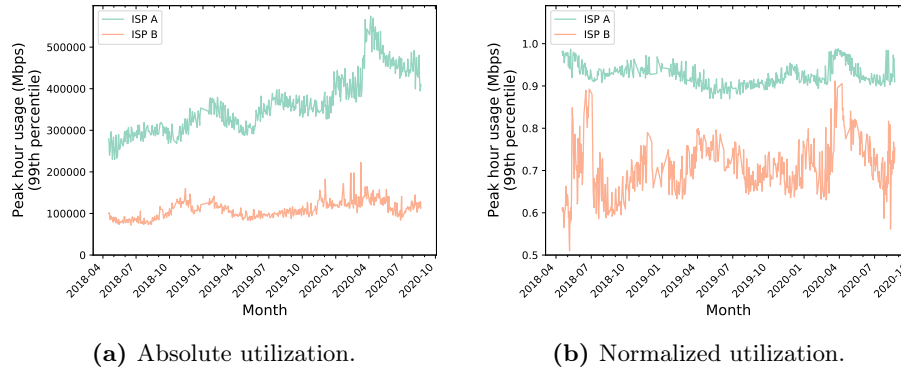
**Traffic Demands and Interconnect Capacity: Internet Connection Measurement Project.** We leverage a dataset that includes network interconnection statistics for links between 7 anonymized access ISPs and their neighboring

partner networks in the United States [11]. These access networks contain about 50% of broadband subscribers across all states within America. At each interconnect interface connecting a neighboring partner network, the access ISP collects IPFIX data. The dataset contains roughly 97% of links (paid peering, settlement-free peering, and ISP-paid transit links) from all participating ISPs. All of the links represented in the dataset are private (i.e., they do not involve public IXP switch fabrics). The dataset consists of flow-level statistics over five-minute intervals, including: timestamp, region (as access ISPs may connect to a partner network in multiple geographic regions), anonymized partner network, access ISP, ingress bytes, egress bytes, and link capacity. In terms of either bytes or packets over a period of time, each five-minute interval provides the sum of the utilization of traffic flows that were active during that interval. We also calculate secondary statistics from the dataset, including: timestamp for the peak ingress and egress hour for each day on each link in terms of usage, ingress/egress peak hour bytes, and daily 95th and 99th percentile usage.

**Performance Data: Federal Communications Commission Measuring Broadband America (MBA).** We analyze the FCC’s ongoing nationwide performance measurement of broadband service in the United States [10]. The raw data is collected from a collection of distributed measurement devices (named Whiteboxes) placed in volunteer’s homes across all states of America and operated by SamKnows. The sample includes tiers composed by the top 80% of the subscriber base for each ISP and is representative. Measurements are conducted on an hourly basis. The dataset includes raw measurements of several performance metrics, such as timestamp, unit ID, target server, round trip time, traffic volume, etc. Each Whitebox also includes information pertaining to its ISP, technology, and state where it is located. We also define dates related to the status of the pandemic response (e.g., stay-at-home orders, state of emergency declaration, etc.). Based on these, we can compute more statistics for specified groups (e.g., break into ISPs): average and standard deviation among Whiteboxes, daily 95th and 99th percentile latency/throughput.

To keep the network capacity consistent and to record eventual changes solely based on utilization factors, we pre-process the MBA dataset with several filters. First, we filter the non-continuous data within the dates of interest (Dec. 1st, 2019 to June, 30th 2020, and the previous year) to capture successive shifts. Then, we eliminate the Whiteboxes which do not aggregate a statistically significant amount of data, such as some states, ISPs, and technologies with limited data (e.g., satellite). Finally, we choose the measurements from Whiteboxes to the top 10 most targeted servers across the United States to represent the overall US performance. We take this decision because servers with less measurements will have higher variance in sample, and introduce unexpected errors when tracked across time. These servers are sparsely located in major cities of the US and they have the most Whiteboxes (over 200 for each ISP) connecting with them.

**IP Prefix Advertisements: RouteViews.** To gain insight into changes in IP address space, we parse Internet-wide BGP information globally from sev-



**Fig. 1:** 99th percentile interconnect link utilization for two ISPs.

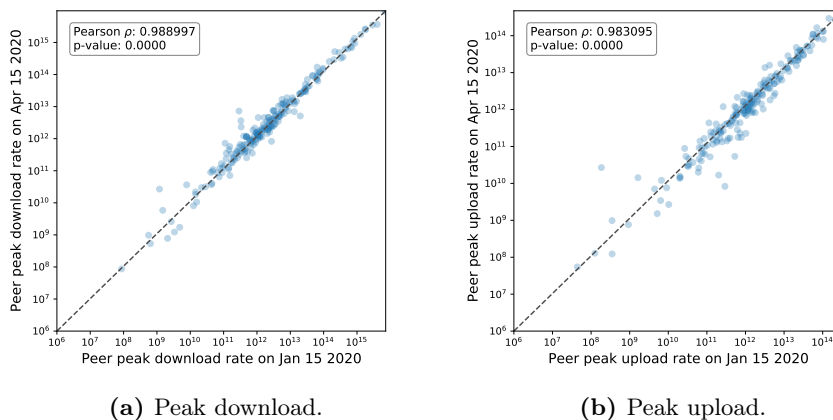
eral locations and backbones via RouteViews. Raw RIBs (Routing Information Bases) files were obtained from RouteViews [22] data on a weekly basis. The average of each Tuesday is computed to represent that week. The RIBs are then parsed to obtain IPv4 Prefix-to-Autonomous System (AS) relationships, including mappings of IP prefix, prefix length, paths of AS numbers. In Section 6.2, we compute the total advertised IPv4 spaces for AS numbers associated with two popular video conferencing applications: Zoom and Cisco WebEx [9].

#### 4 How did traffic demands change?

Because most previous studies [3, 12, 16] focus on Europe, we begin our explorations by validating whether similar traffic changes are observed in the United States. We consider peak hour link utilization from the Interconnect Measurement Project as a measure of traffic demand. We pre-process the interconnect dataset and remove anomalous data points that are caused by failures in the measurement system. In particular, we do not analyze dates that are greater than two standard deviations outside of a 60-day rolling mean for each link. Due to confidentiality reasons, we present the results in aggregation for the United States as a whole.

Figure 1 shows both the absolute utilization and the utilization normalized against the link capacity for two anonymized ISPs. For each ISP, we plot the value corresponding to the 99th percentile link utilization for a given day. We observe from Figure 1a that ISP A saw a dramatic increase in raw utilization at roughly the same time as the initial COVID-19 lockdowns (early March 2020), with values tapering off slightly over the summer of 2020. ISP B, on the other hand, saw a smaller raw increase in utilization for its 99th percentile links. To better understand whether ISP B’s smaller increase is a byproduct of different operating behaviors, we explore possible trends in the normalized data (Figure 1b). Here we see that both ISPs experienced significant increases in utilization in March and April 2020.

We also investigated how traffic patterns changed between ISP A and each of its peers, in both the upstream and downstream directions. For this analysis, we



**Fig. 2:** Peer link utilization for ISP A between January 15 to April 15, 2020.

focused on the dates around the utilization peaks shown in Figure 1. We compared the peak hour download and upload rates on all of ISP A’s interconnects on (1) January 15, 2020, and (2) April 15, 2020 (Figure 2). In general, we see that traffic patterns to peers do not vary greatly between the two dates. We do see, however, that traffic volumes to (and from) some peers change significantly—some by several orders of magnitude. The identities of the peers are anonymous in the dataset, but some patterns are nonetheless clear: For example, some peers show an increase of upstream utilization by two or three orders of magnitude. Such drastic changes may be attributable to users working from home and connecting to services that would cause more traffic to traverse the peer link in the upstream direction. We confirmed these results with the operators at ISP A and report that they observed that streaming video traffic decreased from 67 to 63% of the total traffic, but video conferencing increased from 1% to 4% as a percentage of overall traffic.

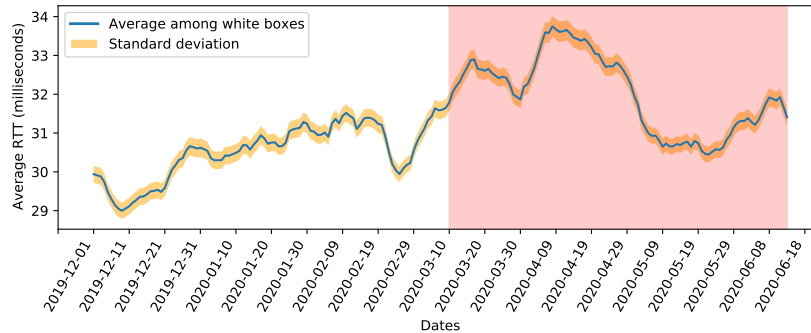
## 5 What was the effect on performance?

The surge in interconnect utilization poses a challenge for service providers, as high utilization of interconnects can potentially introduce high delays for interactive traffic, packet loss, or both. These effects can ultimately be observed through changes in latency (and, potentially, short-term throughput). To examine whether we can observe these effects, we look into the latency and throughput reported by the Measuring Broadband America (MBA) dataset [10]. We explore these effects over the course of several years to understand whether (and how) performance anomalies that we observe during COVID-19 lockdown differ significantly from performance anomalies observed during other time periods.

### 5.1 How performance changed after lockdown.

To better understand how performance changed during the COVID-19 lockdown in the United States, we explored how latency evolved over the course of 2020. To establish a basis for comparison, we show the time period from late 2019





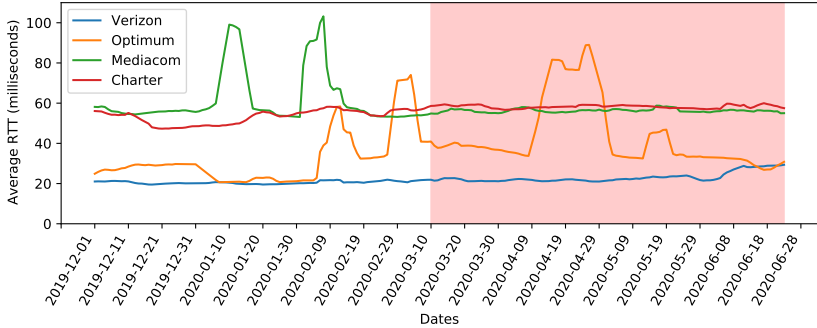
**Fig. 3:** Daily changes of latency from Dec. 2019 to June 2020. The lockdown period is marked in red. Change in average latency across the non-satellite ISPs in the FCC MBA program reflect a small (2–3 ms) but significant increase in overall average latency. (Note: y-axis does not start at zero.)

through mid-2020. The Appendix also contains a similar analysis for the 2018–2019 time period. We compute the average latency per-Whitebox per-day, and subsequently explore distributions across Whiteboxes for each ISP. (As discussed in Section 3, we consider only Whiteboxes in fixed-line ISPs for which there are an adequate number of Whiteboxes and samples.) We use March 10th<sup>4</sup>, the average declaration of emergency date [13], to mark the beginning of the COVID-19 pandemic phase (red shaded for figures).

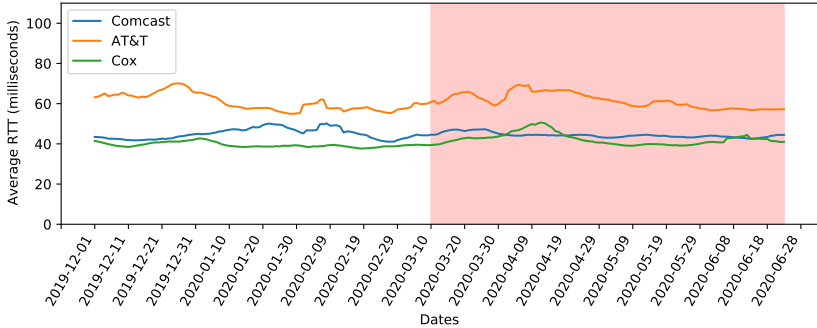
**Longitudinal evolution of aggregate, average round-trip latency.** Figure 3 shows a seven-day moving average of average round-trip latencies between all Whiteboxes in this study. We observe an increase in average round-trip latency by as much as 10%, this increase in mean latency is significant, corresponding to 30x standard deviation among all Whiteboxes. At the end of April, latencies return to early 2020 levels. It is worth noting that, although this increase in average latency is both sizable and significant, similar deviations and increases in latency have been observed before (see the Appendix for comparable data from 2018–2019). Thus, although some performance effects are visible during the COVID-19 lockdown, the event and its effect on network performance are not significantly different from other performance aberrations. Part of the reason for this, we believe, may be the providers’ rapid response to adding capacity during the first quarter of 2020, which we explore in more detail in Section 6.

**Longitudinal evolution of per-ISP latencies.** In addition to the overall changes in performance, we also explored per-ISP latency and throughput effects before and during the COVID-19 lockdown period. Figures 4 and 5 show these effects, showing (respectively) the 95th and 99th percentiles of average round-trip latency across the Whiteboxes. These results show that, overall 95th percentile latency across most ISPs remained stable; 99th percentile latency, on the other hand, did show some deviations from normal levels during lockdown for certain ISPs. Notably, however, in many cases the same ISPs experienced devia-

<sup>4</sup> note that this is also the launch date of Call of Duty Warzone



(a) 95th percentile of ISP latency (Group 1).



(b) 95th percentile of ISP latency (Group 2).

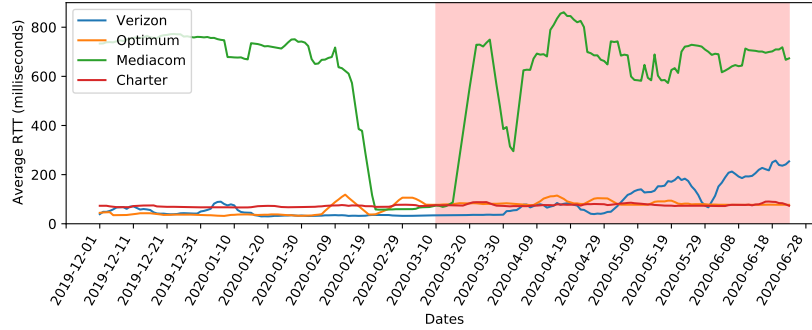
Fig. 4: Latency (95th percentile) for different ISPs.

tions in latency during other periods of time, as well (e.g., during the December holidays).

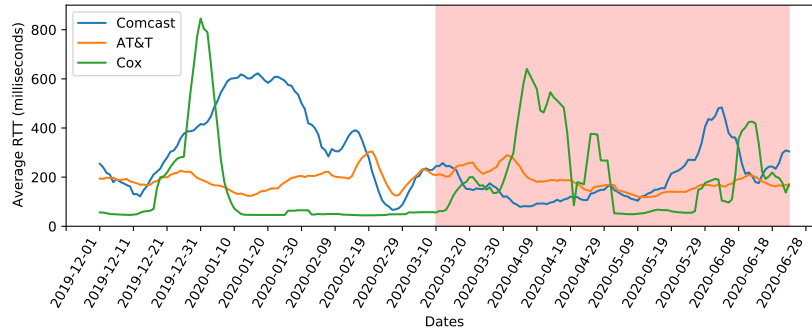
5.2 Throughput-latency relationship

High latencies can sometimes be reflected in achieved throughput, given the inverse relationship between TCP throughput and round-trip latency. To explore whether latency aberrations ultimately result in throughput effects, as well as how those effects manifest at different times of day, we explored the distribution of latencies before COVID-19 emergency declarations (ED), after the ED but before the stay-at-home order (SO). Our hypothesis was that we might see higher latencies (and lower throughputs) during “peak hours” of the day from broadband access networks, with the peak hours effectively expanded to the weekday working hours, in accordance with previous descriptions of these effects [5].

We explored these metrics for a baseline period predating COVID-19, the time between state declaration of emergency and stay at home ordered [13], after stay-at-home declarations were ordered, and two months after stay-at-home ordered. Because these dates differed across states, we used known dates for each state [13] and matched the corresponding dates for each state against the known location of the Whiteboxes.



(a) 99th percentile of ISP latency (Group 1).



(b) 99th percentile of ISP latency (Group 2).

**Fig. 5:** Latency (99th percentile) for different ISPs.

Figure 11 shows the distribution throughput and latency distributions across all Whiteboxes for four time intervals, plotted in four-hour intervals. From Figure 11a, it is clear that the quantiles, median, and maximum latencies all exhibit effects that correlate with these time periods, which are consistent with the latency changes in Figure 3.

The period between ED and SO corresponds to abrupt routing changes, and the latency data thus reflects a corresponding degradation during this time interval, perhaps at least partially due to the fact that providers cannot immediately respond after the initial emergency declaration (we discuss the timeframes during which capacity was added to the networks in Section 6). As the transition continues, SO appears to be a point in time where latency stabilizes. Figure 11b shows that distributions of throughput measurements are more robust, although the upper end of the distribution is clearly affected, with maximum achieved throughputs lower. The median and minimum have negligible changes during time periods in late April suggesting (and corresponding to) aggressive capacity augmentation, which we discuss in more detail in Section 6.

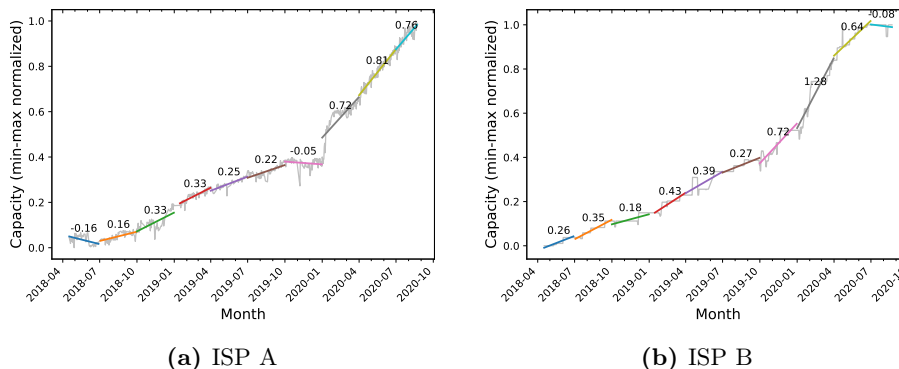


Fig. 6: Normalized interconnect capacity increases for two ISPs.

## 6 How did service providers respond?

In this section, we study how service providers responded to the changes in traffic demands. We focus on the capacity changes during lockdown by inspecting two data sources: (1) to understand how ISPs responded by adding capacity to interconnects, we study the interconnect capacity of two large ISPs in the United States; and (2) to understand how video service providers expanded their network footprints in response to increasing demand, we analyze IPv4 address space from two major video conference providers—WebEx and Zoom—and find that both providers substantially increased advertised IP address space.

### 6.1 Capacity increases at interconnect

We begin by exploring how ISPs responded to changing traffic demands by adding network capacity at interconnect links. To do so, we use the Interconnect Measurement Project dataset. We calculate the total interconnect capacity for each ISP by summing the capacities for all of the links associated with the ISP. To enable comparison between ISPs that may have more or less infrastructure overall, we normalize the capacity values for each using min-max normalization. We again filter out date values that are beyond two standard deviations from a rolling 60-day window mean. To show aggregate infrastructure changes over time, we take all of the data points in each fiscal quarter and perform a least-squares linear regression using SciKit Learn. This regression yields a slope for each quarter that illustrates the best-fit rate of capacity increases over that quarter. We scale the slope value to show what the increase would be if the pace was maintained for 365 days (i.e., a slope of 1 would result in a doubling of capacity over the course of a year). Figure 6 shows the resulting capacity plots.

The overall trend shows how these two ISPs in the United States aggressively added capacity at interconnects—at more than twice the rate at which they were adding capacity over a comparable time period in the previous year. Second, both ISPs significantly added capacity in the first quarter of 2020—at a far greater rate than they were adding capacity in the first quarter of 2019. Recall from the usage patterns shown in Figure 1, ISP A tends to operate their links at nearly full capacity, in contrast to ISP B, where aggregate utilization is well below

90%. Both ISPs witnessed a jump in usage around the lockdown; the response of aggressively adding capacity appears to have mitigated possible adverse effects of high utilization rates. The increase in capacity was necessary to cope with the increased volume: although network performance and utilization ratios returned to pre-COVID-19 levels, the *absolute* traffic volumes remain high.

## 6.2 Increased advertised IP address space

To cope with abrupt changes caused by COVID-19, application service providers also took action to expand their infrastructure. Previous work has observed shifted traffic in communication applications (such as video conferencing apps, email, and messaging) after lockdown [12]. It has been reported informally that many application providers expanded serving infrastructure, changed the routes of certain application traffic flows, and even altered the bitrates of services to cope with increased utilization.

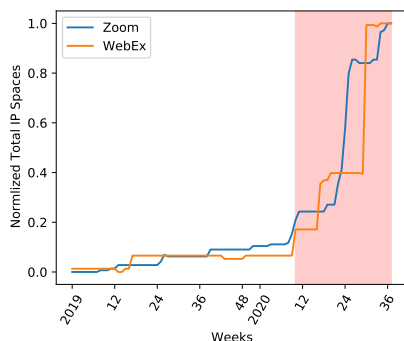
While not all of these purported responses are directly observable in public datasets; however, RouteViews makes available global routing information, which can provide some hints about routes and infrastructure, and how various characteristics of the Internet routing infrastructure change over time. This data can provide some indication of expanding infrastructure, such as the amount of IPv4 address space that a particular Autonomous System (AS) is advertising. In the case of video conference providers, where some of the services may be hosted on cloud service providers or where the video service is a part of a larger AS that offers other services (e.g., Google Meet), such a metric is clearly imperfect, but it can offer some indication of response.

To understand how service providers announced additional IPv4 address space, we parsed BGP routing tables from RouteViews [22]. For each route that originates from ASes of certain application providers, we aggregate IP prefixes and translate the resulting prefixes into a single count of overall IPv4 address space. We focus on two popular video conferencing applications, Zoom and WebEx, since they are two of the largest web conference providers in the United States—as also recognized by the FCC in their recent order for regulatory relief [9]. We track the evolution of the advertised IP address space from the beginning of 2019 through October 2020.

Figure 7 demonstrates how each provider increased the advertised IPv4 address space from before the pandemic through October 2020. After the beginning of the COVID-19 pandemic, both Zoom and WebEx rapidly begin to advertise additional IPv4 address space. Table 7 enumerates the absolute values of advertised IP address space: Zoom and WebEx increased the advertised IP address space by about 4x and 2.5x respectively, as we observe a roughly corresponding 2–3x increase in video conferencing traffic.

## 7 Conclusion

This paper has explored how traffic demands changed as a result of the abrupt daily patterns caused by the COVID-19 lockdown, how these changing traffic patterns affected the performance of ISPs in the United States, both in aggregate and for specific ISPs, and how service providers responded to these shifts in



App	Min	Max
Zoom	9,472	46,336
WebEx	110,080	265,728

**Fig. 7:** Normalized advertised IPv4 space. **Table 1:** Advertised IPv4 space. Red: COVID-19 pandemic phase.

demand. We observed a 30–60% increase in peak traffic rates for two major ISPs in the US corresponding with significant increases in latency in early weeks of lockdown, followed by a return to pre-lockdown levels, corresponding with aggressive capacity augmentation at ISP interconnects and the addition of IPv4 address space from video conferencing providers. Although this paper presented the first known study of interconnect utilization and service provider responses to changes in patterns resulting from the COVID-19 pandemic, this study still offers a somewhat limited viewpoint into these effects and characteristics. Future work could potentially confirm or extend these findings by exploring these trends for other ISPs, over the continued lockdown period, and for other service providers.

## 8 Acknowledgement

This research was funded in part by NSF Award CNS-2028145 and a Comcast Innovation Fund grant. We also thank CableLabs for their help.

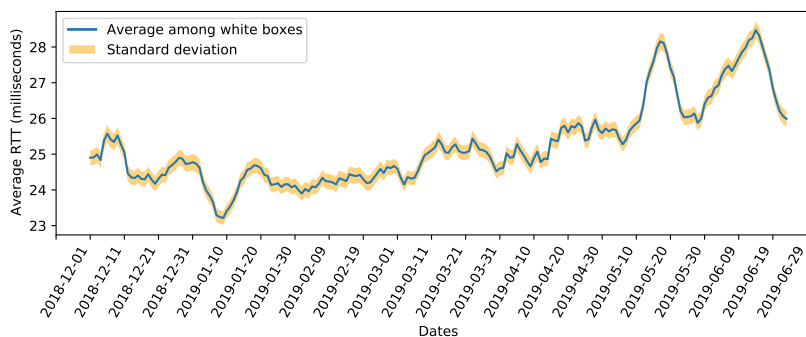
## References

1. AT&T: COVID-19: Our Response (accessed October, 2020), <https://about.att.com/pages/COVID-19.html>
2. Böttger, T., Ibrahim, G., Vallis, B.: How the internet reacted to covid-19 – a perspective from facebook’s edge network. In: ACM IMC (2020)
3. Candela, M., Luconi, V., Vecchio, A.: Impact of the covid-19 pandemic on the internet latency: a large-scale study. arXiv preprint arXiv:2005.06127 (2020)
4. Cho, K., Pelsser, C., Bush, R., Won, Y.: The japan earthquake: the impact on traffic and routing observed by a local isp. In: Proceedings of the Special Workshop on Internet and Disasters. pp. 1–8 (2011)
5. Comcast: COVID-19 Network Update (accessed October, 2020), <https://corporate.comcast.com/covid-19/network/may-20-2020>
6. CTIA: The Wireless Industry Responds to COVID-19: Network Performance (accessed October, 2020), <https://www.ctia.org/homepage/covid-19/#network-performance>
7. Dainotti, A., Squarcella, C., Aben, E., Claffy, K.C., Chiesa, M., Russo, M., Pescapé, A.: Analysis of country-wide internet outages caused by censorship. In: Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference. pp. 1–18 (2011)

8. Favale, T., Soro, F., Trevisan, M., Drago, I., Mellia, M.: Campus traffic and e-learning during covid-19 pandemic. *Computer Networks* (2020)
9. FCC: Keep Americans Connected (accessed October, 2020), <https://www.fcc.gov/keep-americans-connected>
10. FCC: Measuring Broadband Raw Data Releases (accessed October, 2020), <https://www.fcc.gov/oet/mba/raw-data-releases>
11. Feamster, N.: Revealing utilization at internet interconnection points. *TPRC* (2016)
12. Feldmann, A., Gasser, O., Lichtblau, F., Pujol, E., Poese, I., Dietzel, C., Wagner, D., Wichtlhuber, M., Tapidor, J., Vallina-Rodriguez, N., et al.: The lockdown effect: Implications of the covid-19 pandemic on internet traffic. In: *Internet Measurement Conference (IMC '20)* (2020)
13. Foundation, W.: U.S. state and local government responses to the COVID-19 pandemic (accessed October, 2020), [https://en.wikipedia.org/wiki/U.S.\\_state\\_and\\_local\\_government\\_responses\\_to\\_the\\_COVID-19\\_pandemic](https://en.wikipedia.org/wiki/U.S._state_and_local_government_responses_to_the_COVID-19_pandemic)
14. Google: Keeping our network infrastructure strong amid COVID-19 (accessed October, 2020), <https://blog.google/inside-google/infrastructure/keeping-our-network-infrastructure-strong-amid-covid-19/>
15. Heidemann, J., Quan, L., Pradkin, Y.: A preliminary analysis of network outages during hurricane sandy. University of Southern California, Information Sciences Institute (2012)
16. Lutu, A., Perino, D., Bagnulo, M., Frias-Martinez, E., Khangosstar, J.: A characterization of the covid-19 pandemic impact on a mobile network operator traffic. In: *ACM IMC* (2020)
17. McKeay, M.: Parts of a Whole: Effect of COVID-19 on US Internet Traffic (accessed October, 2020), <https://blogs.akamai.com/sitr/2020/04/parts-of-a-whole-effect-of-covid-19-on-us-internet-traffic.html>
18. NCTA: COVID-19: How Cable's Internet Networks Are Performing: METRICS, TRENDS & OBSERVATIONS (accessed October, 2020), <https://www.ncta.com/COVIDdashboard>
19. Nokia: Network traffic insights in the time of COVID-19: April 9 update (accessed October, 2020), <https://www.nokia.com/blog/network-traffic-insights-time-covid-19-april-9-update/>
20. OECD: Keeping the Internet up and running in times of crisis (accessed October, 2020), <https://www.oecd.org/coronavirus/policy-responses/keeping-the-internet-up-and-running-in-times-of-crisis-4017c4c9/>
21. Partridge, C., Barford, P., Clark, D.D., Donelan, S., Paxson, V., Rexford, J., Vernon, M.K., Eisenberg, J.: The internet under crisis conditions: learning from september 11. *Computer Communication Review* **33**(2), 1–8 (2003)
22. RouteViews: University of Oregon Route Views Project (accessed October, 2020), <http://www.routeviews.org/routeviews/>
23. Sandvine: COVID-19 Global Internet Trends (accessed October, 2020), <https://www.sandvine.com/covid-19-trends/>
24. Vu, A.V., Hughes, J., Pete, I., Collier, B., Chua, Y.T., Shumailov, I., Hutchings, A.: Turning up the dial: the evolution of a cybercrime market through set-up, stable, and covid-19 eras. In: *ACM IMC* (2020)
25. Zakaria, C., Trivedi, A., Chee, M., Shenoy, P., Balan, R.: Analyzing the impact of covid-19 control policies on campus occupancy and mobility via passive wifi sensing. *arXiv preprint arXiv:2005.12050* (2020)

## Appendix A Longitudinal Latency Evolution for 2018–2019 (Previous Year)

This section provides a basis for performance comparison in Section 5. Following the same analysis, we choose the exact same time period in the previous year (i.e., late 2018 to mid-2019) in the United States. We compute the average latency per-Whitebox per-day, and subsequently explore distributions across Whiteboxes for each ISP.



**Fig. 8:** Daily changes of latency from Dec. 2018 to June 2019. (Note: y-axis does not start at zero.)

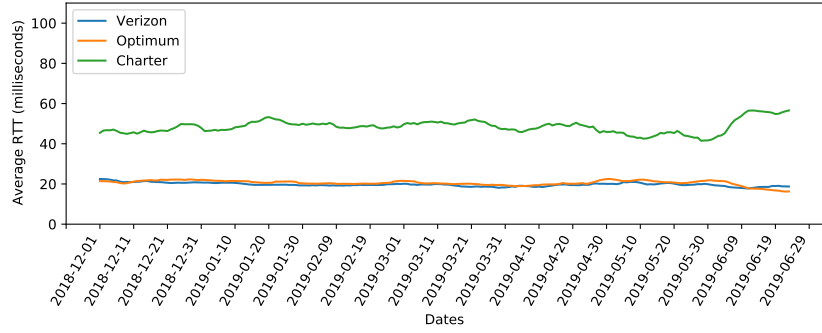
**Longitudinal evolution of aggregate, average round-trip latency.** Figure 8 shows the aggregate average latency per-Whitebox per-day. The previous year has an overall latency of about 6ms lower than 2020. We observe that the latency keeps stable until the end of April, where a deviation of about 2ms is shown. The rate of increase is of about 10%, echoing similar effects around lockdown.

**Longitudinal evolution of per-ISP latencies.** We further break the aggregate results into the granularity of ISPs. We report both 95th and 99th percentile latencies here. Note that in the 95th percentile plot, we show the groups differently, mainly because of major differences of latency for Mediacom and AT&T compared to other ISPs. From Figure 9, we find that the majority of ISPs performed stably, while Mediacom has a large variance in the average RTT. They both have a tail that contributes to what we observed in 8. Figure 10 is grouped the same as Figure 5, which shows that for certain ISPs, they experience similar deviations in latency during similar periods of different years.

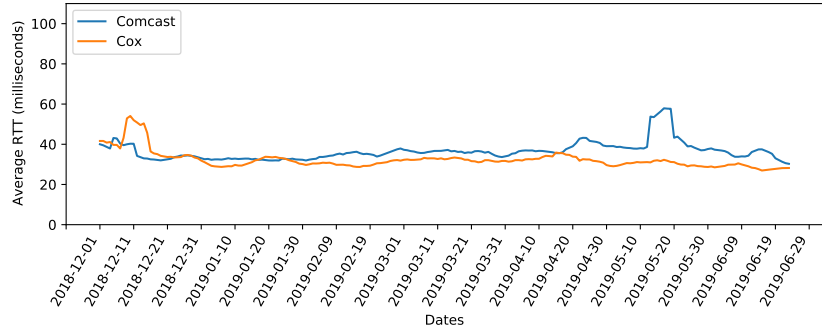
## Appendix B Throughput-latency relationship

We put a supplementary figure referred to in Section 5 in this appendix. It shows the distributional changes in latency and throughput on a 4-hour basis. Detailed explanations are in the main text.

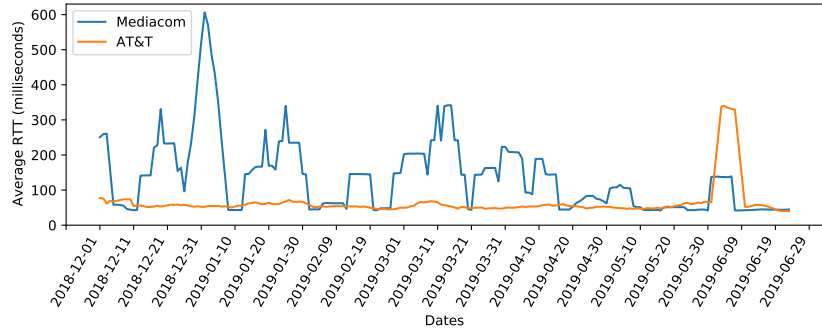




(a) 95th percentile of ISP latency (Group 1)

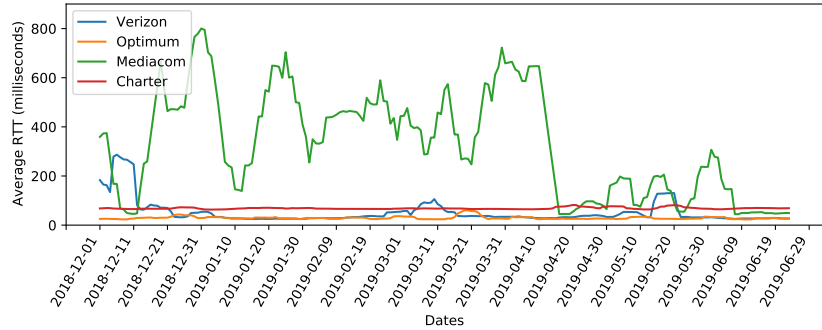


(b) 95th percentile of ISP latency (Group 2)

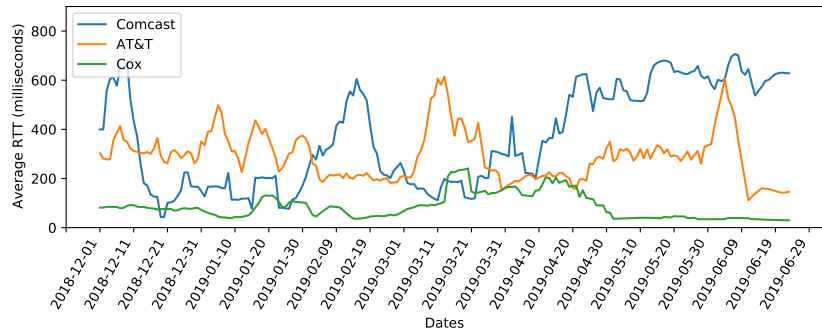


(c) 95th percentile of ISP latency with unstable changes (Group 3)

**Fig. 9:** Latency (95th percentile) for different ISPs.

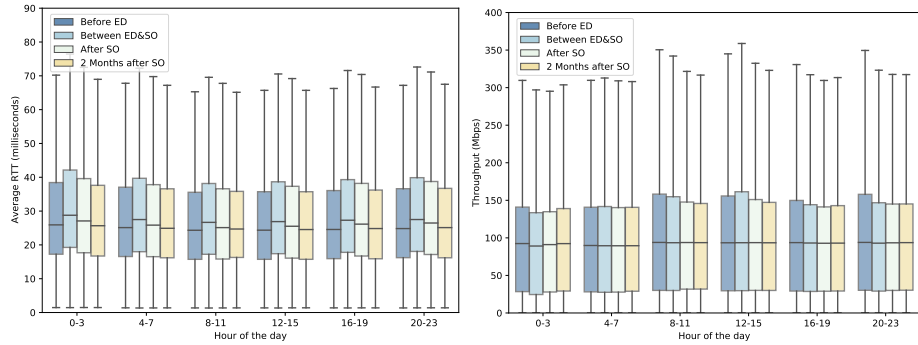


(a) 99th percentile of ISP latency (Group 1)



(b) 99th percentile of ISP latency (Group 2)

Fig. 10: Latency (99th percentile) for different ISPs.



(a) Latency.

(b) Throughput.

Fig. 11: Changes in latency and throughput before and after the lockdown. ED means “Emergency is declared” SO means “Stay-at-home Ordered”.