Internet Measurements

From IPv6 Scanning to the COVID-19 Pandemic

Oliver Gasser

Max Planck Institute for Informatics · Internet Architecture research group

mpii

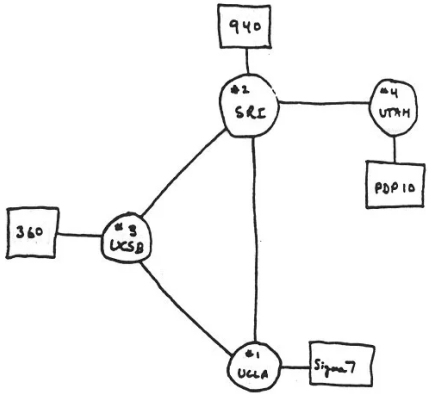

THE ARPA NETWORK

DEC 1969

4 NODES

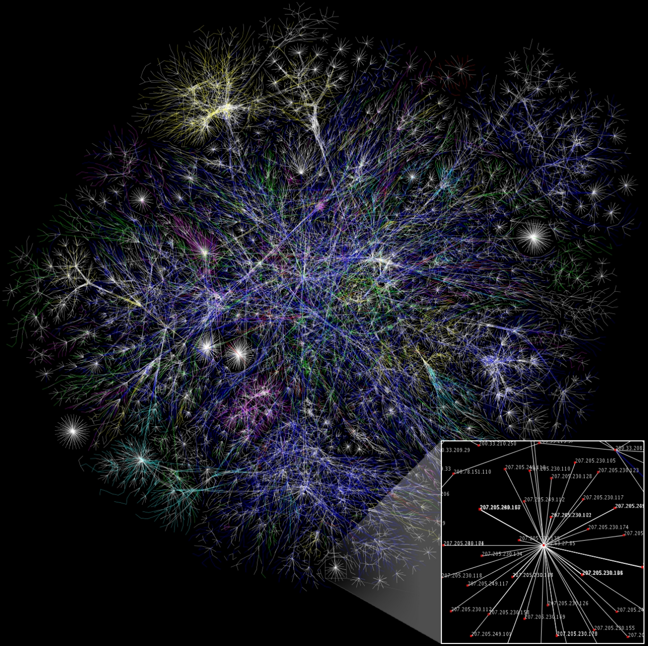
200.33.210.250
208.78.151.110
207.205.230.105
207.205.230.110
207.205.230.128
207.205.230.123
207.205.249.112
207.205.230.117
207.205.230.174
207.205.230.134
207.205.230.118
207.205.249.117
207.205.230.126
207.205.230.112
207.205.230.169
207.205.230.155
207.205.249.101

Internet Measurements

Why Internet measurements?

Analyze deployment of old and new protocols

Why Internet measurements?

Analyze deployment of old and new protocols

Evaluate performance, resilience, security,…

Why Internet measurements?

Analyze deployment of old and new protocols

Evaluate performance, resilience, security,…

Better understand the current state of the Internet

Types of Internet measurements

Active

Types of Internet measurements

Active

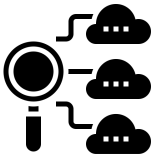

Passive

Active measurements

- **Actively send** probe packets to target hosts

Active measurements

- **Actively send** probe packets to target hosts
- Analyze responses to draw conclusions

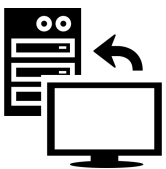

Active measurements

- **Actively send** probe packets to target hosts
- Analyze responses to draw conclusions
- Analogy: *Travel to a country to learn more about it*

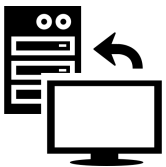

Passive measurements

- **Passively observe** traffic in the network

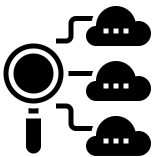

Passive measurements

- **Passively observe** traffic in the network
- Analyze traffic to draw conclusions

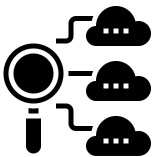

Passive measurements

- **Passively observe** traffic in the network
- Analyze traffic to draw conclusions
- Analogy: *Watch a documentary to learn more about a country*

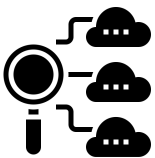

Two measurement projects

1. The Lockdown Effect
2. IPv6 Hitlist

Two measurement projects

1. The Lockdown Effect

2. IPv6 Hitlist

The Lockdown Effect

COVID-19 and the Internet

euronews.

Coronavirus: Half of humanity now on lockdown as 90 countries call for confinement

The New York Times

Working From Home: How Coronavirus Could Affect the Workplace

INSIDE HIGHER ED

Will Shift to Remote Teaching Be Boon or Bane for Online Learning?

REUTERS

Under lockdown, Italy's social and family life goes virtual

COVID-19 and the Internet

euronews.

Coronavirus: Half of humanity now on lockdown as 90 countries call for confinement

The New York Times

Working From Home: How Coronavirus Could Affect the Workplace

INSIDE HIGHER ED

Will Shift to Remote Teaching Be Boon or Bane for Online Learning?

REUTERS

Under lockdown, Italy's social and family life goes virtual

The Internet is essential in all these efforts, but how well does it cope?

Lots of passive measurement data

Lots of passive measurement data

3 IXPs

IXP Central Europe
IXP Southern Europe
IXP US East Coast

Lots of passive measurement data

3 IXPs

IXP Central Europe
IXP Southern Europe
IXP US East Coast

Interconnecting networks

Lots of passive measurement data

3 IXPs

IXP Central Europe
IXP Southern Europe
IXP US East Coast

Interconnecting networks

ISP

Central Europe

Lots of passive measurement data

3 IXPs

IXP Central Europe
IXP Southern Europe
IXP US East Coast

Interconnecting networks

ISP

Central Europe

Residential customers working from home

Lots of passive measurement data

3 IXPs

IXP Central Europe
IXP Southern Europe
IXP US East Coast

Interconnecting networks

ISP

Central Europe

Residential customers working from home

EDU

Madrid region

Lots of passive measurement data

3 IXPs

IXP Central Europe
IXP Southern Europe
IXP US East Coast

Interconnecting networks

ISP

Central Europe

Residential customers working from home

EDU

Madrid region

Service network interconnecting universities and research institutions

Traffic changes in different networks

Traffic changes: Jan. ’19 – Jun. ’20

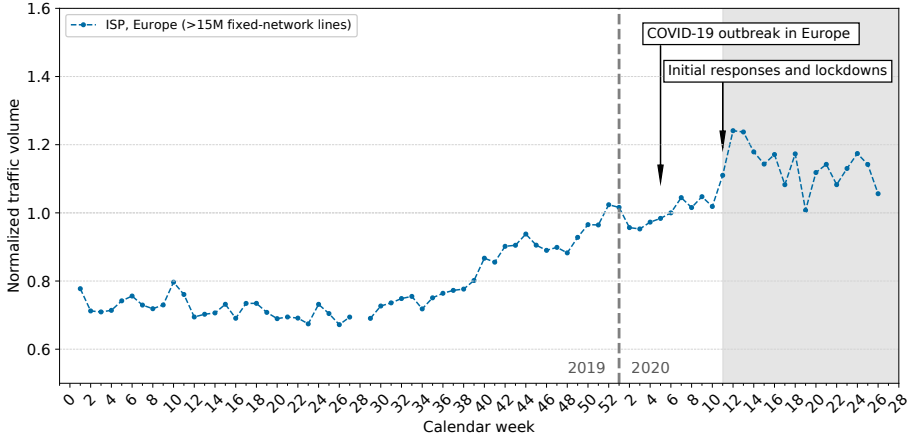

Traffic changes: Jan. ’19 – Jun. ’20

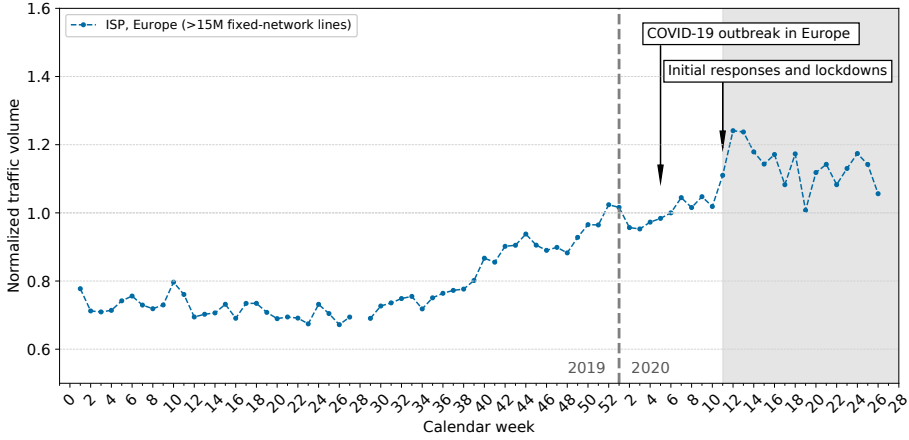

Once the lockdown started the ISP saw a **+30% traffic increase.**

Traffic changes: Jan. ’19 – Jun. ’20

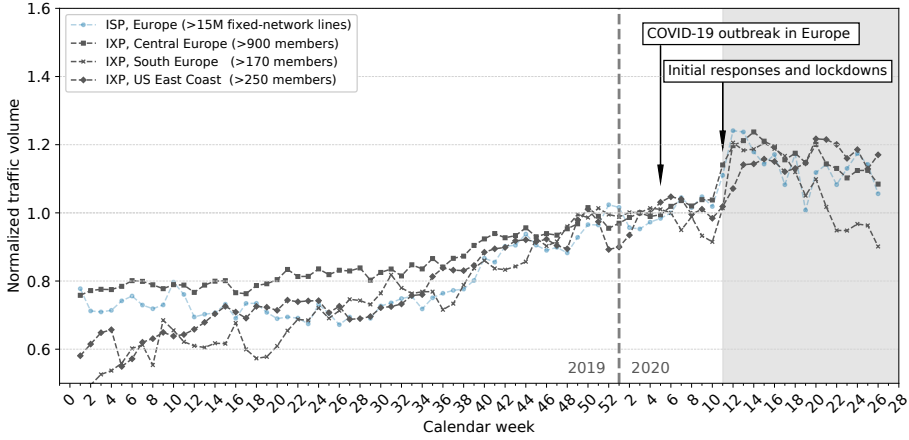

Similar behavior for the three IXPs.

Traffic changes: Jan. ’19 – Jun. ’20

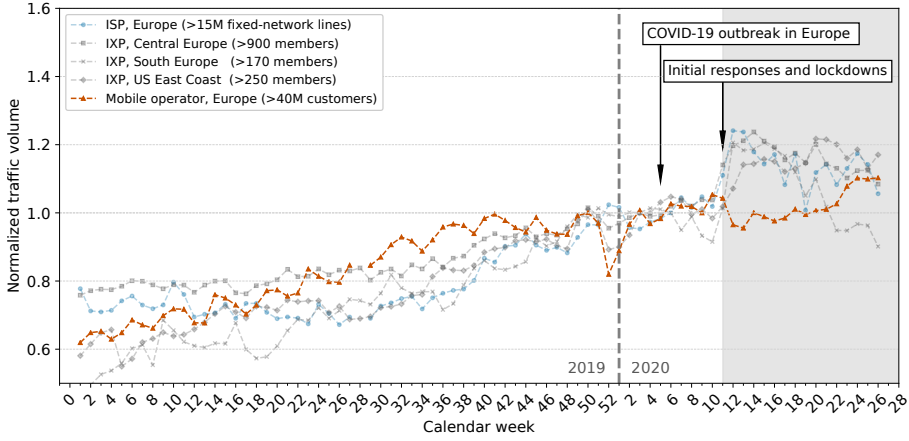

Mobile traffic decreased measurably after the lockdown.

Workday vs. weekend patterns

Workday vs. weekend patterns

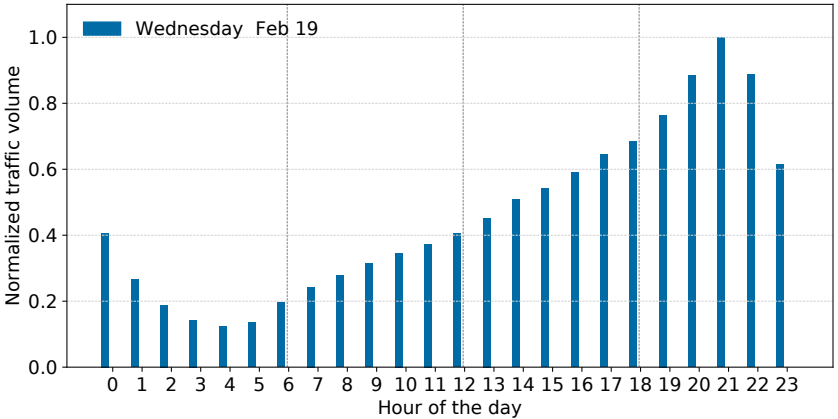

Workday vs. weekend patterns

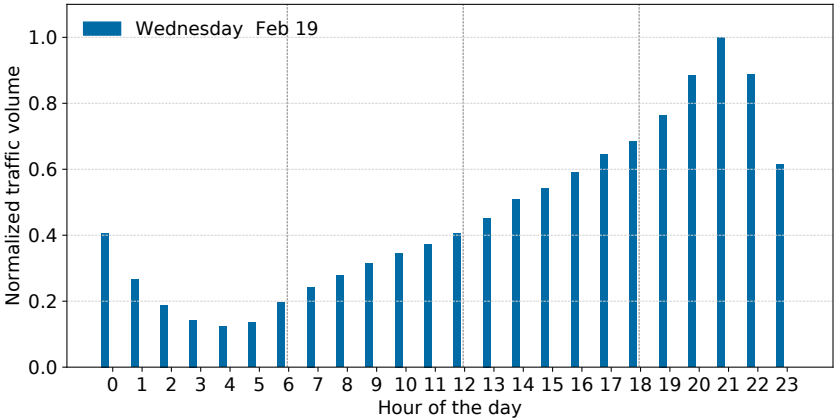

Regular workday: Strong increase in **evening hours**.

Workday vs. weekend patterns

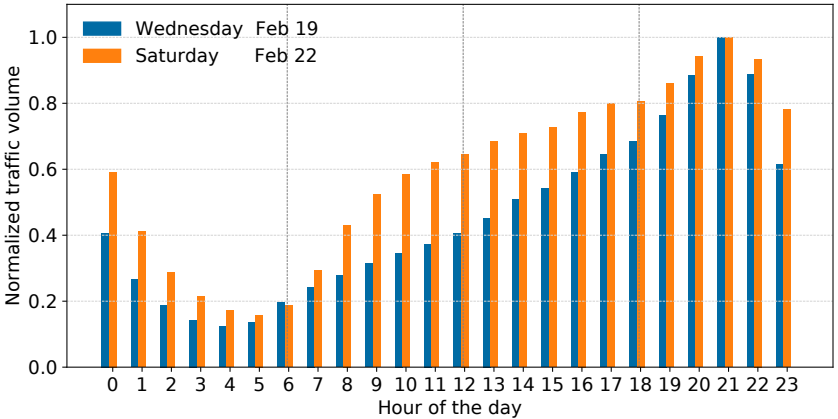

Regular weekend: More traffic during **daytime**.

Workday vs. weekend patterns

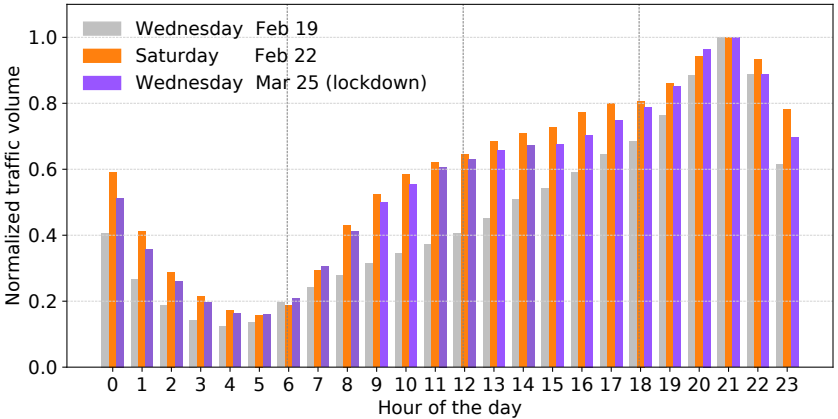

During lockdown: **Workdays look more like weekends.**

Workday vs. weekend patterns

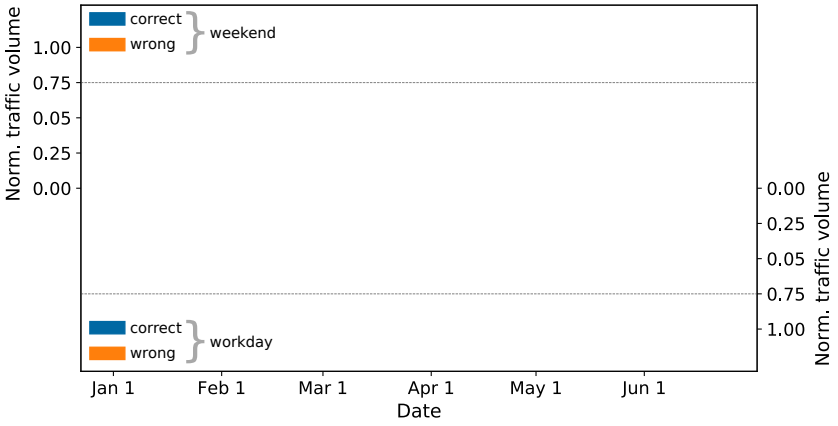

Classify days into workdays or weekends using traffic patterns.

Workday vs. weekend patterns

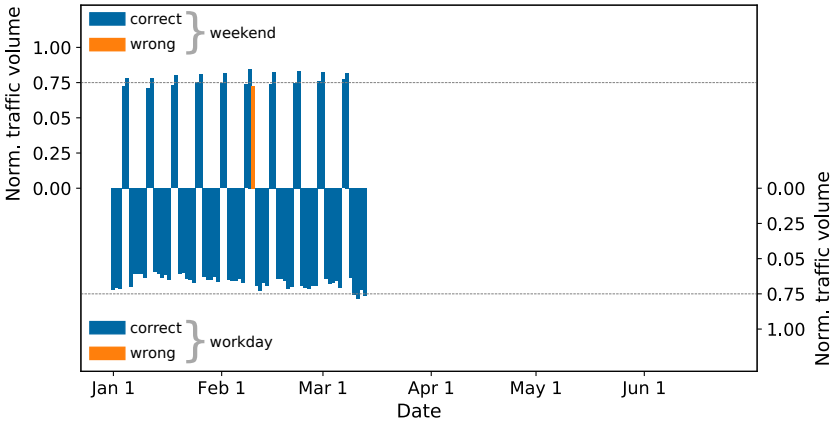

Pre-lockdown: Most days are classified **correctly**.

Workday vs. weekend patterns

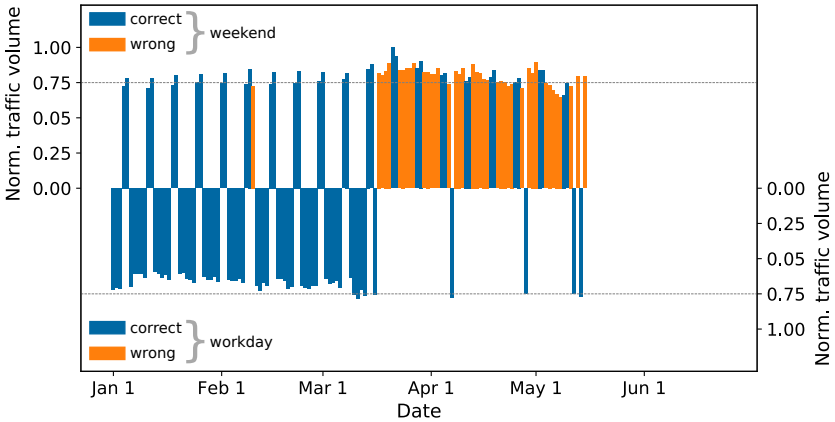

During lockdown: **Workdays are classified as weekends**.

Workday vs. weekend patterns

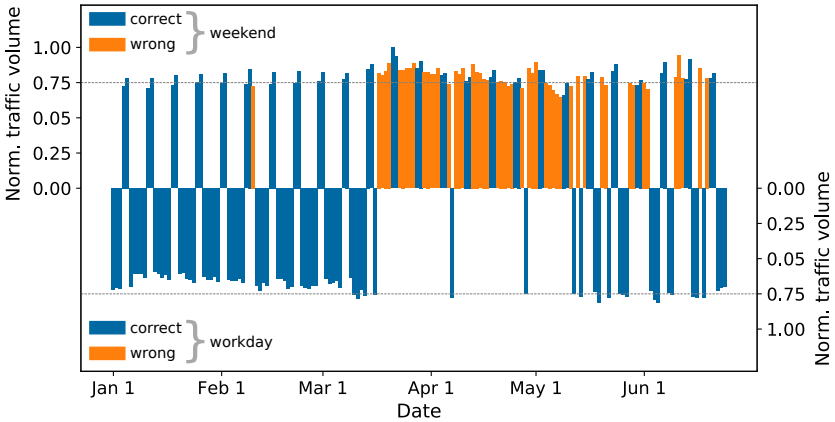

During lockdown: **Workdays are classified as weekends**, recovering mid-May.

Application-level traffic changes

Classify traffic by application class

Classify traffic by application class

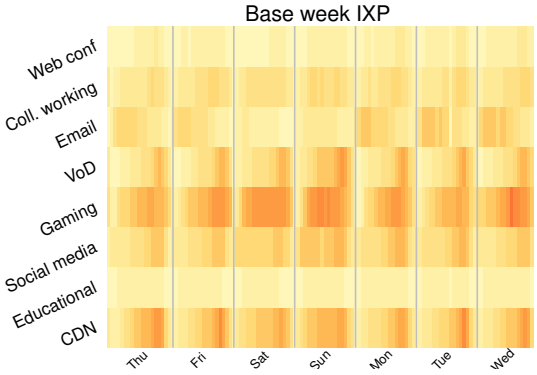

Classify traffic by application class

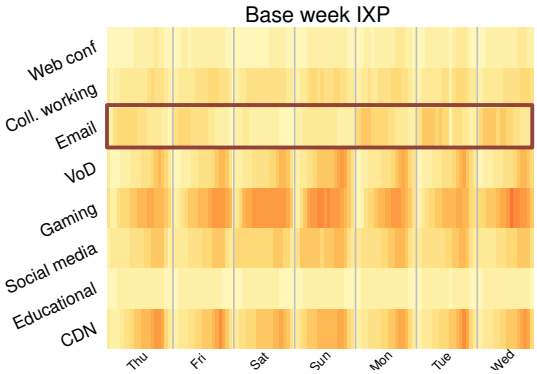

Email during working hours

Classify traffic by application class

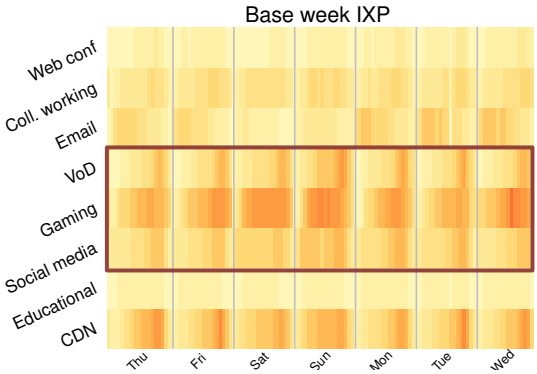

Video, gaming, and social media during evening hours

Classify traffic by application class

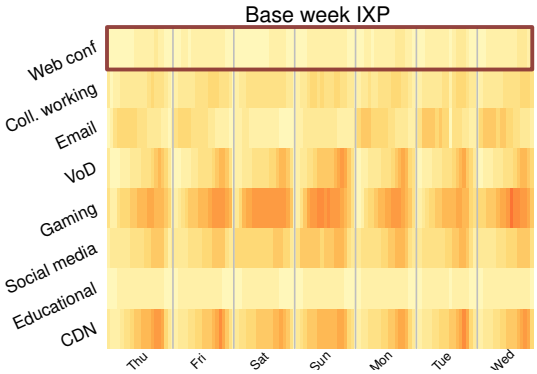

Hardly any web conferencing

Changes in application classes

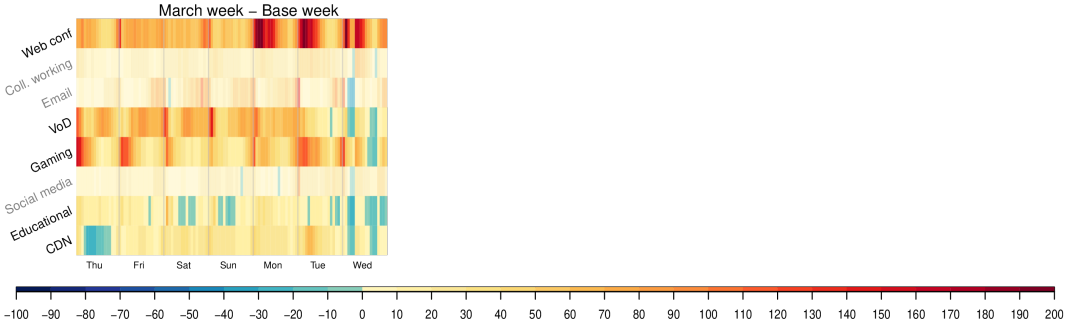

Changes in application classes

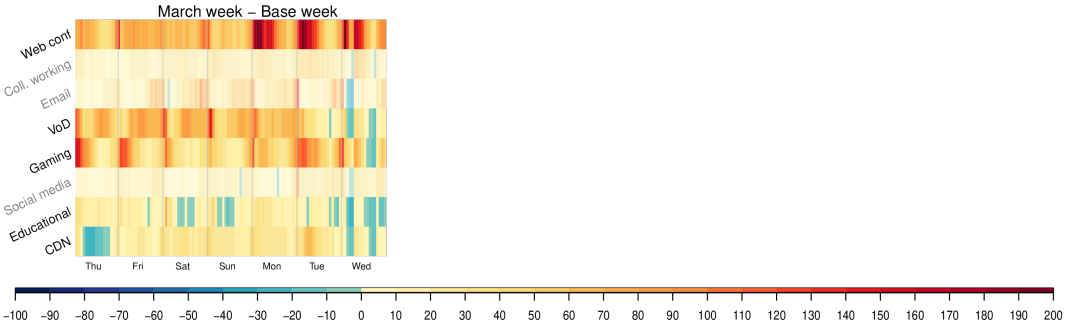

March: Increase in web conferencing, video, and gaming

Changes in application classes

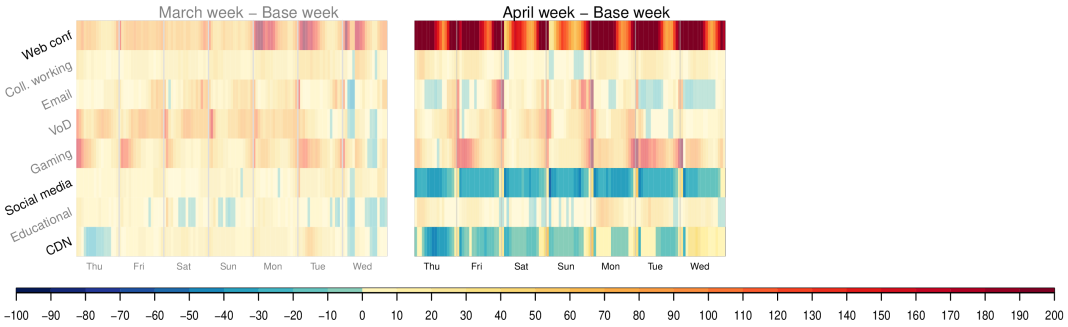

April: Strong increase in web conferencing, decrease in social media and CDN

Changes in application classes

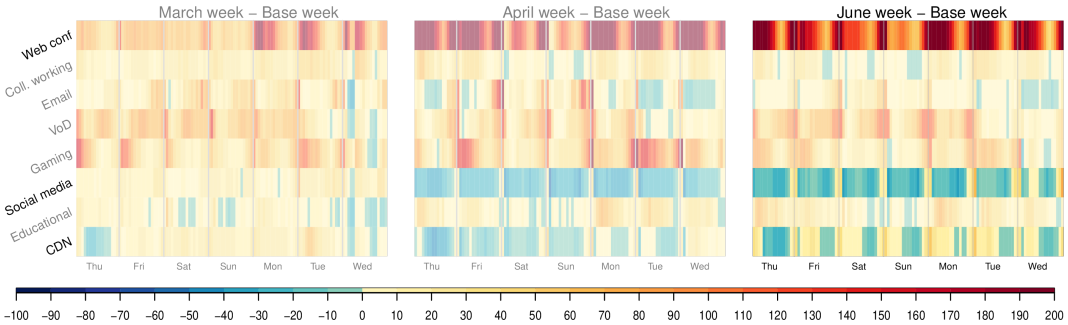

April & June: Strong increase in web conferencing, decrease in social media and CDN

VPN traffic

VPN traffic

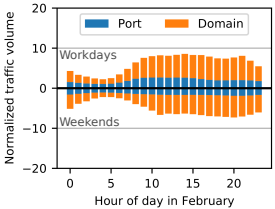
Port
Domain
Workdays
Weekends
Normalized traffic volume
Hour of day in February
20
10
0
−10
−20
0
5
10
15
20

VPN traffic

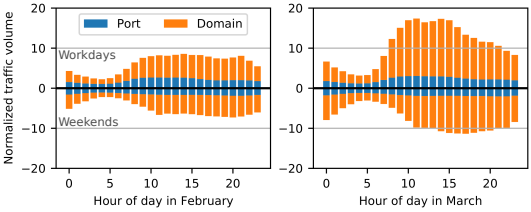

- **200% increase** in VPN traffic in March during working hours

VPN traffic

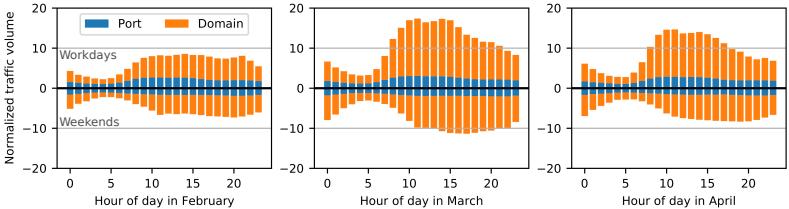

- **200% increase** in VPN traffic in March during working hours
- Slight decrease in April

VPN traffic

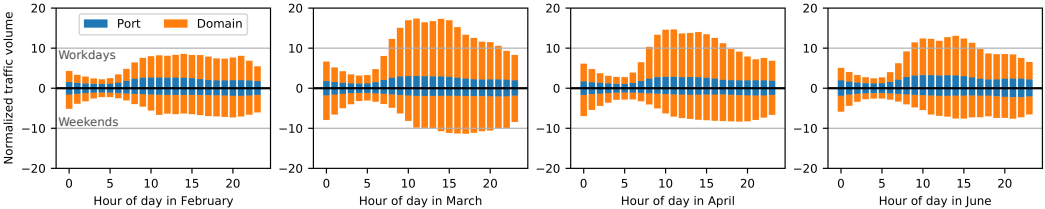

- **200% increase** in VPN traffic in March during working hours

- Slight decrease in April & June

How did educational traffic change?

How did edu traffic change?

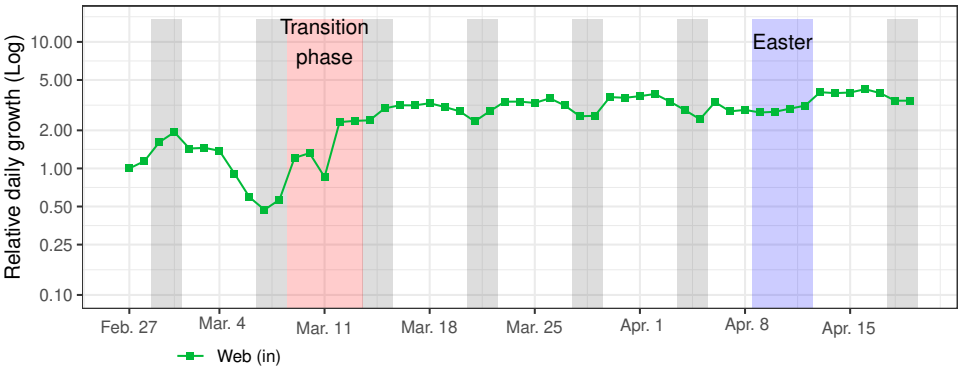

How did edu traffic change?

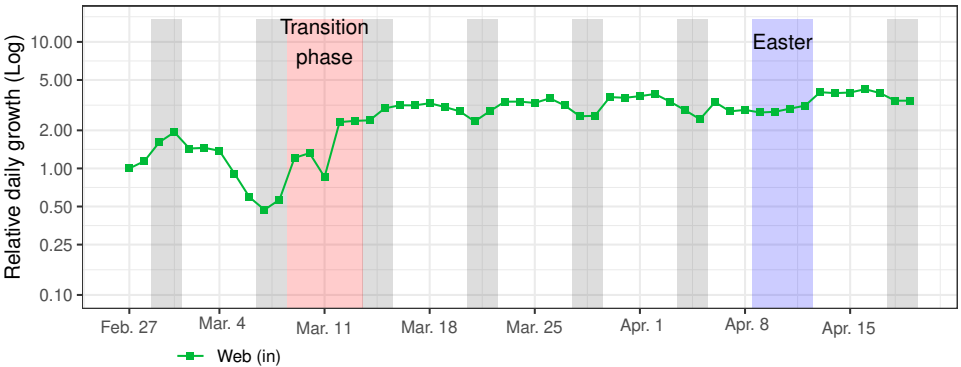

- Increase in incoming web traffic

How did edu traffic change?

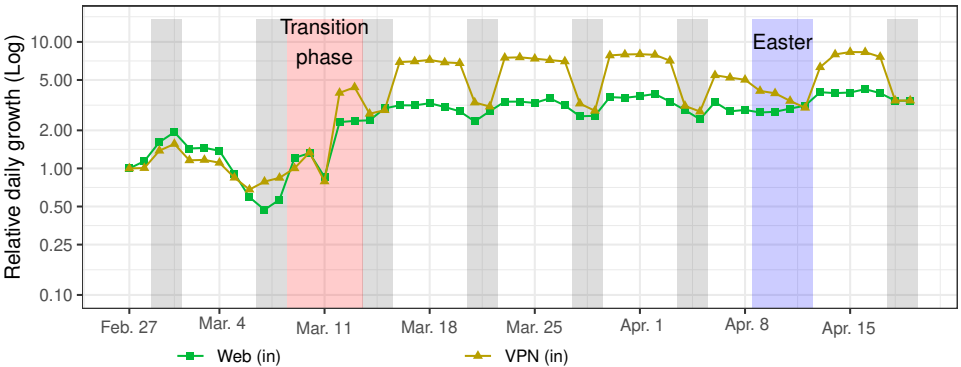

- Increase in incoming web and VPN traffic

How did edu traffic change?

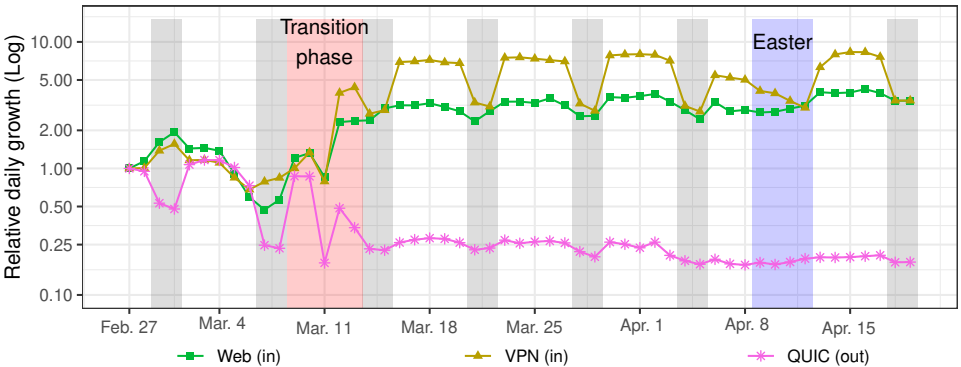

- Increase in incoming web and VPN traffic
- **Decrease** of outgoing QUIC traffic

What about other networks?

RIPE 81 polls

- RIPE 81: Meeting of European network operators

RIPE 81 polls

- RIPE 81: Meeting of European network operators
- October 27, 2020

RIPE 81 polls

- RIPE 81: Meeting of European network operators
- October 27, 2020
- Operators answered polls in the plenary

Traffic changes in their networks

Traffic changes in their networks

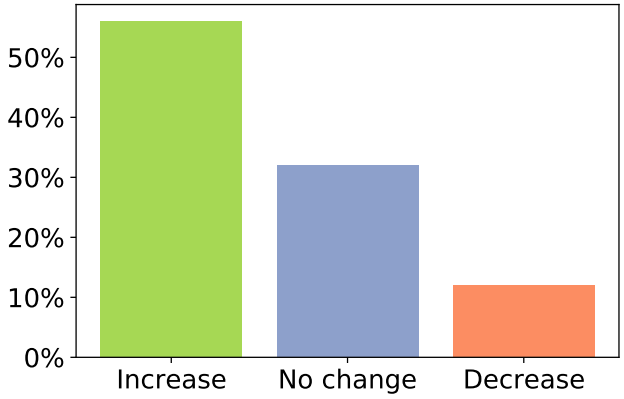

Traffic changes in their networks

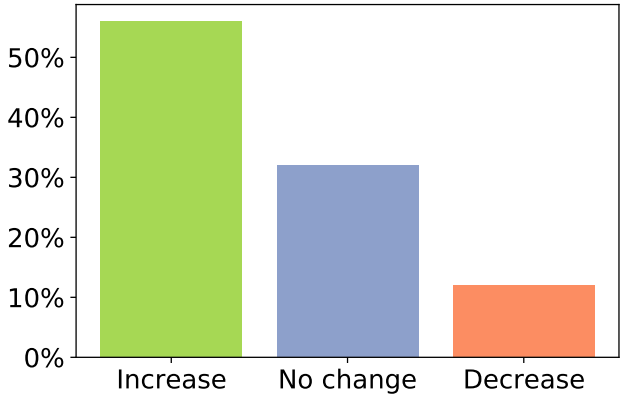

- Majority of networks see a **traffic increase**

Preparedness for 2nd lockdown

Preparedness for 2nd lockdown

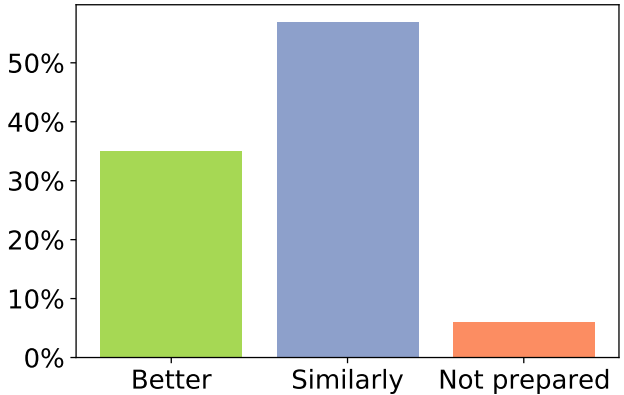

Preparedness for 2nd lockdown

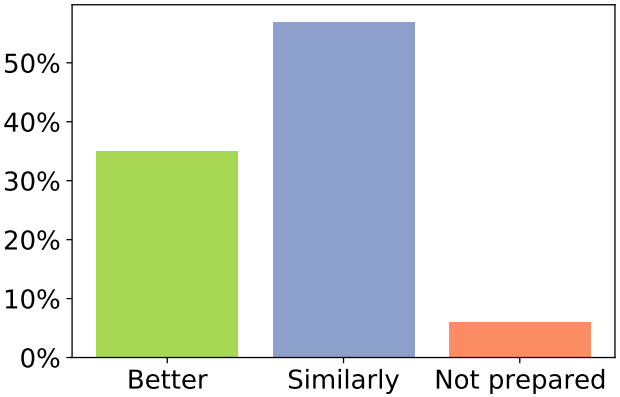

- One third are **better prepared**

What we found

People change → traffic changes

People change → traffic changes

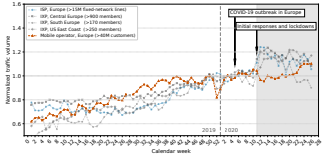

Traffic increase of **15-30%**

People change → traffic changes

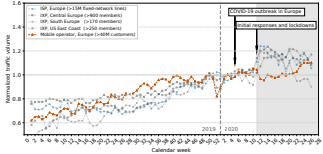

Traffic increase of **15-30%**

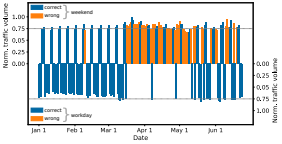

Workdays look like **weekends**

People change → traffic changes

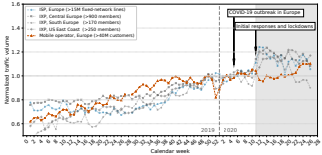

Traffic increase of **15-30%**

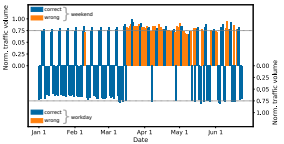

Workdays look like **weekends**

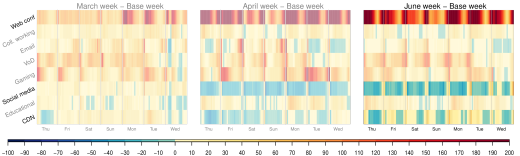

Increase in
remote working traffic

People change → traffic changes

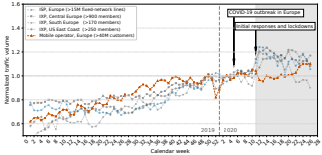

Traffic increase of **15-30%**

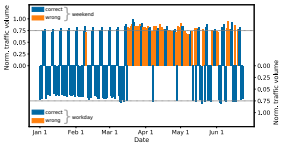

Workdays look like **weekends**

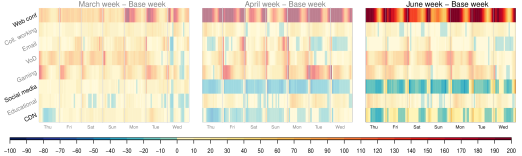

Increase in
remote working traffic

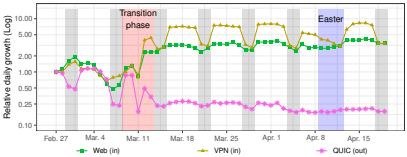

Decrease due to
absence of users

Two measurement projects

1. The Lockdown Effect
2. IPv6 Hitlist

Two measurement projects

1. The Lockdown Effect

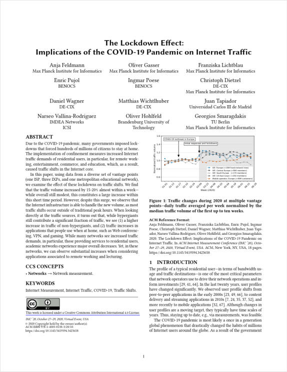

The Lockdown Effect:
Implications of the COVID-19 Pandemic on Internet Traffic

Anja Feldmann, Max Planck Institute for Informatics
Oliver Gasser, Max Planck Institute for Informatics
Franziska Lichtblau, Max Planck Institute for Informatics
Enric Pujol, BENOCS
Ingmar Poese, BENOCS
Christoph Dietzel, DE-CIX, Max Planck Institute for Informatics
Daniel Wagner, DE-CIX
Matthias Wichtlhuber, DE-CIX
Juan Tapiador, Universidad Carlos III de Madrid
Narseo Vallina-Rodriguez, IMDEA Networks, ICSI
Oliver Hohlfeld, Brandenburg University of Technology
Georgios Smaragdakis, TU Berlin, Max Planck Institute for Informatics

ABSTRACT

Due to the COVID-19 pandemic, many governments imposed lockdowns that forced hundreds of millions of citizens to stay at home. The implementation of confinement measures increased Internet traffic demands of residential users, in particular, for remote working, entertainment, commerce, and education, which, as a result, caused traffic shifts in the Internet core.

In this paper, using data from a diverse set of vantage points (one ISP, three IXPs, and one metropolitan educational network), we examine the effect of these lockdowns on traffic shifts. We find that the traffic volume increased by 15-20% almost within a week—while overall still modest, this constitutes a large increase within this short time period. However, despite this surge, we observe that the Internet infrastructure is able to handle the new volume, as most traffic shifts occur outside of traditional peak hours. When looking directly at the traffic sources, it turns out that, while hypergiants still contribute a significant fraction of traffic, we see (1) a higher increase in traffic of non-hypergiants, and (2) traffic increases in applications that people use when at home, such as Web conferencing, VPN, and gaming. While many networks see increased traffic demands, in particular, those providing services to residential users, academic networks experience major overall decreases. Yet, in these networks, we can observe substantial increases when considering applications associated to remote working and lecturing.

CCS CONCEPTS
• **Networks** → Network measurement.

KEYWORDS
Internet Measurement, Internet Traffic, COVID-19, Traffic Shifts.

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/10.1145/3419394.3423658

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.

ACM Reference Format:
Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis. 2020. The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic. In *ACM Internet Measurement Conference (IMC '20), October 27–29, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3419394.3423658

1 INTRODUCTION

The profile of a typical residential user—in terms of bandwidth usage and traffic destinations—is one of the most critical parameters that network operators use to drive their network operations and inform investments [29, 41, 64]. In the last twenty years, user profiles have changed significantly. We observed user profile shifts from peer-to-peer applications in the early 2000s [23, 49, 66], to content delivery and streaming applications in 2010s [7, 24, 35, 37, 52], and more recently to mobile applications [32, 67]. Although changes in user profiles are a moving target, they typically have time scales of years. Thus, staying up to date, e.g., via measurements, was feasible.

The COVID-19 pandemic is most likely a once in a generation global phenomenon that drastically changed the habits of millions of Internet users around the globe. As a result of the government

1

IMC 2020: *The Lockdown Effect*

2. IPv6 Hitlist

Two measurement projects

1. The Lockdown Effect

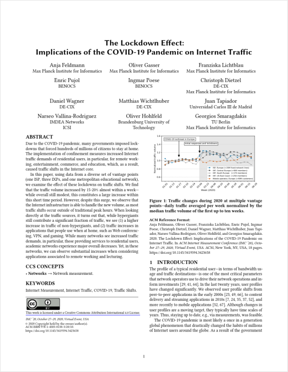

The Lockdown Effect:
Implications of the COVID-19 Pandemic on Internet Traffic

Anja Feldmann, Max Planck Institute for Informatics
Oliver Gasser, Max Planck Institute for Informatics
Franziska Lichtblau, Max Planck Institute for Informatics
Enric Pujol, BENOCS
Ingmar Poese, BENOCS
Christoph Dietzel, DE-CIX, Max Planck Institute for Informatics
Daniel Wagner, DE-CIX
Matthias Wichtlhuber, DE-CIX
Juan Tapiador, Universidad Carlos III de Madrid
Narseo Vallina-Rodriguez, IMDEA Networks, ICSI
Oliver Hohlfeld, Brandenburg University of Technology
Georgios Smaragdakis, TU Berlin, Max Planck Institute for Informatics

ABSTRACT

Due to the COVID-19 pandemic, many governments imposed lockdowns that forced hundreds of millions of citizens to stay at home. The implementation of confinement measures increased Internet traffic demands of residential users, in particular, for remote working, entertainment, commerce, and education, which, as a result, caused traffic shifts in the Internet core.

In this paper, using data from a diverse set of vantage points (one ISP, three IXPs, and one metropolitan educational network), we examine the effect of these lockdowns on traffic shifts. We find that the traffic volume increased by 15-20% almost within a week—while overall still modest, this constitutes a large increase within this short time period. However, despite this surge, we observe that the Internet infrastructure is able to handle the new volume, as most traffic shifts occur outside of traditional peak hours. When looking directly at the traffic sources, it turns out that, while hypergiants still contribute a significant fraction of traffic, we see (1) a higher increase in traffic of non-hypergiants, and (2) traffic increases in applications that people use when at home, such as Web conferencing, VPN, and gaming. While many networks see increased traffic demands, in particular, those providing services to residential users, academic networks experience major overall decreases. Yet, in these networks, we can observe substantial increases when considering applications associated to remote working and lecturing.

CCS CONCEPTS

• **Networks** → Network measurement.

KEYWORDS

Internet Measurement, Internet Traffic, COVID-19, Traffic Shifts.

IMC '20, October 27–29, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-8138-3/20/10.
https://doi.org/10.1145/3419394.3423658

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.

ACM Reference Format:
Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis. 2020. The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic. In *ACM Internet Measurement Conference (IMC '20), October 27–29, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3419394.3423658

1 INTRODUCTION

The profile of a typical residential user—in terms of bandwidth usage and traffic destinations—is one of the most critical parameters that network operators use to drive their network operations and inform investments [29, 41, 64]. In the last twenty years, user profiles have changed significantly. We observed user profile shifts from peer-to-peer applications in the early 2000s [23, 49, 66], to content delivery and streaming applications in 2010s [7, 24, 35, 37, 52], and more recently to mobile applications [32, 67]. Although changes in user profiles are a moving target, they typically have time scales of years. Thus, staying up to date, e.g., via measurements, was feasible.

The COVID-19 pandemic is most likely a once in a generation global phenomenon that drastically changed the habits of millions of Internet users around the globe. As a result of the government

1

IMC 2020: *The Lockdown Effect*

2. IPv6 Hitlist

Technical University of Munich TUM

IITiS

TU Delft

UNIVERSITÉ Grenoble Alpes

RIPE NCC
RIPE NETWORK COORDINATION CENTRE

UNIVERSITY OF TWENTE.

IPv6 Hitlist

Target selection

Active Measurements

- **Actively send** probe packets to target hosts
- Analyze responses to draw conclusions
- Analog: *Travel to a country to learn more about it*

7

Target selection

Active Measurements

- **Actively send** probe packets to target hosts
- Analyze responses to draw conclusions
- Analog: *Travel to a country to learn more about it*

7

But how do we select the **targets**?

IPv4

- Fast measurement tools allow IPv4-wide scans of $\approx$ 4 billion addresses
- Completes within days or hours
- Fire and forget

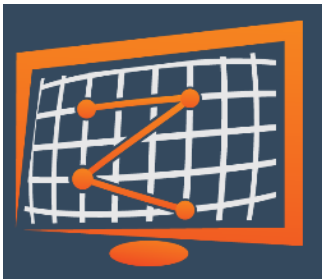

Is IPv4 enough?

Is IPv4 enough?

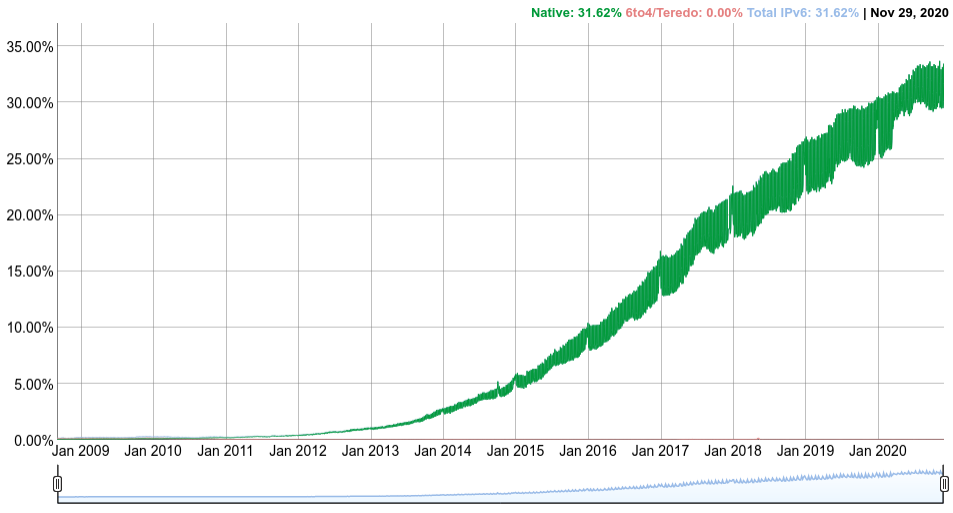

Is IPv4 enough?

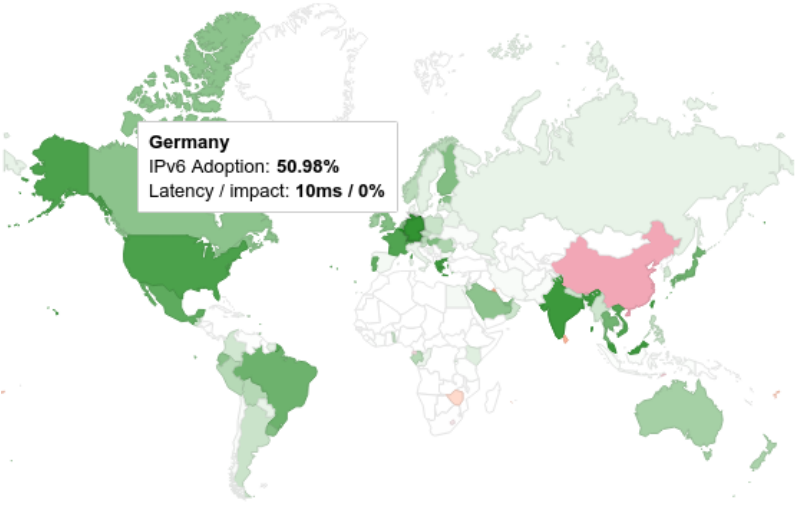

IPv6

IPv6

Vast address space:

10^{38} addresses

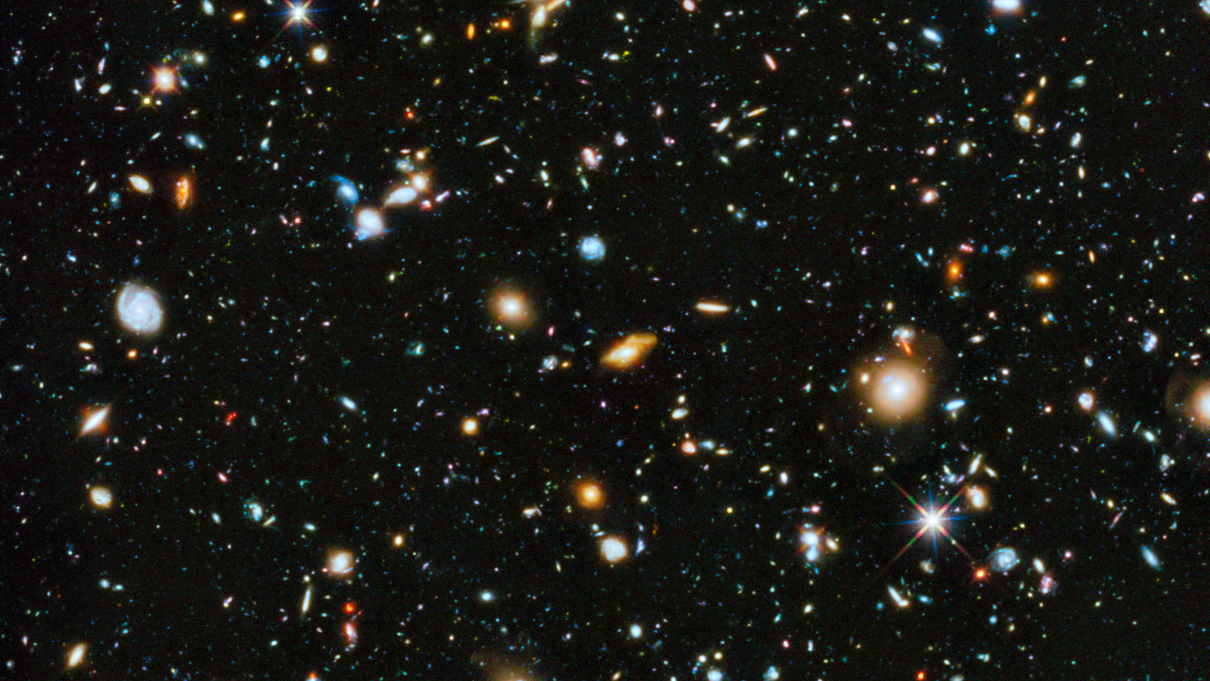


Estimated number of stars in the universe:
10^{24} stars

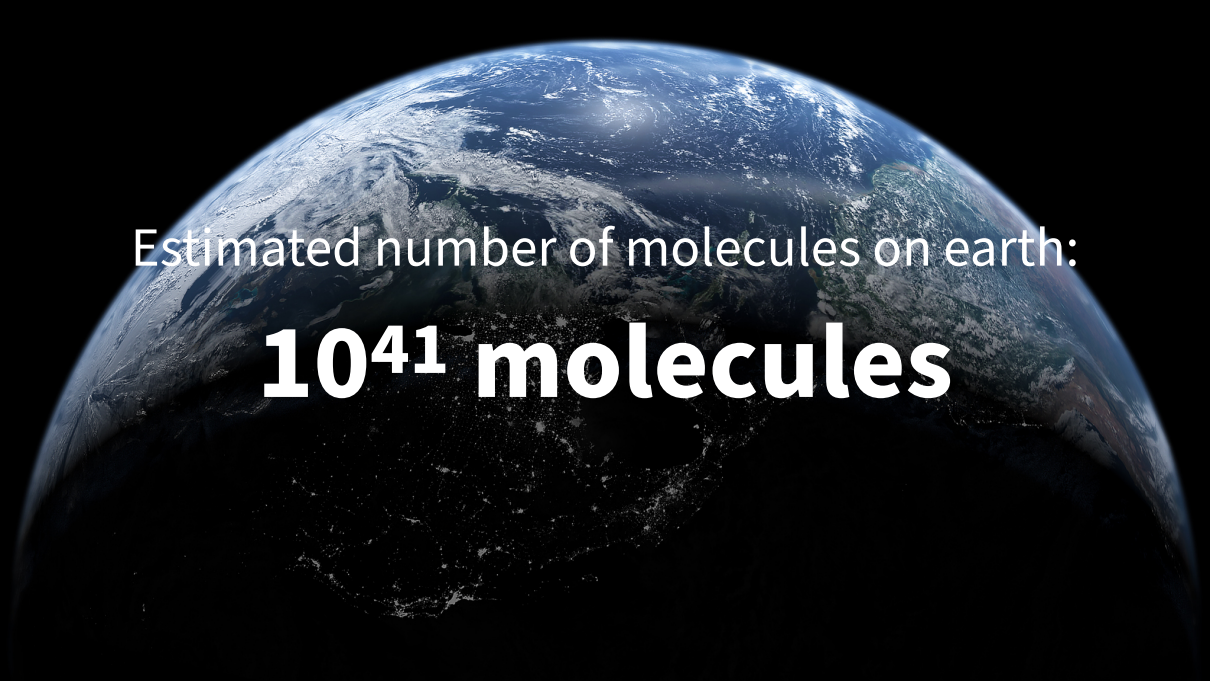
Estimated number of molecules on earth:
10^{41} molecules

IPv6 hitlist

IPv6 hitlist

IPv6 hitlist

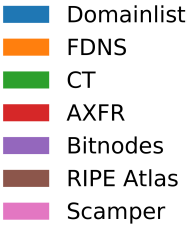

IPv6 hitlist

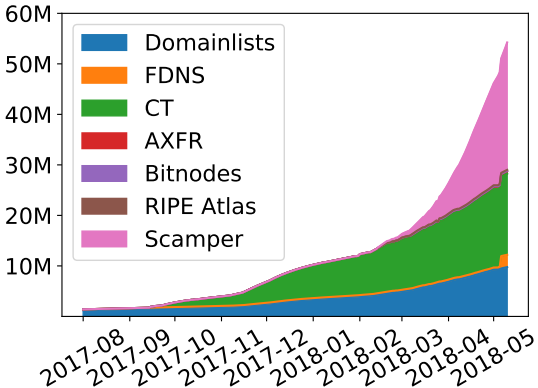

- Many addresses from domainlists, Certificate Transparency, and traceroute

Hitlist biases

Hitlist biases

Biased hitlists can lead to **biased measurement results**

Hitlist biases

Biased hitlists can lead to **biased measurement results**

- Balancedness of sources
- Aliased prefixes

Balancedness of sources

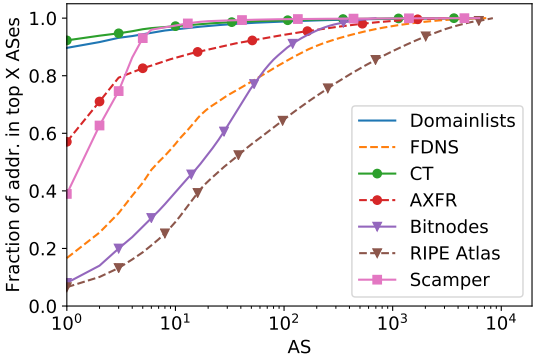

Balancedness of sources

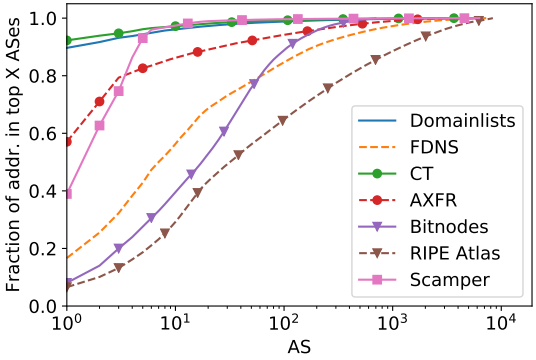

- Domainlists, CT quite unbalanced
- RIPE Atlas and Bitnodes more balanced

Aliased prefixes

- Alias: another address of the same host (`2001:db8:8:4::1`, `2001:db8:8:4::2`)

Aliased prefixes

- Alias: another address of the same host (`2001:db8:8:4::1`, `2001:db8:8:4::2`)
- Aliased prefix: **whole prefix** bound to the **same host** (`2001:db8:8:4::/64`)

Aliased prefixes

- Alias: another address of the same host (`2001:db8:8:4::1`, `2001:db8:8:4::2`)
- Aliased prefix: **whole prefix** bound to the **same host** (`2001:db8:8:4::/64`)
- Bias: some hosts overrepresented due to aliased prefixes

Aliased prefixes

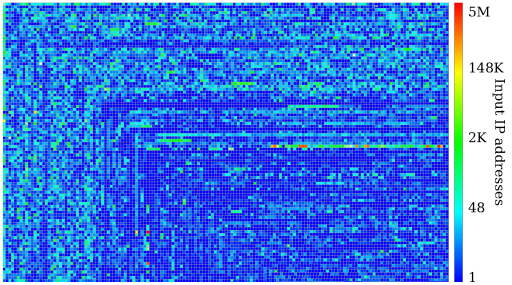

All prefixes covered by hitlist.

Aliased prefixes

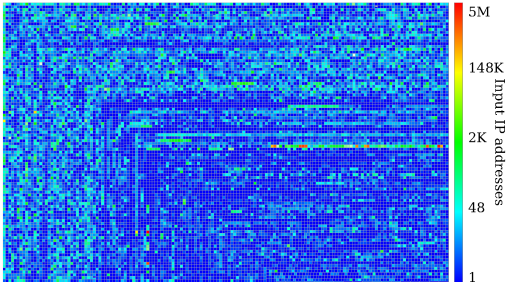

All prefixes covered by hitlist.

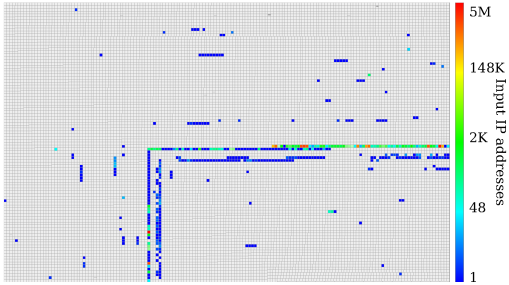

Aliased prefixes.

Aliased prefixes

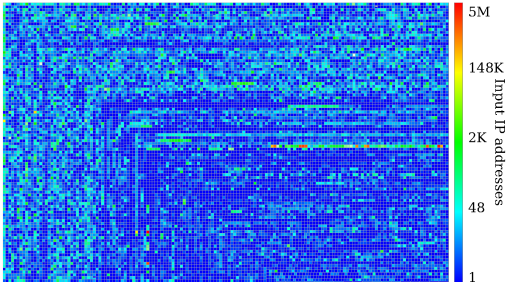

All prefixes covered by hitlist.

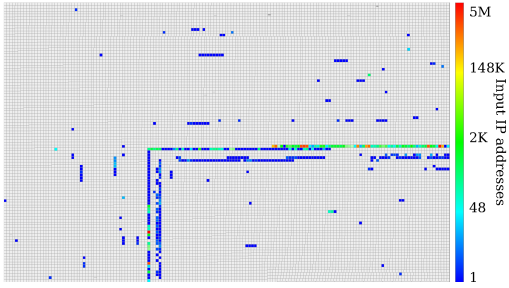

Aliased prefixes.

- Only 3.2% of prefixes are aliased
- But **46.6% of addresses** are in aliased prefixes

IPv6 Hitlist Service

ipv6hitlist.github.io

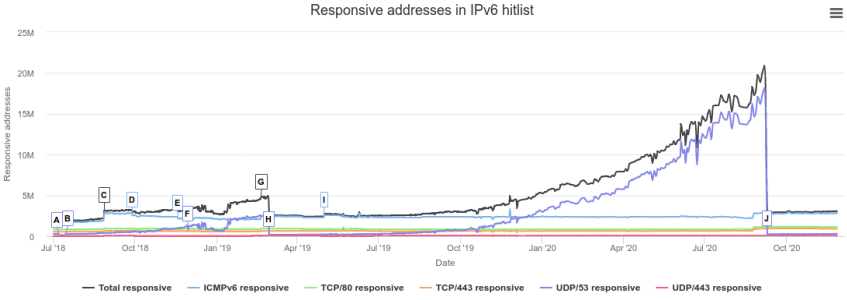

ipv6hitlist.github.io

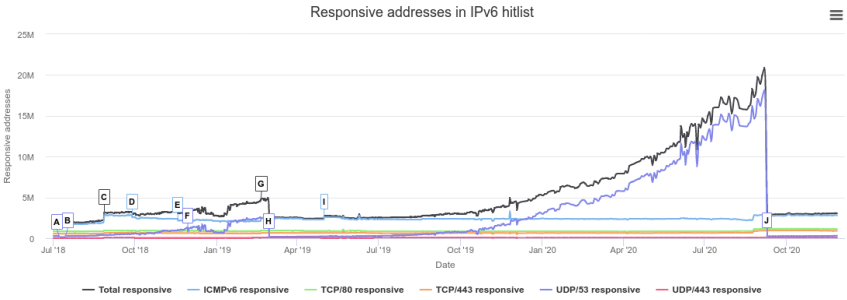

- Daily IPv6 hitlists and aliased prefixes

ipv6hitlist.github.io

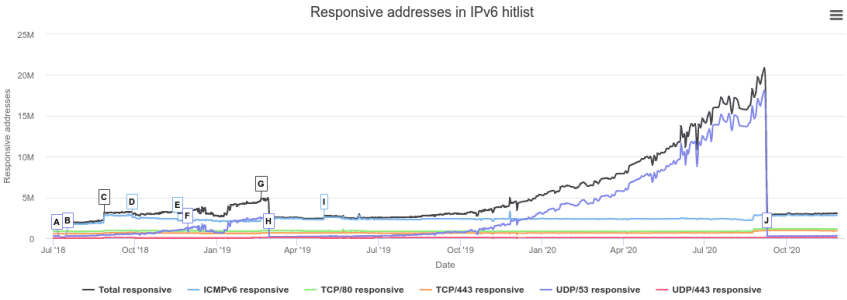

- Daily IPv6 hitlists and aliased prefixes
- Continuously running since July 2018

ipv6hitlist.github.io

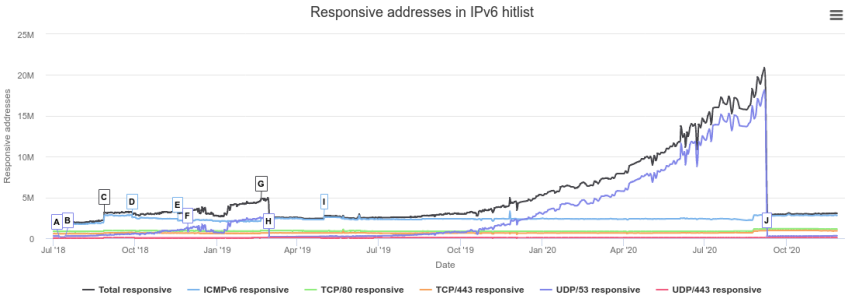

- Daily IPv6 hitlists and aliased prefixes
- Continuously running since July 2018
- Used by dozens of researchers

Users of the IPv6 Hitlist

Users of the IPv6 Hitlist


The Closed Resolver Project:
Measuring the Deployment of Source Address
Validation of Internet Paths

Spoofing detection

Users of the IPv6 Hitlist

Spoofing detection

Host security

Users of the IPv6 Hitlist

Spoofing detection

Host security

DNS security

Users of the IPv6 Hitlist

Spoofing detection

Host security

DNS security

Load balancing

What we found

IPv6 hitlist

IPv6 hitlist

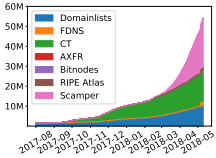

Many addresses from **different sources**

IPv6 hitlist

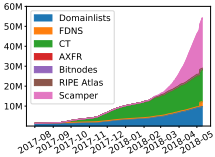

Many addresses from **different sources**

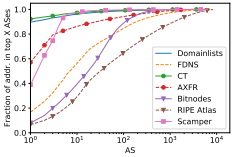

Domainlists and CT quite **unbalanced**

IPv6 hitlist

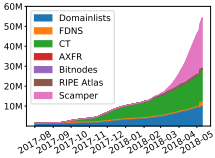

Many addresses from **different sources**

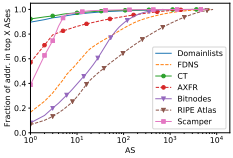

Domainlists and CT quite **unbalanced**

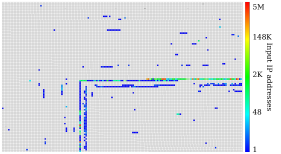

Beware of **aliased prefixes**

IPv6 hitlist

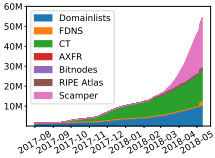

Many addresses from **different sources**

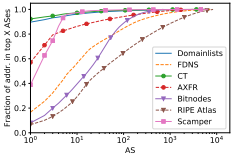

Domainlists and CT quite **unbalanced**

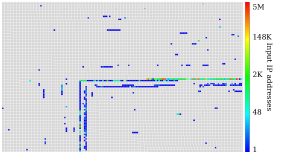

Beware of **aliased prefixes**

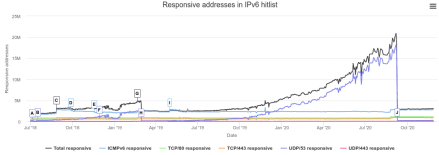

ipv6hitlist.github.io

Two measurement projects

1. The Lockdown Effect

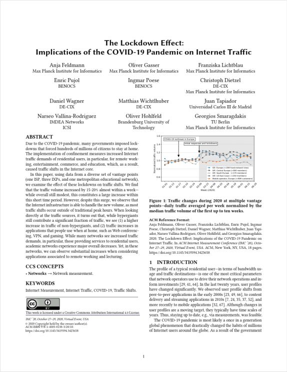

The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic

Anja Feldmann, Max Planck Institute for Informatics
Oliver Gasser, Max Planck Institute for Informatics
Franziska Lichtblau, Max Planck Institute for Informatics
Enric Pujol, BENOCS
Ingmar Poese, BENOCS
Christoph Dietzel, DE-CIX, Max Planck Institute for Informatics
Daniel Wagner, DE-CIX
Matthias Wichtlhuber, DE-CIX
Juan Tapiador, Universidad Carlos III de Madrid
Narseo Vallina-Rodriguez, IMDEA Networks, ICSI
Oliver Hohlfeld, Brandenburg University of Technology
Georgios Smaragdakis, TU Berlin, Max Planck Institute for Informatics

ABSTRACT

Due to the COVID-19 pandemic, many governments imposed lockdowns that forced hundreds of millions of citizens to stay at home. The implementation of confinement measures increased Internet traffic demands of residential users, in particular, for remote working, entertainment, commerce, and education, which, as a result, caused traffic shifts in the Internet core.

In this paper, using data from a diverse set of vantage points (one ISP, three IXPs, and one metropolitan educational network), we examine the effect of these lockdowns on traffic shifts. We find that the traffic volume increased by 15-20% almost within a week—while overall still modest, this constitutes a large increase within this short time period. However, despite this surge, we observe that the Internet infrastructure is able to handle the new volume, as most traffic shifts occur outside of traditional peak hours. When looking directly at the traffic sources, it turns out that, while hypergiants still contribute a significant fraction of traffic, we see (1) a higher increase in traffic of non-hypergiants, and (2) traffic increases in applications that people use when at home, such as Web conferencing, VPN, and gaming. While many networks see increased traffic demands, in particular, those providing services to residential users, academic networks experience major overall decreases. Yet, in these networks, we can observe substantial increases when considering applications associated to remote working and lecturing.

CCS CONCEPTS

• **Networks** → Network measurement.

KEYWORDS

Internet Measurement, Internet Traffic, COVID-19, Traffic Shifts.

IMC '20, October 27–29, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-8138-3/20/10.
https://doi.org/10.1145/3419394.3423658

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.

ACM Reference Format:
Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis. 2020. The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic. In *ACM Internet Measurement Conference (IMC '20), October 27–29, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3419394.3423658

1 INTRODUCTION

The profile of a typical residential user—in terms of bandwidth usage and traffic destinations—is one of the most critical parameters that network operators use to drive their network operations and inform investments [29, 41, 64]. In the last twenty years, user profiles have changed significantly. We observed user profile shifts from peer-to-peer applications in the early 2000s [23, 49, 66], to content delivery and streaming applications in 2010s [7, 24, 35, 37, 52], and more recently to mobile applications [32, 67]. Although changes in user profiles are a moving target, they typically have time scales of years. Thus, staying up to date, e.g., via measurements, was feasible.

The COVID-19 pandemic is most likely a once in a generation global phenomenon that drastically changed the habits of millions of Internet users around the globe. As a result of the government

1

IMC 2020: *The Lockdown Effect*

2. IPv6 Hitlist

Two measurement projects

1. The Lockdown Effect

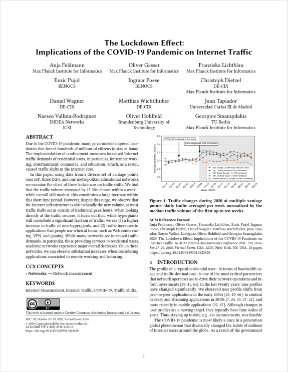

The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic

Anja Feldmann, Max Planck Institute for Informatics
Oliver Gasser, Max Planck Institute for Informatics
Franziska Lichtblau, Max Planck Institute for Informatics
Enric Pujol, BENOCS
Ingmar Poese, BENOCS
Christoph Dietzel, DE-CIX, Max Planck Institute for Informatics
Daniel Wagner, DE-CIX
Matthias Wichtlhuber, DE-CIX
Juan Tapiador, Universidad Carlos III de Madrid
Narseo Vallina-Rodriguez, IMDEA Networks, ICSI
Oliver Hohlfeld, Brandenburg University of Technology
Georgios Smaragdakis, TU Berlin, Max Planck Institute for Informatics

IMC 2020: *The Lockdown Effect*

2. IPv6 Hitlist

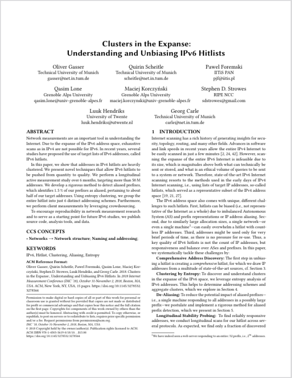

Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists

Oliver Gasser, Technical University of Munich, gasser@net.in.tum.de
Quirin Scheitle, Technical University of Munich, scheitle@net.in.tum.de
Pawel Foremski, IITiS PAN, pjf@iitis.pl
Qasim Lone, Grenoble Alps University, qasim.lone@univ-grenoble-alpes.fr
Maciej Korczyński, Grenoble Alps University, maciej.korczynski@univ-grenoble-alpes.fr
Stephen D. Strowes, RIPE NCC, sdstrowes@gmail.com
Luuk Hendriks, University of Twente, luuk.hendriks@utwente.nl
Georg Carle, Technical University of Munich, carle@net.in.tum.de

IMC 2018: *Clusters in the Expanse*

What’s next?

Current measurement projects

1. Multipath

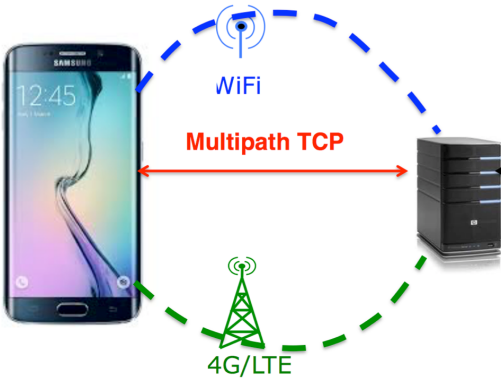

Current measurement projects

1. Multipath

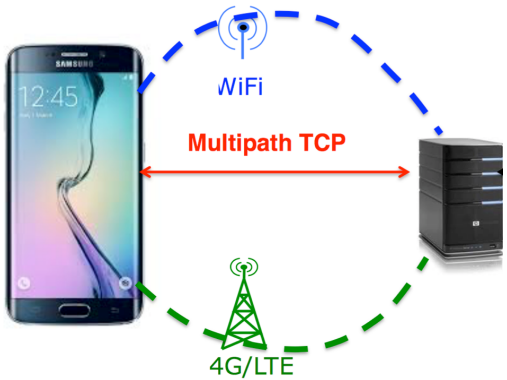

2. Anycast

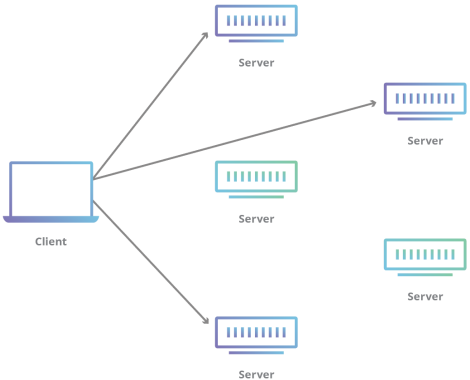

Internet Measurements

Conclusion

- **Internet measurements:** Important tool to quantify phenomena in the network

Conclusion

- **Internet measurements:** Important tool to quantify phenomena in the network
- **Lockdown effect**: Societal changes are directly reflected in the Internet

Conclusion

- **Internet measurements:** Important tool to quantify phenomena in the network
- **Lockdown effect**: Societal changes are directly reflected in the Internet
- **IPv6 hitlists**: Can be used to perform measurements in the vast IPv6 address space

Conclusion

- **Internet measurements:** Important tool to quantify phenomena in the network
- **Lockdown effect**: Societal changes are directly reflected in the Internet
- **IPv6 hitlists**: Can be used to perform measurements in the vast IPv6 address space

oliver.gasser@mpi-inf.mpg.de **mpi-inf.mpg.de/~ogasser**