Internet Interoperability Index: Network Interference

Lily Bhattacharjee, Nick Merrill

Abstract—

Governments, corporations, and the public are embroiled in a debate over to what extent the Internet should be a global network for interpersonal interaction, with some leaning towards a connected space with no borders, while others refer to terms such as the "splinternet" in asserting that the Internet is in fact becoming increasingly nationalized, especially with respect to geographic location. This project seeks to use data science to quantify this fragmentation by analyzing network interference data (e.g. web connectivity, site blocking, etc.) from the Open Observatory of Network Interference (OONI) to assign fractional rankings to each nation based on specific incident occurrence over time. The analysis in developing a formula to specifically address network interference incidents will eventually be displayed on a website along with a real-time globe visualization as well as show reasoning for the country-specific indices. We hope to introduce measurements to drive policy decisions about freedom of information access on the Internet, in America and around the world – especially in the context of FCC regulations on net neutrality – as well to raise public awareness on these issues, and how blocked information can affect the way in which they perceive their governments, companies, and the rest of the world.

1 Introduction

The Internet was created as a medium to facilitate communication between researchers who trusted each other, but since its early development in the 1960s, it has evolved into an interactive network that connects billions around the world. One of its core tenets, as codified in international law according to UNESCO, is the freedom of information, expression, and privacy by users [10]. This project is working on a data-science-driven understanding of how "interoperable"the Internet is across various geographical areas, where **interoperability** is defined as a measure of how uniform the Internet is in the context of different layers of its stack, across countries, specifically available content, protocols, applications, and infrastructure [7].

As some countries, like China, Russia, Romania, and the U.S. (net neutrality) pass regulations that allow the government, or perhaps corporations, more control over information access and the speed at which it is accessed, this project is interested in developing metrics to monitor certain abnormalities in different TCP model layers and correlate these measurements to acts of real-world censorship.

The TCP model of the Internet has four layers – Application (def: how applications create user data and manage resources on the same host device), Transport (def: host-to-host connections), Network (def: packet transport transcending network boundaries), and Link (def: local network communications without routers). This subset of the project focuses on the network layer, which handles packet forwarding, routing, and addressing – how information is sent from a distant server hosting a website to a particular device. Overall, the III project is working on additional metrics to better quantify the impact of censorship on the other layers.

To delve deeper into network interference and its detection, the project analyzed network interference data, provided publicly by the Open Observatory for Network Interference (OONI). OONI has probes stationed in countries across the world, and these collect data when a user runs an OONI-developed Nettest for web connectivity, DNS consistency, etc. as available from direct links from their website. Some examples of active network interference that OONI successfully

detects include HTTP header field manipulation, blocked access to certain apps (e.g. Facebook, Tor, Whatsapp), and shaky connectivity. Note that the last can be the fault of the ISP, but either way, the user is prevented from accessing certain information, so the result is still counted as an act of censorship.

The goal for this segment of the project is to collect, clean, and perform EDA on this data, and calculate a preliminary network interference index for each country based on the rate at which network interference (NI) incidents occur. For this purpose, two possible index formulas were addressed: the strict rate and the loose rate. The strict rate is defined as the following:

$$\frac{\text{number of confirmed NI events}}{\text{total number of tests run}}$$

Similarly, the loose rate is defined as the following:

$$\frac{\text{number of confirmed NI events+number of anomalous events}}{\text{total number of tests run}}$$

It is important to consider anomalous events because some incidents of censorship e.g. IP blocking, DNS poisoning are hard to distinguish from transient network failures. Test results also include a failure designation specifically encompassing failed tests due to probe malfunctions, hardware issues, etc. These rates were designed based on the assumption that a "flat"(optimal) Internet means interference events occur at the same rate regardless of user / location when grouped by test type. Hence, uniform increases worldwide result in a shifting baseline score i.e. if all countries in the world suddenly block Tor, the rates should not change, but should still be calculated relatively and be comparable to one another. Potential designs for future rates may include weighting by number of tests taken by OONI for a particular country to balance out an effect in which countries that have few tests have those tests make larger impact.

2 Data

This research required the prior calculations of strict (SR) and loose (LR) rates by year from the OONI Metadb to ensure that the formulas driving the network interference index were reasonable. OONI has several open-source datasets, and after some experimentation with a smaller rate-limited web API, we used an AWS EC2 instance to set up the complete PostgreSQL database (> 300 million rows), which had more tables of metadata and intermediate calculations available. The tables focused on in this analysis were the following:

- *measurement*: contains information about each test run e.g. start time, type of test, confirmed, anomaly, failure, etc.
- *report*: shares a measurement id with the measurements table, and it is needed because it includes a column called probe_cc, a variable that associates each measurement with a country code
- *input*: domain access attempt by a particular measurement id e.g. website names, IP addresses

Initially, attempting to generate rate calculations with the ooexpl_wc_input_counts table (a smaller, aggregated version of the measurement table) yielded seemingly valid results but looking at the numbers of confirmed / total events revealed that it was not suited for the purpose and appeared to only include tests associated with specific domains (some types of tests do not register domains, and some measurements have NA values).

Using SQL queries to join the relevant measurement, report, and input tables, strict and loose rates were calculated for each country over 2019, and we are currently working backwards from 2018 to 2012, when OONI first started collecting data.

2.1 Exploratory Analysis

Other than the computation aspect of the project, we wanted to analyze the dataset to hopefully shed some more light on a couple of interesting questions that could also be linked to how biased the measurements taken were in favor of certain countries or results:

- How do strict and loose rates change over time? Do they mirror real-world events?
- How many measurements does OONI take per day? Per country?

- Which websites or tests trigger confirmed / anomalous flags the most often? Which countries are these requests coming from?

Regarding the second question, during spikes, certain countries tend to contribute significantly, signaling that it is not a worldwide spike in the number of people finding out about OONI or running their tests. These increases in testing also correlate to an increase in the number of confirmed or anomalous tests.

Specifically, increases in measurements over time are highly correlated with peaks in the number of measurements in Russia. This is likely because Russia is the country that runs the largest number of OONI tests over the course of 2019 specifically. China does not appear to run many, if any, tests, despite the existence of the Great Firewall which should yield many confirmed / anomalous events. This may be an indicator that either the probe is blocked, failed, unable to be placed there for collecting data (China is notoriously careful and even keeps many Google products out), or that China may not even have access to the OONI website under the conditions of the firewall.

In terms of question 3, we notice for the particular countries included in this report (China, Russia), China tends to block websites that run counter to the current Communist government's agenda e.g. the democracy activist group Tiananmen Mothers. More recently in April 2019, China also blocked all language editions of Wikipedia, which explains why we have some entries from the Dutch and French versions in the word cloud. The graphic itself is designed so that the size is correlated to the number of confirmed / anomalous events.

Russia and Romania, in addition to blocking politically controversial sites, also block casinos. Romania specifically does not allow online gambling operations without a license, so only 18 organizations have licenses right now. Russia and Romania actually block some of the same websites e.g. the domain europacasino.

Figure 1: Number of OONI Tests (2019)

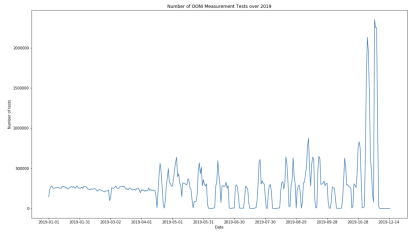

Figure 2: Blocked Websites (CN)

Figure 3: Blocked Websites (RU)

Looking at the figure below, most data points until late August (8/28) are real and calculated from the SQL tables but the ones lying along the dotted line segments are projections based on an autoregressive model. It was automatically fitted to the data using a Python library called statsmodels to check out where indices are likely to trend in the future using past data. AR models work similarly to linear regression, except instead of being dependent on x-value e.g. y as a function of x, AR models are dependent upon a set of past values e.g. $t-1$, $t-2$, etc. Simple AR models tend to have the same spacing between t-values so the differences will be symmetric that way, and this model is based on up to two timesteps in the past. It should not be taken as an accurate representation of a forecast, per se, but rather where indices are likely to go absent sudden shocks (which are almost guaranteed). It does help us notice long-term propensities like Iran's increasing loose rate (and thereby censorship) over time.

Figure 4: Loose Rate Predictions (2019)

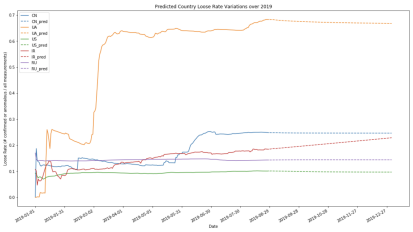

3 Possible Models

These figures show the variations of country-specific loose and strict rates for Russia, China, U.S., Iran, and Ukraine. U.S. was picked as a baseline – as can be seen – while there are no confirmed events throughout the year of 2019, there are some anomalous events mostly associated with web connectivity / speed. This is most likely due to ISPs, and it would be interesting to check out the calculated rates for 2017, when Ajit Pai announced a repeal of net neutrality rules. While some states have passed regulations protecting net neutrality and some ISPs e.g. Verizon have promised to honor it across the nation, it would be curious if there was a spike. Examining the trends of LRs and SRs per country more closely, we notice the following for Iran specifically:

- spikes at 1/15 (Broad DNS hijacking campaign could originate in the Middle East)
- spike beginning of Feb (Iran begins marking 40th anniversary of Islamic Revolution)
- peak mid-June (Twitter removes nearly 4,800 accounts linked to Iranian government, Russia and Iran Plan to Fundamentally Isolate the Internet)
- spike 8/15 (US issues warrant to seize Iranian tanker off Gibraltar)

Another interesting point is that the sudden rate drop-offs in Figures 6 and 8 near the end of the year when the data was calculated likely correspond to dates when data stopped being collected for each country's probes, although these do appear to be staggered. This implies that a completely "live"display of these rates may not be possible due to delay between data collection and logging

3.1 Strict Rates

Figure 5: Strict Rates (2019)

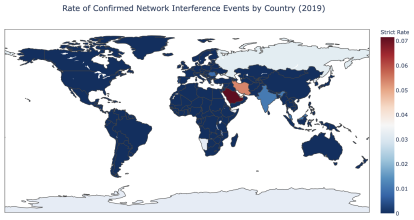

Figure 6: SR Fluctuations by Country (2019)

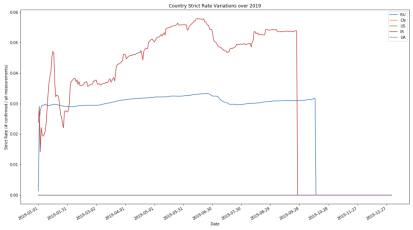

3.2 Loose Rates

Figure 7: Loose Rates (2019)

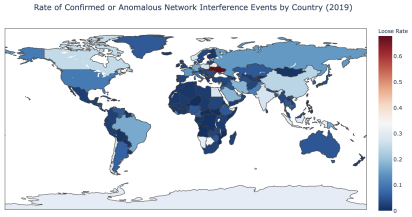

Figure 8: LR Fluctuations by Country (2019)

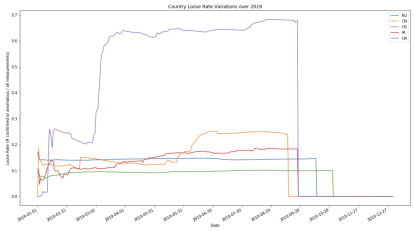

4 Impact

The plan for this project is to work on hosting a website that updates the overall Internet Interoperability Index (of which Network Interference is a part of) corresponding to each country in real-time. The colored maps of the world in Figures 5 and 7 are preview to what that might look like. We hope, by making the

data and our reasoning publicly accessible, to introduce data to the public debate about Internet censorship and its limits, to inform civilians about how their country is doing compare to the rest of world in matters that concern the freedom of information, and to bolster the legitimacy of policy initiatives to change laws in highly-censored nations.

However, before this goal for the future is achieved, a few steps need to be taken to make the analysis in this document more airtight:

- OONI probes only collect data from users who know about the Nettest and manually run their tests, which can result in a biased / incomplete dataset. We are still looking into how to mitigate this bias, potentially by integrating another variable (in the past, we were considering GDP, average income, or other demographics that can be associated with better knowledge of technology and therefore web testing) but there is a possibility that more variables from these datasets can introduce more noise into our index calculations?
- Can we use OONI metadata from other tables in the database to find a culprit for confirmed or anomalous events? Can we differentiate between transient failures and malicious blocking activity?
- How do we handle days when probes fail or do not collect data? My graphs and current calculations keep the rates constant over those days?
- How accurate are the strict / loose rates in telling us about the state of censorship on the Internet across the world? Qualitatively, we can associate current events with spikes / gradual increases in certain rates, and make educated guesses but we're still working on quantifying this in order to find out how to improve our indices and measure that improvement.

References

[1] Lopez, P. (2019, May 14). Wikipedia blocked in China in all languages. Retrieved December 1, 2019, from https://www.bbc.com/news/technology-48269608.

[2] Xynou, M. (n.d.). Open Observatory of Network Interference. Retrieved December 1, 2019, from https://ooni.org/post/next-generation-ooni-explorer/.

[3] Filastò, A. (2019, May 21). OONI MetaDB Sharing. Retrieved December 1, 2019, from https://github.com/ooni/sysadmin/blob/master/docs/metadb-sharing.md.

[4] Merrill, N. (2019, November 29). Internet Fragmentation: Beyond "free"and "closed". Retrieved December 1, 2019, from https://medium.com/cltc-bulletin/internet-fragmentation-beyond-free-and-closed-cb8b1dfcd16a?sk=58aa0e0436f0ccfa8464f972dda9ec12.

[5] Economy, E. C. (2018, June 29). The great firewall of China: Xi Jinping's internet shutdown. Retrieved December 1, 2019, from https://www.theguardian.com/news/2018/jun/29/the-great-firewall-of-china-xi-jinpings-internet-shutdown.

[6] Doffman, Z. (2019, May 1). Putin Signs 'Russian Internet Law' To Disconnect Russia From The World Wide Web. Retrieved December 1, 2019, from https://www.forbes.com/sites/zakdoffman/2019/05/01/putin-signs-russian-internet-law-to-disconnect-the-country-from-the-world-wide-web/2e25fea81bf1.

[7] Frystyk, H. (1994). The Internet Protocol Stack. Retrieved December 1, 2019, from https://www.w3.org/People/Frystyk/thesis/TcpIp.html.

[8] Judicial Official Announces Order to Block Instagram Less Than a Year After Banning Telegram Messaging App. (2019, January 4). Retrieved December 1, 2019, from https://iranhumanrights.org/2019/01/judicial-official-announces-order-to-block-instagram-less-than-a-year-after-banning-telegram-messaging-app/.

[9] Onjn.gov.ro. (2019). Approved – Oficiul National pentru Jocuri de Noroc. Retrieved December 1, 2019, from http://onjn.gov.ro/approved/

[10] Freedom of Expression on the Internet. (2019, March 4). Retrieved December 9, 2019, from https://en.unesco.org/themes/freedom-expression-internet.