

An Observatory for Realtime Monitoring and Analysis of Internet Blackouts Caused by Censorship

1 PROPOSAL

1.1 PROPOSAL NARRATIVE

1.1.1 SUMMARY

Episodes of politically-motivated interference with Internet access are widespread and frequent. In particular, large-scale connectivity disruptions are often used by repressive regimes to restrict the free flow of information. These events are often undocumented, unverified, surrounded by uncertainty about their timing, extent, cause, and specific mechanisms of execution. Objective measurement data are indispensable for establishing facts and responsibilities in order to inform citizens, policymakers, the media, and activists.

Based on diverse types of Internet measurements and peer-reviewed scientific methods, at the Center for Applied Internet Data Analysis (CAIDA), University of California San Diego, we have developed and deployed a platform called IODA (Internet Outage Detection and Analysis) that monitors the Internet 24/7 in order to detect and visualize information about Internet connectivity disruptions. IODA provides a public dashboard, providing timely information about network outages in near-realtime, while it enables registered users to further inspect and explore our data as an aid to investigate disruptive events. The overarching goal of this project is to further enhance the role and impact of IODA as an observatory for Internet shutdowns caused by censorship.

Over the last year, we have dramatically increased and diversified IODA's user base through active engagement with groups interested in and affected by Internet shutdowns (including technical collaborators and regional activists). We advertised and explained how IODA can be used to analyze shutdowns at community events such as the Citizen Lab Summer Institute and RightsCon. We provided active support to users through video calls, and produced screencast tutorials. Additionally, we often reported and shed light on detected events through Twitter and blog posts, sometimes in collaboration with other groups (e.g., OONI).

Our experience with IODA and our interaction with the community has helped us identify and prioritize new key methodological and usability improvements in the detection of Internet shutdowns and in IODA specifically. We request funding to address these gaps and further improve IODA as an observatory for Internet shutdowns caused by censorship. Specifically,

we propose to (i) address open challenges in detecting and analyzing Internet connectivity shutdowns, (ii) improve usability of our outage detection platform (IODA), (iii) further increase its utilization by civil society and other non-research organizations, and (iv) explore the use of IODA to identify in a timely manner which means of connectivity (e.g., telecom operators, IPv6) continue to be available during a shutdown.

Project site: <https://www.caida.org/projects/ioda/>

Demo dashboard: <https://ioda.caida.org>

1.1.2 PROJECT DESCRIPTION

Problem overview

Internet shutdowns reached mainstream attention during the Arab Spring in 2011. Such shutdown events appear to be growing more numerous as hundreds of similar events, involving dozens of different countries worldwide (Syria, Iraq, Libya, Egypt, China, Bahrain, India, Gabon, ...), have been registered in the following years. In July 2016, the United Nations Human Rights Council unanimously passed a resolution condemning Internet shutdowns. Despite this and many other authoritative efforts, the number of government orders to temporarily shut down or restrict access to Internet services seem to be on the rise.

Mounting effective responses to shutdowns is significantly hampered by a fundamental problem: Internet shutdown events are very challenging to measure. There is a scarcity of objective independent sources providing evidence and contextual information about these episodes. While projects such as OONI, Oracle's Internet Intelligence Map, NetBlocks, and IODA provide much-needed visibility into such episodes, these projects' views into Internet shutdown events is by no means complete. For example, state-of-the-art systems still lack visibility into shutdown events in IPv6, events affecting some ASes (but not others) in a country, events affecting primarily cellular users etc. Better insight into these phenomena will help understand the mechanisms through which shutdown events are implemented, and importantly, have the potential to reveal censorship evasion techniques.

Another problem faced by the user-community interested in Internet shutdowns is the lack of timely and publicly accessible measurements of Internet shutdown episodes. The IODA system has played a crucial part in filling this gap by providing publicly accessible interfaces that users can leverage to study the current (and historical) state of Internet connectivity in various areas and networks around the world. However, these interfaces are still in the prototype stage, and are thus complex and can present a barrier to entry to new users. Further, in spite of IODA's rapid uptake over the last year by a diverse set of actors, there remain several members in the community who are currently unaware of IODA's capabilities and how IODA can help them. Additionally, the availability of different projects that measure Internet shutdowns may confuse users thus efforts to fuse the insights of these complementary projects can result in significant benefit to users.

Summary of proposed contributions

The proposed project will make the following contributions to address the above challenges:

- We will improve IODA's technical capabilities to obtain measurements and insights into previously unexplored aspects of Internet shutdowns (such as IPv6 connectivity during

shutdowns). We will also perform usability enhancements to IODA which will enable more advanced analyses, such as the comparison of connectivity of Autonomous Systems in a country during a shutdown.

- We will continue our outreach efforts and provide guidance to enable IODA's usage by diverse actors. Further, we will continue to complement existing systems and data sources, contributing to obtain a more comprehensive picture of Internet shutdowns.

Technical background:

Over the last six years, CAIDA has developed an Internet measurement, data analysis, and automated detection methodology that combines diverse types of data to identify macroscopic Internet outages. The IODA system infers and publicly reports outages significantly affecting a country, a region within a country (e.g., province) or an Internet operator (i.e., Autonomous System) in near-realtime. The IODA system presents several elements of novelty that distinguishes it from other efforts:

1. *Internet Background Radiation (IBR)* is background noise traffic generated by hundreds of thousands of Internet hosts worldwide. IODA is **the only outage detection system using IBR**, thanks to a methodology published by CAIDA in collaboration with RIPE NCC, which we demonstrated in the context of the outages happened in Egypt and Libya during the Arab Spring. We collect IBR traffic through the UCSD Network Telescope, a large block of IPv4 address space estimated to observe 1/256th of all the IBR generated in the Internet. Our inferences have a granularity of 1 minute.
2. IODA separately analyzes and then **combines inferences from three different types of measurement data**. This is highly relevant because each data sources can "see" outages that the others do not observe. In addition to *IBR*, IODA makes inferences using:
 - *Active probing*. We run our measurements from twenty nodes at UC San Diego. We periodically (every 10 minutes) probe approximately 3.5 million networks worldwide and adaptively send more probes to eliminate uncertainty, using a methodology developed by USC.
 - *Global routing information (BGP data)*. For this data source, we leverage the collection infrastructure operated by the RouteViews and RIPE RIS projects. Every 5 minutes, we process BGP data from hundreds of operational routers in different locations of the Internet topology, each providing their view of the current status of global routing on the Internet (i.e., which networks appear reachable on the Internet control plane).
3. IODA is one of a few systems operating continuously and detecting events in **near-realtime**, with a delay in the order of minutes.
4. IODA emphasizes the use of **interactive visual interfaces** both for intuitively communicating the results of complex inferences to non-technical users and to enable effective data exploration. The public dashboards enable a user to "see" the alerts and the corresponding outages. Registered users have access to historical data and additional datasets (e.g., intermediate data from the processing pipeline providing richer information, or data from sources not yet integrated in the automated detection system) as well as to a data exploration interface ("Explorer") that allows to jointly **inspect** multiple time series and apply mathematical transformations.

Description of proposed contributions

Here we provide an overview of the proposed work. These sections provide context for our Project Objectives and Project Activities/Deliverables.

a) Proposed technical improvements: Objective 1 and 2

IODA's current methodology for outage detection is state-of-the-art but it lacks visibility into certain types of censorship events (e.g.: events that affect some cellular ASes, events that affect IPv6 etc.). Gaining the ability to measure and shed light upon these events will significantly improve the community's understanding of *how* shutdowns are implemented and has the potential to unlock new censorship evasion measures.

An orthogonal direction in which we intend to make technical progress is in the enhancement of IODA's usability through its visual interfaces. IODA is in the prototype stage; as its uptake among users increases, we learn more about usability features that they are interested in. We also come to know about bugs and other issues with the visual interfaces that users are facing.

In this proposed work, we will achieve the following technical improvements:

1. We will investigate IPv6 connectivity during Internet shutdown events. During some prior events, we observed that IPv6 connectivity persisted even when IPv4 connectivity was blocked. A potential explanation for these events is that the censoring entity may have neglected IPv6 when implementing the shutdown. Analyzing IPv6 connectivity in addition to IPv4 can therefore help shed light on how shutdowns affect the two protocols. Further, such analysis can determine if censored users can continue to access the Internet over IPv6 and can perhaps help them circumvent censorship implemented solely on IPv4.
2. IODA is one of the few systems that can detect outages at the AS-level. This enables novel analyses, such as identifying which ASes in a country are (not) experiencing a shutdown. To further leverage this capability, we will enhance IODA's user interface to enable deeper and more detailed AS-level outage analyses. We envisage that such analyses can be particularly useful, since knowledge of which ASes have not been affected by a shutdown could be used to circumvent censorship.
3. We will investigate adding data sources to IODA that can help provide visibility into outages affecting cellular networks. Studies of Internet shutdowns suggest that governments in many countries are especially likely to target cellular Internet Service Providers (ISPs). However, IODA currently has limited visibility into the connectivity of cellular networks. This is because cellular ISPs often have distinct properties (like the use of Carrier Grade NAT and the blocking of responses to active probes) compared to other networks, necessitating new sources of data to analyze shutdowns in cellular networks.
4. We will perform several enhancements to IODA's visual interfaces, both to fix known bugs, and also to enable new features. A key task for this purpose involves upgrading IODA's interface to use a new version of the user interface framework that we developed as part of another CAIDA project. As part of this task, we will also take feedback from OTF's Usability Lab.

5. We will continue to produce screencasts and tutorials to help users analyze outages. Screencasts will focus upon newer features that we will add (such as more detailed AS-level outage analyses) or help users gather a more comprehensive (e.g., over a longer period of time) view of a country's connectivity.

b) Outreach and Collaboration Efforts: Objective 3 and 4

We envision IODA to be used by a diverse audience that includes civil society actors, policy makers, and activists. Further, IODA provides a platform that will gradually bridge the gap between technologists and civil society actors, encouraging collaboration between these two spheres. During the previous year, we have observed progress towards this vision, with adoption and use of IODA by diverse users and the creation of new collaborations. We have created more than 30 IODA accounts for users with varying backgrounds, including users from other groups performing censorship-related measurements such as OONI, users from activist organizations such as Tibet Action Institute, users from advocacy groups such as the Software Freedom Law Centre, India (internetshutdowns.in), and even the Office of the United Nations High Commissioner for Human Rights (OHCHR). The number and diversity of actors who can benefit from using IODA is large enough that we envision the need for significant efforts for continued outreach over the coming year as well, especially focusing on making users autonomous in their use of our platform. Further, we have gained considerable experience through our interaction with the community and this experience has taught us how to better direct our efforts.

Our efforts over the last year have led to several fruitful collaborations with related projects. Our goal with these collaborations is to complement existing systems and data sources, contributing to obtain a more comprehensive picture of Internet shutdowns, and increasing the amount, granularity, accuracy, and timeliness of available data. Towards this goal, we have been actively collaborating with OONI on performing more detailed analyses of Internet shutdown events, with each project contributing complementary insights. We have begun collaborations with Psiphon to explore flows and anomalies in Psiphon traffic that often correspond directly to shutdown events identified in IODA. We have had several interactions with groups interested in exploring sub-national Internet outages, such as the Tibet Action Institute and internetshutdowns.in.

This year, we plan to continue to engage in outreach efforts to advertise IODA, teach people how to use it, and also engage in collaborative outreach efforts involving the use of IODA and other related projects such as OONI and Psiphon. Our team will engage in outreach efforts to maximize IODA's use to the general public. We will continue to release "public information announcements" to communicate events in near real time through Twitter and will also continue our efforts to inform the KeepItOn mailing list. We will occasionally publish reports targeting both the general public and technical audience, sometimes in collaboration with related projects such as OONI. Finally, we will participate in events where we can share our work and exchange ideas with other technical experts and civil society actors, most importantly the Citizen Lab Summer Institute and RightsCon. During these events, we will investigate opportunities to form collaborations with other interested groups.

1.1.3 PROJECT OBJECTIVES

We spent the last five years of IODA's project history on research and development. We are now approaching a point where we have enough faith in IODA's design and methodologies for it to be introduced to the public, and improve it iteratively, with the help of our users. We believe that the following objectives are the most important milestones towards making IODA useful to the broad public.

1. Objective 1 Improve IODA's technical capabilities
2. Objective 2 Make IODA more useful by improving its user interface
3. Objective 3 Release timely reports about shutdown events
4. Objective 4 Engage in outreach and collaboration

1.1.4 PROJECT DELIVERABLES/ACTIVITIES

Below is a detailed list of our objectives and the corresponding tasks.

Objective 1 (Improve IODA's technical capabilities)

- 1.1 Monitoring IPv6 connectivity:
 - 1.1.1 Extend our current BGP-based outage detection approach to work with IPv6. This task requires adding IPv6 support to our libIPMeta framework, the work of our research developer will be overseen by Alistair who is the original developer of libIPMeta. We believe this task is well understood.
 - 1.1.2 Investigate theoretical and practical challenges in developing active probing techniques for the purpose of detecting outages in the IPv6 Internet. This task represents a preliminary investigation in order to focus future development efforts.
- 1.2 Investigate adding data sources to IODA that can help provide visibility into outages in cellular ASes. This task is exploratory and involves collaboration with external parties (e.g., researchers at Internet Initiative Japan). While this type of engagement carries risks, we believe the potential benefit of illuminating shutdowns to cellular connectivity outweighs such risks.
- 1.3 Add a new dashboard to enable users to investigate a country of interest. This dashboard will list time series for all ASes/regions in a country specified by the user (including ASes/regions which have no ongoing outages). This is mainly a development task based on data sources already present in IODA, we therefore do not expect particular issues in carrying out this task.
- 1.4 Improve the stability of the liveness signal extracted from the IBR data source. The network telescope signal is polluted by both erratic and spoofed traffic. We plan to refine and update our methods to filter out this traffic, which greatly reduce the number of false positives. We understand this task well, since we periodically need to perform this task to keep the IBR signal usable.

Objective 2 (Make IODA more useful by improving its user interface)

- 2.1 Produce screencasts and tutorials to help users analyze outages with IODA. Screencasts will focus upon newer features that we will add (such as more detailed AS-level outage analyses) or will help users identify anomalies and artifacts in our datasets, as well as gathering a more comprehensive (e.g., over a longer period of time) view of a country's connectivity, helping them to use IODA more effectively.

In the past year, we have produced screencasts describing the IODA Dashboard and Explorer interfaces and we have received very positive feedback from users. Users with technical backgrounds, such as our collaborators from OONI and Psiphon, have informed us that the screencasts helped them perform more advanced analyses. New IODA users with non-technical background, such as members from the Tibet Action Institute, have provided feedback that the screencasts helped them consolidate details about how to use IODA after our initial demo of IODA to them. We have also received positive feedback from IODA users at events like RightsCon (some of them learned to use IODA on their own with the help of the screencasts). Thus, we see the screencasts as an important component towards our goal of empowering users to adopt IODA for their analyses.

We understand this task well as we performed this type of task in our previous OTF funded project receiving extremely positive feedback.

- 2.2 Upgrade the visual interface codebase and introduce usability improvements. We will work with the OTF's Usability Lab to identify the most pressing improvements to IODA's user interface in order to lower the barrier for non-expert users. Examples of the improvements we will carry out are:
 - IODA currently has no notion of time zones – everything is based on UTC. We want to provide a menu that lets users configure what time zone they are interested in.
 - IODA's time series tree is complicated and confusing to navigate. We want to replace it with a more intuitive way for users to select the time series that they are interested in.

However, the codebase of IODA's visual interface needs a major upgrade in order to simplify the job of performing modifications and introducing new functionalities. At CAIDA we have developed a new version of the Charthouse Javascript/PHP framework that IODA was originally based on and used it to support other projects. The upgrade of IODA to the new framework will require substantial work by a UI/UX expert with support of a software developer. Such an upgrade is also required to support an efficient (and less time consuming) implementation of the visual interfaces needed for IPv6 and AS-level analyses (Tasks 1.1.1, 2.3).

- 2.3 Extend user selection of aggregations to support regions of a country and autonomous systems. This capability supports users looking for signs of a suspected disruption, for example in the context of protests limited to a specific region, or rumors about a certain ISP being shutdown. We understand this task well and it requires only implementation by a UI developer.
- 2.4 IODA's user interface shows "alert bands" to draw attention to an outage. Sometimes these alert bands are not displayed correctly; presumably due to a change in how

browsers render SVG images. We will look into this problem and make alert bands rendering consistent. Furthermore, alert bands that cover a small time span may not be displayed in IODA's interface if the user zooms out, causing the data points to be aggregated. We plan to "bubble up" the alerts to the current zoom level. We had a preliminary look at this issue and concluded that it is not trivial and may require a significant rework of our visualization framework.

- 2.5 When clicking on an IODA link, the redirection occasionally breaks if the user is not logged into IODA. While this is not mission-critical, it is a technical issue that negatively affects usability by often confusing users and slows them down unnecessarily. We understand this task well and it only requires implementation.

Objective 3 (Release timely reports about shutdown events)

- 3.1 We will continue to use the IODA Twitter account for timely dissemination of detected outages. We will use this account to tweet outage visualizations and corresponding links. Since December 2018, we have been (re)tweeting regularly about Internet outages detected by IODA, averaging 1-2 tweets per week. During this time, the number of followers of our Twitter account has increased to 370 from less than 10 initial followers. Some of these tweets are from IODA users who have been using IODA independently to detect and report upon Internet shutdowns. Consistently with our experience we expect to (re)tweet once a week on average.

We will also publish detailed analysis providing insight into interesting events either using the CAIDA blog or in collaboration with other groups (e.g., in the past we published such analyses in collaboration with OONI on their blog). Since December 2018 we have (co)authored three blogposts offering in-depth analyses of shutdown events and are actively working on two more. Consistently with our experience we expect to (co)author a blogpost once every 2 to 3 months (depending upon the occurrence of shutdown events).

We understand this task well and it only requires execution.

- 3.2 IODA is one of the few systems that can detect outages at the AS-level. We will use the enhancements to AS-level user interfaces to perform in-depth analyses of AS-level outages during major shutdowns and report upon our findings. We understand this task well and it only requires execution.

Objective 4 (Engage in outreach and collaboration)

First, we would like to make a key clarification: *users do not need accounts to use the main IODA interface and functionalities (the public "IODA Dashboard")*. We make all the essential components of IODA interfaces public in order to maximize outreach and transparency. We require users to create an account only when they need to perform advanced analyses using IODA's Explorer interface. Users with these needs and expertise are rare. We apologize for not having made this clear in our earlier submission and in the next paragraph provide details about user uptake mostly based on Google Analytics data.

Over the past year, we have seen a very large increase in user base for both the interfaces. The Dashboard interface has tremendously grown in popularity during this time: from 2 new users in the month of Jan 2018, we now see over 200 new users of the dashboard in each month

of 2019. June 2019—when we participated in RightsCon—saw 875 new users! Moreover, the new users are from diverse geographical locations, across all continents: of 875 new users in June 2019, 392 are from Europe, 282 from the Americas, 166 from Asia, 23 from Africa, and 11 from Oceania. IODA “power-users”, i.e., the users of the Explorer interface (which do require accounts), have grown in number to 30 from 5.

These numbers are in keeping with our overarching goal to enable IODA to be used more widely. With our various outreach efforts, we have gained a significant user-base who are familiar with IODA and use the dashboard. A small fraction of these users are interested in performing additional advanced analyses: these are the users for whom we create IODA accounts (which lets them use the Explorer interface), and typically, also provide additional support through video calls and email. We know that some of these users actively use IODA since we have observed that they tweet and/or release blogposts about events using IODA.

- 4.1 Alberto Dainotti, Ramakrishna Padmanabhan, and Alistair King have been on the KeepItOn mailing list. We will continue to participate in the mailing list discussion and provide insights, references, and tutorials. We understand this task well and it only requires execution. We are also working with third parties that use IODA visualizations as supporting evidence for their work. These analyses may result in reports, albeit not co-authored by our team.
- 4.2 We already attended the Citizen Lab’s Summer Institute 2018 and RightsCon 2019 and plan to attend additional Internet freedom forums, including the next editions of CLSI, RightsCon and the Internet Freedom Festival.
- 4.3 We will continue to collaborate with other groups operating in this space and foster new collaborations. We have already collaborated with OONI on the analysis of shutdown events in several countries and we have discussed plans for further joint work. We discussed avenues for collaboration also with Keith McManamen from Psiphon and the Software Freedom Law Centre, India (internetshutdowns.in). Recently we have provided support to the United Nations (www.ohchr.org) to use IODA interfaces and data.
- 4.4 Over the last year, we have created more than 30 IODA accounts for users with diverse backgrounds. We occasionally assist these users in the analyses they are interested in. Depending on the complexity, it can take a couple minutes or up to several hours to help them solve their problem. We plan to keep assisting our users to make the best use of IODA. We understand this task well, and it only requires execution.