

How the Internet reacted to Covid-19 – A perspective from Facebook’s Edge Network

Timm Böttger, Ghida Ibrahim and Ben Vallis
Facebook

ABSTRACT

The Covid-19 pandemic has led to unprecedented changes in the way people interact with each other, which as a consequence has increased pressure on the Internet. In this paper we provide a perspective of the scale of Internet traffic growth and how well the Internet coped with the increased demand as seen from Facebook’s edge network.

We use this infrastructure serving multiple large social networks and their related family of apps as vantage points to analyze how traffic and product properties changed over the course of the beginning of the Covid-19 pandemic. We show that there have been changes in traffic demand, user behavior and user experience. We also show that different regions of the world saw different magnitudes of impact with predominantly less developed regions exhibiting larger performance degradations.

CCS CONCEPTS

• **Networks** → **Network measurement**; • **Information systems** → **Social networks**.

ACM Reference Format:

Timm Böttger, Ghida Ibrahim and Ben Vallis. 2020. How the Internet reacted to Covid-19 – A perspective from Facebook’s Edge Network. In *ACM Internet Measurement Conference (IMC ’20)*, October 27–29, 2020, Virtual Event, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3419394.3423621>

1 INTRODUCTION

The Covid-19 pandemic is a global crisis without precedent in recent history. The only other event which comes close is the 1918 Spanish flu pandemic. Many countries world-wide have imposed lock-down measures of varying degrees, leading to closures of offices, schools, restaurants, factories and other venues.

This sudden and unpredictable change in people’s behavior also changed the way Internet products are consumed and used. In this paper we study how changes in user behavior affected demand for Internet egress traffic. We also discuss implications of these changes on the network and on user perceived Quality of Experience (QoE).

The main contributions of this paper are the following:

- We show that the pandemic caused a sharp uptake in traffic, but that this uptake was limited to a short period of time only. This uptake was followed by a phase of increased but stable




Figure 1: Relative change of global edge traffic. Vertical lines mark the implementation dates of lockdown measures in the largest country per region (c.f. Table 2).

request volume. The initial traffic surge exhibited regional differences both in terms of timing and growth.

- We show a significant change in user behavior translating into new traffic trends across products and access types. We observe a surge in popularity of livestream services, although the contribution of this growth to overall traffic is small. We likewise observe a surge in popularity of messaging with variable traffic implications across regions. On the other hand despite relatively lower growth of more traditional social media services like video, the high initial volume of those services led to a significant increase of global traffic.
- We show that the Internet did not cope with this increase in traffic in the same way globally. While North America and Europe did not show any signs of stress in their networks, India, parts of Sub-Saharan Africa and South America did witness signs of network stress coinciding with the traffic surges in the second half of March. Nevertheless, measures taken by operators (like traffic rate limiting or video bitrate capping) and the eventual stabilization of network traffic did allow networks to recover to their pre-Covid-19 performance levels relatively quickly.

2 VANTAGE POINT

For this paper we use Facebook’s global edge network as a vantage point. That is, we only consider user-facing traffic and disregard any internal traffic like intra- or inter-datacenter traffic. Hence we will refer to this traffic as edge traffic. This network serves Facebook’s over 2.5 billion monthly active users distributed around the world. This network comprises a series of PoPs and off-net cache servers with interconnections spread across six continents. This network maintains interconnections with all major ISPs in all regions and at peak serves traffic in excess of 100Tbps. We observe tens of trillions of traffic flows every day. Although this size and footprint allow us to see a significant fraction of the Internet, we acknowledge

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
IMC ’20, October 27–29, 2020, Virtual Event, USA
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8138-3/20/10.
<https://doi.org/10.1145/3419394.3423621>

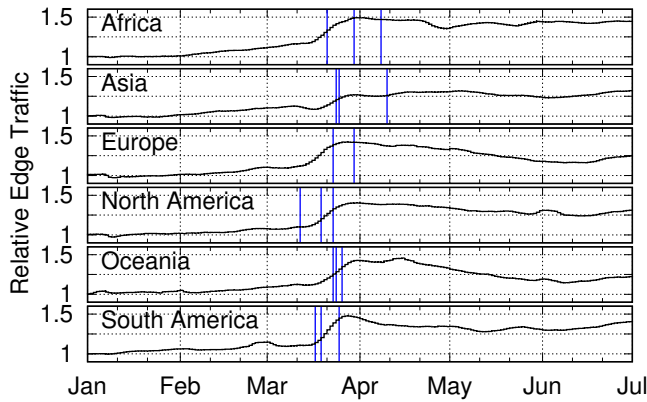


Figure 2: Relative traffic growth per continent. Vertical blue lines mark the implementation date of lockdown measures in the largest three countries per region (c.f. Table 2).

that this paper nevertheless only provides the view from this single network.

3 TRAFFIC PERSPECTIVE

We start by looking at the global traffic footprint. Data is obtained via sFlow sampling on all edge locations at which traffic leaves our network. Figure 1 shows the network’s world-wide total traffic throughput relative to the traffic of January 01, 2020. We depict the data as a time series, showing strong diurnal and weekly patterns. As daily and weekly fluctuations make it hard to discern longer term trends from these shorter term fluctuations, we additionally calculate average traffic rates over the preceding seven days using a rolling window. This is depicted as the smoothed line in the same figure. This way we provide the detailed time series which allows for the differentiation of different days along with the smoothed time series which is better suited to identify the long term trend in the same figure.

In the figure we can see a steady growth in traffic until the second half of March 2020. This traffic growth is expected as it reflects the organic growth of the underlying platform and social network. We observe a significant increase in traffic rate over the second half of March, which then plateaus over April until the end of July. Over the observation period, the seven day period with the highest egress rate has 38.7% more traffic egress compared to the seven day period with the lowest average traffic rate. Toward the end of the observation period, traffic volumes start to decrease again. This is remarkable, as it shows two things. First, from a global perspective we quickly reached a new normal. Once people adapted to new behaviors and routines traffic growth plateaued. Second, we see small decreases in traffic volume, especially toward the end of the observation period. This indicates that at least some fraction of the initial traffic growth was caused by the adaptation to life under the pandemic and its lock-down conditions.

To get a perspective on how and when Covid-19 has impacted traffic in different regions of the world, we segment our global egress data by region. We choose this segmentation to understand how individual regions have fared as a whole, while at the same time averaging away differences between countries in the same region. Figure 2 shows the observed egress rates for the different




Figure 3: Week over week comparison of relative edge throughput growth. Vertical lines mark the implementation date of lockdown measures in the largest three countries per region (c.f. Table 2).

regions. For the sake of readability, we only show seven day rolling averages without the underlying time series. The figure shows a growth in egress in the second half of March, followed by a phase of stability from April to May.

The growth phases in the different regions correlate with the spread of Covid-19. Until the middle of March, North America and Europe show about the same traffic volumes, from then on Europe sees an expedited growth in traffic while traffic levels in North America remain stable. Europe then begins to plateau, while about two weeks later North America sees expedited traffic growth at comparable rates to those seen in Europe. This coincides with the major outbreaks and countermeasures taken by governments in both regions. Interestingly, both regions converge to the same traffic levels before and after their growth phase. South America shows similar, albeit delayed behavior. In contrast, Asia exhibits a small dip in traffic volume prior to the phase of strong growth.

To further assess differences in growth in the different regions, we show relative week over week changes in edge egress traffic volume in Figure 3. We chose to show only four continents as these regions are the main traffic destinations for our network. For this figure we calculated average rates over seven days rolling windows to compensate for weekly and daily shifts in traffic patterns.

The figure highlights two findings. First, in line with the previous figure, continent level traffic growth follows the development of the pandemic: The peak of growth in Europe occurs first, followed by North- and South America and then Asia. Second, for all continents, we observe a clear peak in traffic growth. This suggests that the pandemic-related traffic growth is limited to a short period of time, which seems to coincide with novel response measures being put in place.

4 TRAFFIC AND USER BEHAVIOR

In this section we shift focus to the users. We try to understand what the traffic shift from the previous section tells us about changes in user behavior during lockdown. We use access logs from our CDN infrastructure. In contrast to the sFlow data from the previous section, this allows a more granular dissection of individual requests. The drawback of this approach is that we do not account for a small

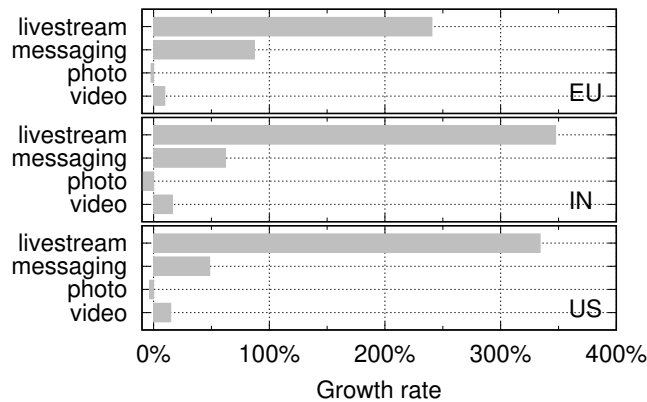


Figure 4: Growth rates of main products in top EU countries (Italy, France, Spain, UK, Switzerland and Belgium), India and USA

amount of traffic which is not served from the CDN (e.g., DNS traffic). The difference in peak traffic throughput between both datasets is 6%.

Product-based overview: We first want to understand which products people consumed more during lockdown. We particularly focus on the four product categories messaging, livestreaming, video and photo as these represent the main services provided by the Facebook family of apps. We use traffic growth of these products pre- and post-Covid-19 as an approximate for changes in user behavior. While we acknowledge that traffic growth per se is not the only measure of changes in user behavior, we use it as an approximate measure that can easily be computed solely through CDN logs.

We compute traffic growth per product type and geography as follows: For each geography we leverage observations from the previous section to identify the dates around which a surge in overall traffic (in bps) for all products combined has been observed.¹ Table 3 lists the dates we used for this paper. Traffic growth rate per product and geography then is the percentage difference between average daily traffic peak rates in the week just before and the week just after the traffic surge.

We compute global, regional and local traffic growth rates for the four product categories. Figure 4 showcases averaged growth rates of these products across Europe, USA and India. Despite regional variations, livestreaming products witnessed an exponential surge in popularity while their contribution to the overall traffic, which we cannot quantify any further, remains minimal.

The usage of messaging services grew globally with high traffic growth rates particularly in Europe and parts of South America, followed by India and Sub-Saharan Africa where messaging apps were already popular pre-Covid-19. Messaging traffic contribution to overall CDN traffic post-Covid-19 varies per country and region with the highest traffic impact observed in India, Sub-Saharan Africa and Spain. Video products grew around 5% globally, with higher rates in North America, Europe and India and lower rates in Africa.

¹Computing traffic growth per product and geography requires identifying the dates of the traffic surge. If the considered geography is a region or continent rather than a country, a precondition is to have different countries within this region witnessing the traffic surge around the same dates (which is the case for instance for the EU).




Figure 5: Relative traffic growth per access type. Vertical lines mark the implementation date of lockdown measures in the largest country per region (c.f. Table 2).

Despite a lower global growth rate than livestreaming products, growth in video products proved to have higher traffic impact given that, in terms of volume, video represents a higher share of overall traffic.

Access-based overview: In this section, we try to understand how lock-down measures have changed the way content is consumed. More specifically we are interested in whether we see more requests using broadband or mobile access technology. CDN logging infers access type based on client or in the case of CGNAT carrier IP, which we use to divide CDN requests into two groups: One group which contains all requests made via the networks of mobile operators, and the second group containing all requests made via fixed lines or broadband. Figure 5 shows the relative growth of CDN traffic, grouped by mobile and broadband networks. As with the previous figures, we show the raw timeline as well as the seven days rolling averages. The figure shows that it is largely broadband consumption where we see traffic growth. We observe a small increase in traffic on mobile networks in the second half of March, but the overall consumption level remains relatively stable. For both access types we observe the largest growth on March 28, 2020. For broadband we observed more than 1.41x as much traffic as in the preceding period, for mobile this peak growth is 1.24x the value of the preceding period.

5 HOW ISPS HANDLED INCREASED TRAFFIC

In this section we attempt to answer whether the Internet was able to cope with increased traffic. We look at user-centric indicators like video QoE as well as network-centric indicators like traffic overflow to public exchanges or transit and round-trip times. We provide these measures as non-exhaustive indicators of path congestion. See Table 4 for the country abbreviations used in this Section.

Video engagement overview: We analyze video traffic as it is one of the products with high traffic weight globally. We compute video watchtime growth in a given country by assessing the percentage difference between the average daily video watchtime in the week pre-Covid-19 traffic surge and the average daily video watchtime in the week post-Covid-19 traffic surge. We use the same methodology for assessing video viewership growth. Video engagement growth is the average of video watchtime and video viewership growth. We calculate video engagement growth for a number of countries that we select based on a mix of factors including Covid-19 exposure (high number of Covid cases), geographic

Country	Video traffic growth	Video engagement growth
IT	30%	35%
FR	20%	20%
ES	30%	40%
CH	20%	15%
UK	18%	20%
US	20%	25%
IN	10%	60%
ZA	3%	35%
NG	3%	10%
EC	10%	43%
CO	10%	40%
PE	20%	40%
BR	20%	25%

Table 1: Video traffic growth and video engagement growth for selected countries.




Figure 6: Global Average Bad Session Rate (BSR) and country rates for US and India.




Figure 7: BSR evolution in main European countries.

diversity (making sure every continent is represented) and size (including major markets for Facebook like India and the US).

As shown in Table 1, we notice a significant gap in growth rates between peak video traffic and daily video engagement in a number of countries. While we cannot pinpoint the exact reason for this gap, it could be an indicator of deteriorating network conditions and bandwidth limitations forcing a dynamic adaptation of video bitrates.

Video QoE Overview: To better understand these gaps in growth rates we now look at changes in video Quality of Experience (QoE). Note that our usage of QoE here is in line with existing research [4, 11, 12, 24].

In order to assess video QoE, we use a composite metric called bad session rate (BSR). A video session is considered bad if it satisfies one or more of the following conditions: it has a slow start (> 1 sec), it witnesses frequent stalls (mean time between rebuffering < 1 min) or if the video encoding resolution is poor given the used




Figure 8: BSR evolution for selected South American countries.

screen. BSR is the ratio of video sessions classified as bad of all video sessions delivered in a given timeframe or geography. The higher BSR the more significant is the number of users witnessing a poor video QoE.

Figures 6, 7 and 8 showcase a surge in the percentage of bad video sessions globally during the second half of March, with the highest percentage (about 8%) around March 25.² Global BSR eventually recovered to pre-Covid-19 levels beginning of April and continued to decline afterwards. Looking at regional and country-level curves we notice that video QoE degradation did not happen in all countries and regions. Namely, video QoE degradation mainly happened in India, some countries in Sub-Saharan Africa, and some South American countries while North America and Europe did not witness major video QoE regressions. We infer that BSR surge, when applicable, was driven by Covid-19-induced traffic growth for the following reasons. First, BSR surges and traffic surges happened simultaneously (around the same dates) in impacted countries. Second, while BSR degradation happened before, the level of BSR degradation in the most impacted countries (India and South Africa) is unprecedented which parallels the fact that Covid-19-induced traffic growth is also unprecedented. Finally, BSR recovered to its normal pre-Covid-19 values in the most impacted countries only after operator intervention (e.g., video bitrate capping or rate limiting) and traffic volumes eventually stabilized. We also note that the countries where BSR degraded have a significant gap between their video traffic and video engagement growth rates, which confirms our hypothesis of network stress.

Overflow to transit and public peering: After looking at the impact to user perceived quality of experience, we now look at the underlying network and how it performed over the same period of time. For a variety of economic and performance reasons, many ASes have a preference for direct interconnections over interconnections via a public exchange or via a transit intermediary. While there are business-strategic reasons to not use the most cost-effective interconnection in some cases, this order of preference is valid in the vast majority of cases. This is important as increased usage of indirect interconnection links can be seen as a sign for congestion on the preferred, direct connection links.³ For a detailed discussion

²Note that all three figures show a correlated peak across all metrics in the second half of April. While we cannot pinpoint the exact cause, we acknowledge that this specific peak might have been caused by our systems rather than being Covid-19 induced.

³While direct congestion measurements might yield better results, we did not have such available for this paper.




Figure 9: Indirect traffic Growth Global




Figure 10: Growth of indirect traffic for selected countries.

of routing and overflow policies and their technical implementation see the papers by Schlinker et al. [28, 29].

In Figure 9, we showcase the growth of indirect traffic globally, computed as the ratio between the daily observed peak of indirect traffic post March 01, 2020 and peak indirect traffic on March 01, 2020. At a global level, we see between 5% and about 25% growth in indirect traffic flowing through transit and public peering during the second half of March, coinciding with the global surge in traffic. While not shown in this figure, we equally see that the indirect traffic contributes more to global egress in the second half of March, although this additional contribution is less than 1% globally. This indicates that, due to congestion, traffic started to overflow from direct links towards public peering and transit routes. The growth in indirect traffic eventually stabilizes in April.

When looking at per country figures in Figure 10, we observe variable growth rates of indirect traffic. For instance, countries that experienced degradation in QoE for video – like India and South Africa – show higher growth rates for traffic over transit and public peering compared to the growth of similar traffic in the USA and other countries where video QoE remained stable. We observe similar tendencies when we look at different South American or European countries.

This comparison reveals that globally the Internet was able to cope with the increased demand. While we see more traffic overflowing to indirect links, the additional contribution of indirect traffic to overall traffic did not exceed 1%. For those individual countries where we saw a non-negligible impact on user experience, we also see a higher growth in indirect traffic. India is a clear case with indirect traffic almost doubling early April with respect to beginning of March. However, even for India, the extra contribution of indirect traffic to overall traffic remained less than 1%. Traffic overflow to indirect paths hints at congestion on the preferred direct




Figure 11: Average round-trip time globally and for selected countries. Normalized against values of March 01, 2020. GI is global RTT.

traffic links which, along with a possible access network congestion, is likely to be one of the factors contributing to the reduced user experience we observed.

Round Trip Times: Path congestion typically goes hand in hand with increased round trip times. Figure 11 shows observed round-trip times of client connections measured from our servers.

Similar to the other metrics we observe, we see an increase in average global round trip time in the second half of March. From April onwards RTT values start to decrease again. At the end of the observation period, the average global round trip time is only slightly elevated at 1.1x the value of the beginning of March.

On a country level we see differences between those countries that showed regression in video quality versus those that showed less pronounced regressions. Server-side round trip times are reasonably stable for the USA and Spain, which is in line with their stable video performance. Italy shows slightly more pronounced variation of RTTs, which again is in line with the relatively small degradations in video performance we observed. The last two countries in this figure, South Africa and India, show significant increases in RTT. And again, these are two countries in which we also observed significant degradations in user-perceived video quality. This reinforces our finding that the degradations in video performance we observed in some countries can be attributed to limited capacity and thus congestion in the country’s networks.

Discussion: For all the countries in our observation set we see that degraded video experience always coincides with an increase in network metrics like RTT and the amount of traffic overflowing to indirect links. While RTT and video QoE degradation could be attributed to traffic rerouting to secondary CDN locations via indirect links, we do believe that these metrics are strong pointers towards network congestion, both from the CDN side and on the last-mile network for the following reasons. First, traffic rerouting to indirect and distant CDN locations is usually triggered by a congestion of direct traffic links. Second, the amount of degradation in video QoE and in RTT cannot be explained by the relatively small additional contribution of indirect traffic to overall traffic. Therefore, the hypothesis of a last-mile congestion causing video QoE degradation and longer RTTs, followed by a congestion on the CDN direct peering links triggering traffic overflow to indirect links, is the most plausible for the most impacted countries.

In summary, while Covid-19 did not cause widespread congestion, the induced additional traffic load did indeed cause localized congestion in some countries.

6 RELATED WORK

The Covid-19 pandemic is not the first event that forced the Internet to react. It has always adapted to predictable one-off events like New Year's Eve and major broadcasts as well as to unpredictable events like flash crowds. Ari et al. and Stading et al. study flash crowds and how flash crowds can be characterized and modeled [3, 30]. There are also works providing a characterization of a wider range of traffic anomalies, including denial of service attacks, network outages, and traffic engineering [21].

Beyond these organic outages caused by user behavior, there are also works studying outages caused by externalities. These studies cover for example (severe) weather conditions like Hurricanes [13], the impact of rain [25] or power outages on the Internet [5].

Although the Covid-19 pandemic is a recent event only, its impact on Internet infrastructure has already been widely been discussed through blog posts from individual companies [1, 2, 10, 15, 18, 31] as well as at network operator meetings [8, 19].

The academic community has also turned to studying the impact of Covid-19. Ribeiro et al. study how Wikipedia received significantly more user requests and how the user's interest shifted towards medical topics with the outbreak of Covid-19 [27]. Vu et al. investigate how cybercrime has risen during the pandemic [32]. Zakaria et al. rely on passive WiFi sensing in order to study the impact of Covid-19 policies on campus occupancy and mobility [38]. Favale et al. utilise a campus network dataset to study how e-Learning has changed their traffic profile [16]. Similar to our work, Feldmann et al. and Lutu et al. [17, 23] also study how the pandemic has changed Internet traffic. These studies use IXP or ISPs with a local (albeit country-wide) footprint, whereas our study draws conclusions from the vantage of a global network. We believe (and encourage the reader to do so) that these studies should be read in conjunction, as their perspectives differ because of the different vantage points that were used.

7 DISCUSSION

The changes in user behavior and increased demand caused by the Covid-19 pandemic have put unprecedented stress on the Internet. In fact, it is the largest traffic surge we have ever observed on a global level. In this section we discuss what we have learned from an operator's perspective and suggest measures the Internet might take to better cope with future events.

At the beginning of the pandemic Facebook, as well as many other organizations, was able to quickly add additional network capacity in sufficient amounts to mitigate the largest increases in user demand. This capacity augmentation was only possible due to network gear already in place inside many networks, which now could quickly be utilized to activate more capacity. This highlights the importance of building network infrastructure with a long enough outlook that also factors in headroom to quickly react to changes.

These capacity augmentations however mostly helped at peering points and traffic exchanges that are on the edge of the Internet but not the actual last mile to the end user. Adding a similar amount of

capacity to last mile access networks is significantly more challenging and expensive. During this pandemic we have seen that, while the core of the Internet handled traffic increases relatively well, the middle and last mile access networks especially in less-developed regions have struggled. It is also clear these aspects of the Internet are crucial for good product performance.

Large hypergiants [7, 20] have long realized and tried to combat this issue. One way is to embed off-network caching servers deeply in the access network to relieve middle-mile pressure from their peering exchange point to the end user [6, 9]. Last mile networks have their own set of challenges, and companies like Google and Facebook have long running initiatives to help develop open, cost effective solutions to help with this aspect of providing network service [14, 22, 26]. These initiatives exist as the burden to build infrastructure to cope with large increases is often too high for ISPs in less developed countries. We believe that open standards and open technology are one way to decrease this burden.

Lastly, this pandemic has shown that the Internet is an ecosystem that thrives through the cooperation of all stakeholders. It is this cooperation that made the Internet scalable and reliable in face of the pandemic. We therefore believe that open communication and discussion between all stakeholders, e.g., traffic consumers, traffic producers and intermediaries is vital for the success story of the Internet to continue.

8 CONCLUSION

In this paper we used Facebook's edge network serving content to users across Facebook's family of apps to provide a perspective on how the Internet coped with and reacted to the surge in demand induced by Covid-19.

We showed that the increase in traffic demand was substantial but that this surge was limited to a short period of time, with traffic subsequently stabilizing at heightened levels.

We then studied how these significant changes in user behavior translated into new traffic trends across products and access types. While a surge in the popularity of livestream and messaging products was accompanied by significant traffic increases for those products, the largest traffic impact resulted from the relatively lower traffic growth of video products. Moreover, we found that traffic increases occurred mainly on broadband networks.

Finally, we assessed the impact of this traffic surge on network stress and performance, where we observed an uneven regional distribution. While North America and Europe did not show any signs of stress in their networks, India and parts of Sub-Saharan Africa and South America did witness signs of network stress translating into degraded video experience, higher amount of traffic overflowing to indirect links and secondary CDN locations, and higher network round trip times. While we cannot pinpoint the exact causes of network stress, we do know that it can be caused by a variety of factors including congestion of direct CDN peering links, overutilization of CDN servers and congestion of ISPs access networks, particularly mobile carriers in emerging markets. Nevertheless, measures taken by operators (such as capacity additions, rate limiting, or capping video bitrates) and the eventual stabilization of network traffic did allow networks to recover to their pre-Covid-19 performance levels relatively quickly.

REFERENCES

- [1] Akamai. 2020. The Building Wave of Internet Traffic. <https://blogs.akamai.com/sitr/2020/04/the-building-wave-of-internet-traffic.html>
- [2] Akamai. 2020. Parts of a Whole: Effect of COVID-19 on US Internet Traffic. <https://blogs.akamai.com/sitr/2020/04/parts-of-a-whole-effect-of-covid-19-on-us-internet-traffic.html>
- [3] Ismail Ari, Bo Hong, Ethan L. Miller, Scott A. Brandt, and Darrell D. E. Long. 2003. Managing Flash Crowds on the Internet. In *11th International Workshop on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2003)*, 12–15 October 2003, Orlando, FL, USA. IEEE Computer Society, 246–249. <https://doi.org/10.1109/MASCOT.2003.1240667>
- [4] Athula Balachandran, Vyas Sekar, Aditya Akella, Srinivasan Seshan, Ion Stoica, and Hui Zhang. 2013. Developing a predictive model of quality of experience for internet video. In *ACM SIGCOMM 2013 Conference, SIGCOMM'13, Hong Kong, China, August 12–16, 2013*. ACM, 339–350. <https://doi.org/10.1145/2486001.2486025>
- [5] Ryan Bogutz, Yuri Pradkin, and John S. Heidemann. 2019. Identifying Important Internet Outages. In *2019 IEEE International Conference on Big Data (Big Data)*, Los Angeles, CA, USA, December 9–12, 2019. IEEE, 3002–3007. <https://doi.org/10.1109/BigData47090.2019.9006537>
- [6] Timm Böttger, Félix Cuadrado, Gareth Tyson, Ignacio Castro, and Steve Uhlig. 2018. Open Connect Everywhere: A Glimpse at the Internet Ecosystem through the Lens of the Netflix CDN. *Computer Communication Review* 48, 1 (2018), 28–34. <https://doi.org/10.1145/3211852.3211857>
- [7] Timm Böttger, Félix Cuadrado, and Steve Uhlig. 2018. Looking for hypergiants in peeringDB. *Computer Communication Review* 48, 3 (2018), 13–19. <https://doi.org/10.1145/3276799.3276801>
- [8] Samuel Burke, Graham Kinsey, Robert Rockell, and David Temkin. 2020. Networks’ responses to Covid-19. In *NANOG 79*. <https://youtu.be/aolGV9eBwro>
- [9] Matt Calder, Xun Fan, Zi Hu, Ethan Katz-Bassett, John S. Heidemann, and Ramesh Govindan. 2013. Mapping the expansion of Google’s serving infrastructure. In *Proceedings of the 2013 Internet Measurement Conference, IMC 2013, Barcelona, Spain, October 23–25, 2013*. ACM, 313–326. <https://doi.org/10.1145/2504730.2504754>
- [10] Comcast. 2020. COVID-19 Network Update. <https://corporate.comcast.com/covid-19/network/may-20-2020>
- [11] Giorgos Dimopoulos, Ilias Leontiadis, Pere Barlet-Ros, and Konstantina Papagiannaki. 2016. Measuring Video QoE from Encrypted Traffic. In *Proceedings of the 2016 ACM on Internet Measurement Conference, IMC 2016, Santa Monica, CA, USA, November 14–16, 2016*. ACM, 513–526. <http://dl.acm.org/citation.cfm?id=2987459>
- [12] Florin Dobrian, Vyas Sekar, Asad Awan, Ion Stoica, Dilip Antony Joseph, Aditya Ganjam, Jibin Zhan, and Hui Zhang. 2011. Understanding the impact of video quality on user engagement. In *Proceedings of the ACM SIGCOMM 2011 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Toronto, ON, Canada, August 15–19, 2011*. ACM, 362–373. <https://doi.org/10.1145/2018436.2018478>
- [13] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. 2013. ZMap: Fast Internet-wide Scanning and Its Security Applications. In *Proceedings of the 22th USENIX Security Symposium, Washington, DC, USA, August 14–16, 2013*. USENIX Association, 605–620. <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/paper/durumeric>
- [14] Facebook. 2020. Facebook Connectivity - Bringing more people online to a faster internet. <https://connectivity.fb.com/>
- [15] Facebook. 2020. Keeping Our Services Stable and Reliable During the COVID-19 Outbreak. <https://about.fb.com/news/2020/03/keeping-our-apps-stable-during-covid-19/>
- [16] Thomas Favale, Francesca Soro, Martino Trevisan, Idilio Drago, and Marco Mellia. 2020. Campus traffic and e-Learning during COVID-19 pandemic. *Comput. Networks* 176 (2020), 107290. <https://doi.org/10.1016/j.comnet.2020.107290>
- [17] Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poesse, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis. 2020. The Lock-down Effect: Implications of the COVID-19 Pandemic on Internet Traffic. In *Proceedings of the Internet Measurement Conference, IMC 2020, Pittsburgh, PA, USA, October 27–29, 2020*. ACM.
- [18] Google. 2020. Keeping our network infrastructure strong amid COVID-19. <https://www.blog.google/inside-google/infrastructure/keeping-our-network-infrastructure-strong-amid-covid-19>
- [19] Craig Labovitz. 2020. Effects of COVID-19 lockdowns on service provider networks. In *NANOG 79*. <https://youtu.be/lg-YHkWFJIE>
- [20] Craig Labovitz, Scott Iekel-Johnson, Danny McPherson, Jon Oberheide, and Farnam Jahanian. 2010. Internet inter-domain traffic. In *Proceedings of the ACM SIGCOMM 2010 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, New Delhi, India, August 30 -September 3, 2010*. ACM, 75–86. <https://doi.org/10.1145/1851182.1851194>
- [21] Anukool Lakhina, Mark Crovella, and Christophe Diot. 2004. Characterization of network-wide anomalies in traffic flows. In *Proceedings of the 4th ACM SIGCOMM Internet Measurement Conference, IMC 2004, Taormina, Sicily, Italy, October 25–27, 2004*. ACM, 201–206. <https://doi.org/10.1145/1028788.1028813>
- [22] Loon. 2020. Loon. <https://loon.com>
- [23] Andra Lutu, Diego Perino, Marcelo Bagnulo, Enrique Frias-Martinez, and Javad Khangoslar. 2020. A Characterization of the COVID-19 Pandemic Impact on a Mobile Network Operator Traffic. In *Proceedings of the Internet Measurement Conference, IMC 2020, Pittsburgh, PA, USA, October 27–29, 2020*. ACM.
- [24] Hyunwoo Nam, Kyung-Hwa Kim, and Henning Schulzrinne. 2016. QoE matters more than QoS: Why people stop watching cat videos. In *35th Annual IEEE International Conference on Computer Communications, INFOCOM 2016, San Francisco, CA, USA, April 10–14, 2016*. IEEE, 1–9. <https://doi.org/10.1109/INFOCOM.2016.7524426>
- [25] Ramakrishna Padmanabhan, Aaron Schulman, Dave Levin, and Neil Spring. 2019. Residential links under the weather. In *Proceedings of the ACM Special Interest Group on Data Communication, SIGCOMM 2019, Beijing, China, August 19–23, 2019*. ACM, 145–158. <https://doi.org/10.1145/3341302.3342084>
- [26] Telecom Infra Project. 2020. OpenRAN. <https://telecominfraopenran.com/>
- [27] Manoel Horta Ribeiro, Kristina Gligoric, Maxime Peyrard, Florian Lemmerich, Markus Strohmaier, and Robert West. 2020. Sudden Attention Shifts on Wikipedia Following COVID-19 Mobility Restrictions. *CoRR abs/2005.08505* (2020). arXiv:2005.08505 <https://arxiv.org/abs/2005.08505>
- [28] Brandon Schlinker, Ítalo S. Cunha, Yi-Ching Chiu, Srikanth Sundaresan, and Ethan Katz-Bassett. 2019. Internet Performance from Facebook’s Edge. In *Proceedings of the Internet Measurement Conference, IMC 2019, Amsterdam, The Netherlands, October 21–23, 2019*. ACM, 179–194. <https://doi.org/10.1145/3355369.3355567>
- [29] Brandon Schlinker, Hyejeong Kim, Timothy Cui, Ethan Katz-Bassett, Harsha V. Madhyastha, Ítalo Cunha, James Quinn, Saif Hasan, Petr Lapukhov, and Hongyi Zeng. 2017. Engineering Egress with Edge Fabric: Steering Oceans of Content to the World. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication, SIGCOMM 2017, Los Angeles, CA, USA, August 21–25, 2017*. ACM, 418–431. <https://doi.org/10.1145/3098822.3098853>
- [30] Tyron Stading, Petros Maniatis, and Mary Baker. 2002. Peer-to-Peer Caching Schemes to Address Flash Crowds. In *Peer-to-Peer Systems, First International Workshop, IPTPS 2002, Cambridge, MA, USA, March 7–8, 2002, Revised Papers (Lecture Notes in Computer Science, Vol. 2429)*. Springer, 203–213. https://doi.org/10.1007/3-540-45748-8_19
- [31] Telegeography. 2020. COVID-19 Network Impact. <https://www2.telegeography.com/network-impact>
- [32] Anh Viet Vu, Jack Hughes, Ildiko Pete, Ben Collier, Yi Ting Chua, Ilia Shumailov, and Alice Hutchings. 2020. Turning Up the Dial: the Evolution of a Cybercrime Market Through Set-up, Stable, and Covid-19 Eras. In *Proceedings of the Internet Measurement Conference, IMC 2020, Pittsburgh, PA, USA, October 27–29, 2020*. ACM.
- [33] Wikipedia. 2020. COVID-19 pandemic in Canada. https://en.wikipedia.org/wiki/COVID-19_pandemic_in_Canada
- [34] Wikipedia. 2020. COVID-19 pandemic in Egypt. https://en.wikipedia.org/wiki/COVID-19_pandemic_in_Egypt
- [35] Wikipedia. 2020. COVID-19 pandemic in Ethiopia. https://en.wikipedia.org/wiki/COVID-19_pandemic_in_Ethiopia
- [36] Wikipedia. 2020. COVID-19 pandemic in Indonesia. https://en.wikipedia.org/wiki/COVID-19_pandemic_in_Indonesia
- [37] Wikipedia. 2020. National responses to the COVID-19 pandemic. https://en.wikipedia.org/wiki/National_responses_to_the_COVID-19_pandemic
- [38] Camellia Zakaria, Ameer Trivedi, Michael Chee, Prashant Shenoy, and Rajesh Balan. 2020. Analyzing the Impact of Covid-19 Control Policies on Campus Occupancy and Mobility via Passive WiFi Sensing. *CoRR abs/2005.12050* (2020). arXiv:2005.12050 <https://arxiv.org/abs/2005.12050>

A COVID-19 LOCKDOWN DATES

For each region we want to capture the dates of lockdown measures affecting the most people. Hence we select per region the top three countries by population. The only exception is Asia where we did not consider China due to limited user traffic from this country. Note that this covers countries coping relatively well with respect to case- and death numbers (e.g., Australia, Germany, Ethiopia) as well as those that coped less well (e.g., Brazil, United Kingdom, USA). When no lockdown date was listed on the overview page, we extracted the lockdown dates from the per country pages. When multiple regions in a country went into lockdown, we chose the earliest date of any region (e.g., the lockdown date of California for the USA). When a country did not go into lockdown officially, we use the date of when social distancing measures were introduced first as the lockdown date. Countries marked with an * are the ones for which we also show dates on the global figures.

Region	Country	Lockdown Date	Source
Africa	Nigeria*	2020-03-30	[37]
Africa	Ethiopia	2020-04-08	[34]
Africa	Egypt	2020-03-21	[35]
Asia	India*	2020-03-25	[37]
Asia	Indonesia	2020-04-10	[36]
Asia	Pakistan	2020-03-24	[37]
Europe	Russia*	2020-03-30	[37]
Europe	Germany	2020-03-23	[37]
Europe	United Kingdom	2020-03-23	[37]
North America	USA*	2020-03-19	[37]
North America	Canada	2020-03-23	[33]
North America	Mexico	2020-03-12	[37]
Oceania	Australia*	2020-03-23	[37]
Oceania	Papua New Guinea	2020-03-24	[37]
Oceania	New Zealand	2020-03-26	[37]
South America	Brazil*	2020-03-17	[37]
South America	Colombia	2020-03-25	[37]
South America	Argentina	2020-03-19	[37]

Table 2: Lockdown dates per region. Lockdown dates of countries marked with * are also shown in the global plots.

B COVID-19 TRAFFIC SURGE DATES

For the calculation of product growth rates in Section 4 we used the following date ranges.

Geography	Days before surge	Days after surge
EU	Mar 05 - Mar 14, 2020	Mar 15 - Mar 24, 2020
IN	Mar 10 - Mar 19, 2020	Mar 20 - Mar 29, 2020
US	Mar 06 - Mar 15, 2020	Mar 16 - Mar 25, 2020

Table 3: Date ranges we used for the calculation of product growth rates in Section 4.

C COUNTRY ABBREVIATIONS

In this paper we use the following abbreviations for countries and regions.

Abbreviation	Country
BR	Brazil
CO	Colombia
CH	Switzerland
EC	Ecuador
EU	Europe
ES	Spain
FR	France
GB	United Kingdom
IN	India
IT	Italy
PE	Peru
US	United States of America
ZA	South Africa

Table 4: Country and region abbreviations used in this paper.