

PingER and Covid-19 by Region

Introduction

This addresses the question of whether the [PingER](#) data can identify the impact of COVID-19 on the Internet. According to <https://thrivenextgen.com/covid-19-impact-on-internet-performance/> Most Internet Service Providers (ISPs) oversubscribe their bandwidth and networks as typical inbound and outbound traffic are bursty and often don't sustain high levels on a continuous basis. Fortunately, most of the larger National Providers have had sufficient time to prepare for the impact of COVID-19 and plan for the possibility of business workloads shifting to the home. Regardless of this capacity planning, Internet traffic patterns are about to change drastically based on schools across the country opting for virtual learning and business work from home migrations. Also, Time Magazine April 5/ April 15, 2020, reports that "traffic worldwide is up 35%...Demand is highest in the evening in the past two weeks, says networking firm Century Link."

For example, one might expect that as schools shut, people self-isolate, are sent home from, work from home, lockdowns are imposed etc., physical person-to-person communications would migrate to the Internet and will increase the use of the Internet e.g. by virtual learning, streaming, communicating, gathering information and entertaining. The thought is that this will lead to different Internet patterns and possibly lead to Internet congestion. A study of the impact of COVID-19 on Internet speeds by Ookla ("[TRACKING COVID-19'S IMPACT ON GLOBAL INTERNET PERFORMANCE](#)") is directly relevant, however, it is hard to identify any dramatic impact on the latencies in the report. There is also [COVID-19 impacts on Internet traffic: Seattle, Northern Italy, and South Korea](#) which looks at the impact on Internet traffic volumes. Also companies such as Netflix, Amazon, Apple, Disney, Google have reduced the quality of their videos to help reduce traffic and congestion of the Internet (see <https://www.traffic-masters.net/covid19-web-traffic-statistics/>).

Using [PingER data](#) and comparing the [various metrics](#) it provides including: average Round Trip Time (RTT), Conditional Loss Probability (CLP), Inter Packet Delay Variability (IPDV), Inter Quartile Range (IQR) of the round trip times, and derived throughput (see [Tutorial on Internet Monitoring and Pinger at SLAC](#)); we ascertained that the most stable and yet sensitive metric in detecting changes caused by Covid-19 interventions was the [Inter Packet Delay Variation](#) (IPDV). The data analyzed are for 120 days starting around mid to late January 16th and ending around mid to late May 2020. For each country, there is 1 point per day. Initially only weekday data was used in order to reduce the variability of the data and to focus more on the effect of interventions such as closing the workplaces, schools, universities, people working from home, or being out of work. The weekday was determined based on the Universal Time Coordinated (UTC).

The data was divided up by region in order to reduce the number of countries in a single chart. This is done so as to allow separation by eyeball of the IPDV lines for the various countries in a chosen region. It also keeps the data more self-consistent in terms of time zones, development, economy, customs etc. For each region, we show one or more charts of the daily median IPDV. The intervention data is from Wikipedia (e.g. https://en.wikipedia.org/wiki/COVID-19_pandemic_in_Egypt).

We are looking for a significant rise or fall in IPDV possibly correlated with an intervention or growth in cases. The links below point to charts of the IPDV (x-axis in msec.) versus the date. There is no data for 1/30/20 thru 2/2/20 (n.b. 2/1/30-2/2/30 are weekends).

Weekday data:

There are two Excel spreadsheets for the weekday filtered data (two files since the master ran out of tabs):

- [Master](#) most of the analysis for IPDV 120days weekday data
- [Overflow](#) for the remainder of the analysis

Alldays data

Since the weekday's data when plotted versus date left gaps that either resulted in gaps in the curves or extended the curve across a weekend thus making the peaks look much wider, it made the data harder to understand by simple visualization. Therefore, we decided to include all days.

Data by host

The analysis with the data aggregated by the country was generally unsuccessful ([see here](#)) in identifying changes in performance correlated with the interventions. This is partially due to the number and variety of hosts involved. By variety we mean whether or not they have high-speed Research and Education network links, whether the sites are educational, commercial, etc. We, therefore, decided to study the impact of interventions in more detail by reviewing the data within a country by host. This enabled fewer trajectories to track by eye, and also the separation of R&E and educational links from others. The results below are for all days and by host.

Results by region, country and by host

Asia

E Asia

Middle East (to be analyzed)

South Asia (to be analyzed)

Africa

Central Africa

East Africa

North Africa

Republic of South Africa

Southern Africa excluding the Republic of South Africa

West Africa

America

Caribbean

Central America

North America (to be analyzed)

South America

Europe

Baltics

Balkans

Eastern Europe

Northern Europe

Southern Europe

Western Europe

UK and Ireland

Summary

PingER was originally set up for monitoring connections between High Energy Physics (HEP) sites that were mainly at National Labs and Universities in the US, Canada, Europe, and Japan. As a result, most hosts being monitored were and still are involved in Research and Education and have high-performance network links from organizations such as: Canarie, ESnet, Internet2, GEANT, DFN etc. Such network connections have sufficient capacity to accommodate changes in traffic patterns. These networks are deliberately well provisioned to enable research and explore the future etc.

Thus looking at the overall performance for a country in these regions there is little evidence of the impact of Covid-19 interventions, See, for example, the [analysis for Spain and Italy](#), two countries that were badly impacted by Covid-19. Also, see the [UK](#).

On the other hand, almost all hosts monitored in [Africa](#) (apart from 3 in the Republic of South Africa), and [South America](#) (apart from some in [Brazil](#) and [Argentina](#)) are on public networks that are probably engineered to be good enough for the typical traffic, and thus more likely to be impacted by unexpected changes in traffic patterns caused by Covid-19 interventions.

The types of interventions that we may expect to make noticeable changes in Internet traffic patterns, and possible changes in ping jitter due to changes in path utilization, include:

- Lockdowns;
- Closing/suspending schools, colleges, sporting events, public meetings etc;
 - Thus we might expect to see impacts on hosts at universities. However, as pointed out by Wade Hong of Carleton University "Due to the students in residence being asked to vacate for social distancing. With many university LMSes (Learning Management Systems) now cloud-based, and the use of cloud media services for streaming of recorded content, the potential bottlenecks have moved from the institution to these services."
- Imposing a curfew;
- Calling a state of emergency;
- It is probable that the increases in cases and deaths that are publicized will result in changed behaviors and people voluntarily sheltering in place working from home etc. however we do not regard this as an intervention.

The typical impacts are:

- A sustained (for a week or more) significant increase or decrease in IPDV values for some or all hosts in the country correlated with an intervention that may be expected to change traffic patterns.
- A change in the difference in the weekday vs the weekend performance (bear in mind all times are UTC) correlated with an intervention.
- An increase or decrease in the variability day to day of the IPDV

To get a better idea of the overall impact we looked in more detail at 26 major countries in Africa that included 76 monitored hosts in Southern Africa ([Top Level Domains](#): ZA, LS, BW, NM), East Africa (BI, ER, ET, KY, MG, ML, MZ, RW, SD, TZ, UG, ZM, ZW), and Central Africa (AO, CM, CD, GA)

- There are 10 hosts (1 host in Mozambique, 2 in Botswana, 1 in Cameroon, 1 in Eritrea, 1 in Ethiopia, 1 in Sudan, 1 in Tanzania, 1 in the Republic of South Africa, and 1 in Zambia) that had insufficient PingER data to draw any conclusions.
- We could find no interventions of the above kind for Burundi (2 hosts), and Mozambique (4 hosts).
- Three hosts had high speed (TENET) connections
- Of the remaining 57 hosts:
 - 17 had immediately recognizable IPDV changes within a week of the start of an intervention,
 - 9 had associated changes that required further investigation (typically averaging the IPDV for a period of 14 days before and after the interventions)
 - there were 3 high performance connected (TENET) hosts
 - the remaining 31 either had no significant (i.e. to be significant the IPDV needed to have a difference that lasted for at least a week) changes or the change did not match an intervention.
- We categorized the remaining hosts into the site's purpose, i.e. gov (13 hosts), net (10 hosts), com (15 hosts), edu (17 hosts), org (1 host)
 - The breakdown of host category versus IPDV changes is below.

	Gov	Net	Com	Edu	Org	Total
Strong	3	3	4	7	0	17
Medium	4	0	2	3	0	9
Unclear	7	7	9	7	1	31
Total	14	10	15	17	1	57

Possible Future work

Add data for South Asia.

Analyze more countries in Africa and South America.

References

Spread sheets:

- [PingER daily IPDV from SLAC to over 700 hosts in over 155 countries](#)
- [Correlations between Covid-19 interventions and PingER IPDV changes.](#)
- [PingER daily IPDV versus Covid-19 interventions summary](#)
 - Africa
 - [Central Africa](#)
 - [East Africa](#)
 - [Southern Africa](#)
 - [Angola](#)
 - [Ethiopia](#)
 - Europe
 - [Ireland](#)
 - [Italy](#)
 - [UK](#)
 - South America
 - [Argentina](#)
 - [Brazil](#)
 - [Bolivia](#)
 - [Chile](#)
 - [Colombia](#)
 - [Ecuador](#)
 - [Paraguay](#)
 - [Peru](#)
 - [Uruguay](#)
 - [Venezuela](#)