

Implications of the COVID-19 Pandemic on the Internet Traffic

Anja Feldmann^a, Oliver Gasser^a, Franziska Lichtblau^a, Eric Pujol^b, Igmarr Poese^b, Christoph Dietzel^c, Daniel Wagner^c, Matthias Wichtlhuber^c, Juan Tapiador^d, Narseo Vallina-Rodriguez^e, Oliver Hohlfeld^f, Georgios Smaragdakis^{a,g}

^a *Max Planck Institute for Informatics*, Saarbrücken, Germany

^b *BENOCS GmbH*, Berlin, Germany

^c *DE-CIX*, Cologne, Germany

^d *Universidad Carlos III de Madrid*, Spain

^e *IMDEA Networks*, Madrid, Spain and *ICSI*, Berkeley, USA

^f *Brandenburg University of Technology*, Germany

^g *Technische Universität Berlin*, Germany

Abstract

In this paper, we report on a measurement study by researchers from several institutions that collected and analyzed network data to assess the impact of the first wave of COVID-19 (February-June 2020) on the Internet traffic. The datasets from Internet Service Providers, Internet Exchange Points, and academic networks, primarily in Europe, provide a unique view on the changes of Internet traffic due to pandemic and the lockdown that forced hundreds of millions of citizens to stay and work from home. The analysis shows that the increase of Internet traffic was about 15-20 % within a couple of weeks, an increase that is typically spread over multiple months under typical operation. However, traffic during peak hours does not increase by more than 5 %. The increase was noticeably higher for specific applications, e.g., remote work applications, teleconferencing, video on demand; in some cases up to 200 %. However, overall, the Internet reacted well to these unprecedented times.

1 Introduction

As a result of the ongoing COVID-19 pandemic, the global population had to depend on residential Internet connectivity for work, education, social activities, and entertainment. This opens questions on how traffic characteristics changed during this period and if these changes challenged the operation of the Internet infrastructure. In this paper, we summarize a measurement study on Internet traffic shifts in the first wave of the COVID-19 pandemic, i.e., February 2020 to June 2020. Our study provides an empirical and multi-provider perspective on traffic shifts by using data from a diverse set of vantage points: one major Tier-1 Internet Service Provider (ISP), three Internet Exchange Points (IXPs) of which two are located in Europe and one in the US, one metropolitan area educational network, and a mobile operator. The major observations are summarized as follows:

- (1) The traffic volume changes follow demand changes, causing a “moderate” traffic surge of 15-20 % during lockdown for the ISP/IXPs in our study, but decreases up to 55 % at the education network. Even after the lockdown, an increase of about 20 % at one IXP, but only 6 % at the Tier-1 ISP, are still visible.
- (2) Traffic increase takes place, typically, during non-traditional peak hours. Daily traffic patterns are moving to weekend-like patterns.

- (3) Demand for online entertainment contributes the most to the surge of traffic originated by Hypergiants, including cloud and content providers.
- (4) Traffic related to remote working applications, such as VPN connectivity applications and video-conferencing applications, surge by more than 200 %.
- (5) At the IXPs, we observe that port utilization increases. This phenomenon is mostly explained by a higher traffic demand from residential users.
- (6) Traffic changes across networks differ. For example, in an educational network in our study, there was a significant drop (by up to 55 %) in traffic volume on workdays after the lockdown. At the same time, remote working and lecturing cause a surge in incoming traffic, e.g., for email and VPN connections.

2 Datasets

For our study we utilize multiple vantage points in the Internet. These vantage points are located at the backbone and peering points of a major Tier-1 Internet Service Provider (ISP), at the core of the Internet (IXPs) around the globe, and at the edge (a metropolitan university network, a mobile operator). At each location we collected network flow data (NetFlow or IPFIX). The vantage points are the following:

ISP: It is a large Central European ISP that provides service to more than 15 million fixed line subscribers and also operates a transit network (Tier-1).

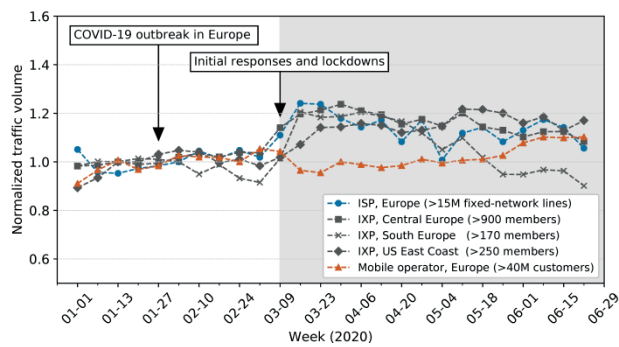


Figure 1 Traffic changes during 2020 at multiple vantage points – daily traffic averaged per week normalized by the median traffic volume of the first up to ten weeks.

IXPs: We consider three major IXPs in our study. The first one has more than 900 members, is located in Central Europe (IXP-CE), and has peak traffic of more than 8 Tbps. The second one has more than 170 members and is located in Southern Europe (IXP-SE). The third one has 250 members and is located at the US East Coast (IXP-US).

Educational: The REDImadrid academic network interconnects 16 independent universities and research centers in the region of Madrid. It serves nearly 290,000 users including students, faculty, researchers, student halls, WiFi networks (including Eduroam), and administrative and support staff.

Mobile operator: The mobile operator is also located in Europe and has more than 40 million customers.

From each vantage point we analyze traffic flows to reason about COVID-19 related traffic shifts.

3 Traffic Volume Changes

In **Figure 1** we show the changes in traffic at all the vantage points in our study from January 1, 2020 (01-01) to June 29, 2020 (06-29). We annotate the date when the COVID-19 outbreak started in Europe – around the date that the World Health Organization declared the novel coronavirus outbreak as a public health emergency of international concern. We also annotate the initial responses and first lockdown dates in Europe.

We observe a significant increase of traffic at all vantage points (except the mobile operator) after the initial lockdowns. The opposite trend is visible after May, when the lockdown restrictions were relaxed and the economy gradually re-opened. Traffic demands for broadband connectivity, as observed at the ISP in Central Europe as well as at the IXP-CE and the IXP-SE increased slowly at the beginning of the outbreak and then more rapidly (within days) by more than 20 % after the lockdowns started. Recall, a 30 % increase in traffic is expected to take place within a year, under typical network operation conditions. The traffic increase at the IXP-US is a bit shifted as the lockdown in the US occurred several weeks later. Apart from the observed similarities in the traffic changes at the ISP and IXP vantage points, the relative traffic increase at the IXPs

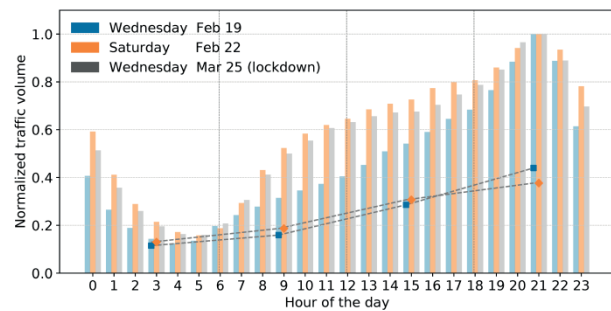


Figure 2 Workday vs. Weekend patterns before and after the lockdown as observed in the ISP.

seems to persist longer while traffic demand at the ISP decreases quickly towards May. This correlates well with the first partial re-opening of the economy, including shop reopenings in this region in mid-April and further relaxations including school openings in May. Reports by Google [1], Comcast [2], Nokia [3], TeleGeography [4], and a blog post by Akamai [5] align with our insights.

The decrease in mobile traffic can be explained by the fact, that people did not go out that frequently and would therefore use their home Wi-Fi more often instead of their phone’s mobile data plan. However, after May, mobile traffic increased significantly. Google reports similar trends in mobile activity [6].

4 The Eternal Weekend Pattern

In light of the global COVID-19 pandemic, a surge of network traffic is expected. However, the way the traffic increases may have different implications. For example, increase of traffic during the peak hour may have significant implications for the dimension of networks and traffic engineering. In **Figure 2** we plot the traffic pattern at the Central European ISP for three days: Wednesday, February 19 and Saturday, February 22 (before the lockdown), and March 25 (after the lockdown). Before the lockdown, the Internet’s regular workday traffic patterns are significantly different from weekend patterns. On workdays, traffic peaks are concentrated in the “peak hour”, i.e., evening time from 7 pm to midnight. For instance, when we compare the traffic (normalized by the peak traffic in a week) on Wednesday, February 19 vs. Saturday, February 22, 2020, we observe that the activity during the peak time is quite similar. However, the traffic during the non-peak time during a weekday is relatively lower than during the weekend. With the pandemic lockdown in March, this workday traffic pattern shifts towards a continuous weekend-like pattern, as can be seen in the daily pattern on March 25, 2020. More specifically, we call a traffic pattern a workday pattern if the traffic spikes in the evening hours and a weekend pattern if its main activity gains significant momentum from approximately 9 am to 10 am.

On a weekend day (orange bars) the pattern looks different, with a much steeper increase during the morning hours

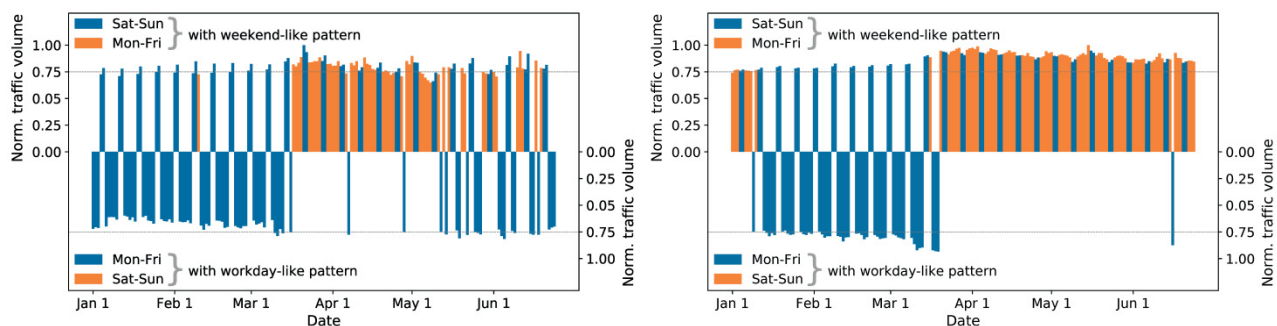


Figure 3 ISP (left) and IXP-CE (right): Workday-like (upper half) vs. Weekend-like pattern (lower half), Jan.-June, 2020.

and a slower growth during the day, again reaching the traffic peak at around 9 pm in the evening. Since more people are staying at home during the day on a weekend compared to a working day, this behavior is affecting the traffic pattern as well. When we investigate the traffic pattern of a working day during lockdown (gray bars), we see that it much more resembles a weekend day than a working day before the pandemic. This nicely visualizes the effect of lockdown measures on Internet traffic patterns. We also noticed that there is very small, if any, increase during the prime time.

We now classify every day based on its traffic pattern to being more workday-like or weekend-like. In **Figure 3** we plot the result of this classification. In the upper part of the graph we show days classified as weekend-like, in the lower part of the graph we show days exhibiting workday-like traffic patterns. If the classification is in line with the actual day (workday or weekend) the bars are colored blue, otherwise they are colored in orange.

For the ISP, see **Figure 3** (left), we find that up to mid-March, most weekend days are classified as weekend-like days and most work days as workday-like days. Between mid-March and mid-May (re-opening of the economy), however, the majority of all days are classified as weekend-like, no matter if they are work days or weekends, and we therefore see a lot of misclassified working days (orange bars). The only exception is the holiday period at the beginning of the year. This pattern changes drastically once the confinement measures are implemented, i.e., almost all days are classified as weekend-like, until the re-opening of the economy around mid-May, 2020.

As previously seen in **Figure 2**, people using the Internet from home during the day exhibits more of a weekend-like traffic pattern. Additionally, as can be seen by the increasing length of the bars starting around mid-March, we see an increase in the overall traffic per day. This increase, however, is not equally distributed over the whole day but is mostly happening in off-peak hours, which can be seen in **Figure 2**. After the opening of the economy, around mid-May, we observe a mixed picture. However, in the IXP-CE, see **Figure 3** (right), the traffic continues to behave like weekend even after the re-opening of the economy. Thus, some changes in the traffic patterns seem to be permanent in the IXP.

5 Application-Level Traffic Analysis

After having analyzed changes in total traffic volume and diurnal patterns, we now turn our attention at traffic of specific application classes. Since port-based classification mixes together a lot of traffic using common ports such as TCP/80 or TCP/443, we use a combination of port-based and AS-based classifications to classify the traffic at the ISP into different groups (see [7] for details on our classification).

In **Figure 4** we plot traffic changes comparing the months of March, April, and June to our base week of February (February 20–27, 2020) for the IXP-CE (top) and the ISP in Central Europe (bottom). We group traffic into seven different traffic classes. Traffic changes are shown for each hour of the day for all days of the week. We highlight the most important changes based on our analysis. We observe a strong increase in the traffic associated to web conferencing, video, and gaming traffic in March as a result of the increasing user demand for solutions like Zoom or Microsoft Teams. Also, as people spend more hours at home, they tend to watch videos or play games, thus, increasing entertainment traffic demands. Interestingly, we also see a decrease in educational traffic in these vantage points. In April and June, web conferencing traffic is still high compared to the pre-pandemic scenario, while we see a slight decrease in content delivery networks (CDNs) and social media traffic. During these months many people are still working from home, but restrictions have been relaxed, which leads to an increase in in-person social activities and a decrease in online ones.

We observed a significant increase in entertainment traffic. In particular, video-on-demand streaming application usage shows high growth rates at the IXP-CE of up to 100%. Interestingly, the ISP only sees a slight growth of about 10% during the lockdown, while in June, i.e., well after the lockdown, the traffic volume drops back to the February level. Recall that the major streaming companies reduced their high resolution streaming in Europe by mid-March for 30 days. In the case of the ISP that covers the March as well as the April weeks. In the US, the trend is the other way around. Notably, this may be a biased measurement, as at the IXP-US the measurement of the video-on-demand

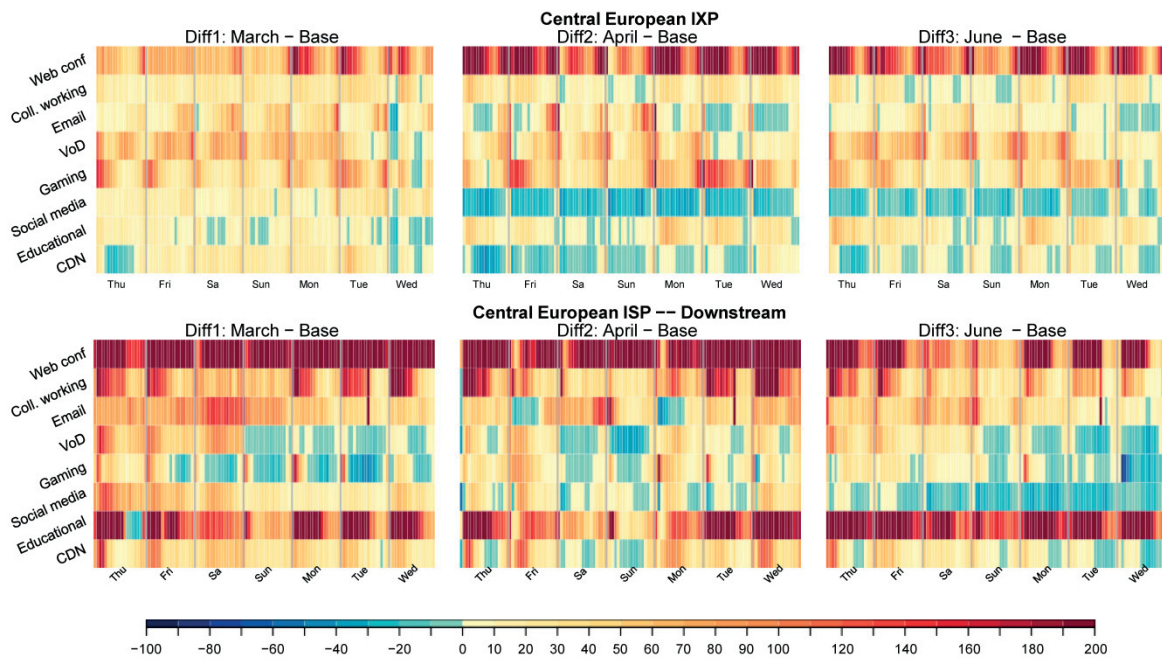


Figure 4 Traffic pattern changes for specific classes of applications during the first wave of the COVID-19 pandemic as observed at the IXP-CE (top) and ISP - downstream traffic (bottom).

class is based on only three ASes, one of which is very large. Consequently, the decrease may reflect a traffic engineering decision of the large AS, e.g., establishing a private network interconnect instead of public peering. The strong growth of gaming applications is more coherent across all three IXP vantage points, especially during the day. While the ISP shows a significant increase during morning hours, it generally leans towards declining. Note, that this effect is mainly caused by unusually high traffic levels in this category in February. Gaming applications, typically used in the evening or at weekends, are now used at any time. The trend starts to flatten in June. This may be attributed to the fact that people going on vacation or spending more time outside. Moreover, we see an increase at the IXPs for social media application traffic during the March week, while the effect quickly diminishes in April. In March the ISP experiences a 70% growth, which slows down in April but not as drastic as at the IXPs. The effects in this class correlate with the gradual de-escalation of the lockdown restrictions in Europe, as people are allowed to leave their homes freely again and resume social life, this traffic decreases. In June, social media usage has returned to figures slightly below the level of March across all vantage points.

Our analysis shows that indeed, the pandemic increased the demand for applications supporting remote teaching and working to guarantee social distancing as shown in our analysis across all vantage points. The Internet could handle this new load due to the flexibility and elasticity that cloud services offer, and the increasing connectivity of cloud providers. Our results confirm that most of the applications with the highest absolute and relative increases are cloud-based. Moreover, the adoption of best practices on designing, operating, and provisioning networks contributed to the smooth transition to the new normal. Due to the

advances in network automation and deployment, e.g., automated configuration management and robots installing cross connects at IXPs with minimal human involvement, it was possible to cope with the increased demand. In summary, our study demonstrates that over-provisioning, network management, and automation are key to provide resilient networks that can sustain drastic and unexpected shifts in demand such as those experienced during the COVID-19 pandemic.

Many employees are working from home during the pandemic. In order to access company infrastructure remotely, some of them use VPN services. Therefore, we expect to see a shift towards more VPN traffic after lockdowns went in place. When looking at the VPN traffic identified with the domain-based technique, we see a significant increase in VPN traffic. During working hours the work-day traffic increases by more than 200% in March compared to the week in February. The increase on weekends is not as pronounced as during workdays, indicating that these traffic shifts occur due to changes in user behavior, i.e., people working from home. When looking at the month of April, we still see a gain in VPN traffic compared to February, although not as high as in March. This is likely due to the gradual lifting of lockdown restrictions in Central Europe, resulting in fewer people working from home in April compared to March. We also note that traffic changes are diverse and highly dependent on the vantage point. For instance, traffic shifts in the REDIMadrid academic network show a complementary behavior. While we observe a 55% drop in traffic volume on workdays even after the lockdown measures loosened, remote working and teaching that continued for the rest of the academic year caused a surge in incoming traffic for email, web, and VPN connections.

6 Conclusion

The COVID-19 pandemic drastically changed working and social habits for the global population. Yet, life continued thanks to the increased digitization and Internet as well as cloud investments that have taken place during the last decade. The Internet indeed played a critical support role for businesses, education, entertainment, and social interactions during these unprecedented times. In this paper, we analyzed network flow data from multiple vantage points, including some at the core, three IXPs located in Europe and the US, and a large academic network and a large ISP at the edge. Together, they allow us to gain a good understanding of the lockdown effect on Internet traffic in more developed countries. Overall, the observed Internet patterns were unique during the first wave of the pandemic, but the Internet reacted well to this challenge. For additional results and details on our methodology, we refer to the full report of our study [7].

Acknowledgement: This work was supported in part by the European Research Council – Starting Grant Resolutio-Net (679158) and by BMBF as BIFOLD – Berlin Institute for the Foundations of Learning and Data (01IS18025A, 01IS18037A).

7 References

- [1] <https://www.blog.google/inside-google/infrastructure/keeping-our-network-infrastructure-strong-amid-covid-19/>
- [2] <https://corporate.comcast.com/covid-19/network>
- [3] <https://www.nokia.com/blog/network-traffic-insights-in-the-time-of-covid-19-june-4-update/>
- [4] <https://www2.telegeography.com/network-impact>
- [5] <https://blogs.akamai.com/sitr/2020/04/the-building-wave-of-internet-traffic.html>
- [6] <https://www.google.com/covid19/mobility>
- [7] A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. Smaragdakis. 2020. The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic. In Proceedings of IMC’20: ACM Internet Measurement Conference.