





























## Summary Review Documentation for

# “Analysis of Country-wide Internet Outages Caused by Censorship”

Authors: A. Dainotti, C. Squarcella, E. Aben, K. Claffy, M. Chiesa, M. Russo, A. Pescapè

### Reviewer #1

**Strengths:** The paper is timely, these events just having happened this year, and with continuing instability in the region, may happen again.

The study uses multiple sets of data, and the authors show a clear understanding of the strengths and weaknesses of each dataset.

**Weaknesses:** It may be tempting to write this paper off as “yet another” because this topic has received plenty of attention, even at the technical level. However, that would miss the point that it is important to have such studies reported in peer-reviewed literature to serve as an archive for future investigators. Only by such efforts will be improve the Internet, and for that matter, human rights.

**Comments to Authors:** In Table 1, it would be very useful to see the intersections between NIC and Geomapped IP addresses located to the country. At a qualitative level it would help inform about how consistent they are. At a quantitative level it might even allow a Petersen-like estimate of how many IPs are missing, e.g. see Matthew Roughan, Jonathan Tuke, Olaf Maennel, “Bigfoot, Sasquatch, the Yeti and other missing links: what we don’t know about the AS graph”, Proc. IMC’08, pp. 325-330, Vouliagmeni, Greece, October 20 - 22, 2008.

Please be more careful about the term “permanently withdrawn”. Presumably you mean withdrawn for the life of the event, but for instance, the before and after states from Figures 3 and 4 don’t seem to quite match, so maybe some prefixes did go away forever?

### Reviewer #2

**Strengths:** A methodology to combine multiple sources of data to detect Internet outages. The paper provides a very good application of large scale monitoring of the Internet.

**Weaknesses:** There is not much technical innovation in the paper. Most of the analyses in this paper use standard techniques for analyzing such data.

**Comments to Authors:** I find the topics of discussion in the paper interesting from historical and political point of view. I also like your careful analyses of the data. However, I find that the technical innovation in the paper is a bit weak for a full IMC paper. Most of the techniques are standard analyses of the data.

The observations on the reachability of IPv6 addresses are interesting as it provides an opportunity for people to by-pass

censorship. It would be great if the authors can elaborate on this point and its potential application to censorship elsewhere.

A more effective detection system should use more targeted active probes. Perhaps, these measurements could be triggered by alarms from BGP and darknet data.

The work can have even bigger impact if the authors use the same analysis for other types of outages such as network failures, natural disasters, not just political censorships.

### Reviewer #3

**Strengths:** This very well done paper could be of interest to more than just the measurement community. Also, its completeness has a chance to make it the definitive work on the subject.

**Weaknesses:** For completeness, I would have liked to better understand the physical topology and the resulting limitations.

**Comments to Authors:** Great paper - timely, well written, and complete. Not sure what more I can add.

A small story that I feel like should be better emphasized: that the backscatter detectors were actually sensitive enough to track both the censorship of these prefixes, but also the DDoS random src-address spoofing.

The fact that Egypt did not withdraw its Ipv6 prefixes is interesting: do the authors believe this was just an oversight?

I realize that there is a lot of information in this paper, but I feel like its still not the “complete picture” because of the lack of physical topology information: how many external links do these countries have and how many operators run them?

I very much appreciated the note about ethics in Appendix A and believe that more researchers should stop and consider the implications of their work, even if in the end they decide not to anonymize their data.

### Reviewer #4

**Strengths:** Timely analysis. Interesting to see that as a community, we have developed or otherwise have publicly available access to tools with adequate visibility into such events.

**Weaknesses:** Not a breakthrough on the technical front. Also, finding what you are looking for is much easier than if you did not know what was happening. In most cases, there are widely published timelines of disruption and reconnection events

and this paper shows those times to be correlated with the feeds they analyzed. However, it is unclear if the same techniques can be used when such public timelines are unavailable. For example, there could be many events causing similar smaller disruptions (misconfigurations, power outages, religious holidays) and seems that the data has little predictive value besides indicating abnormality.

**Comments to Authors:** I enjoyed reading the analysis!

Regarding traffic received by the darknet, there are multiple places where the paper seems to worry too much. If you can distinguish which traffic originated from the observed IP ranges and which was due to backscatter from spoofed probes as you seem to be able to, what is the concern?

I am surprised at the synchronization disconnection across multiple ISPs... I was hoping to see some distance between the complying ASes and those trying to circumvent or resisting the censorship.

Could you tell what fraction of the flat periods in Figure 3, when some prefixes have started disappearing but not all, are due to BGP reconvergence?

For both darknet traffic (Figure 2) and BGP visibility (Figure 3), it would be good to show a longer time-frame control graph. That way, we could figure out the 'natural' variation in these feeds.

The matching from address space to prefixes in 4.1, prefers to find prefixes that are more likely to contain the required set of IPs at the cost of completeness in coverage. It would help to explicitly call that out and use both sets of prefixes: 'accurate but potentially incomplete', and 'complete but potentially inaccurate'.

You seem to have compensated for temporary disappearances of prefixes, but how?

Results are a slight let-down: mostly BGP withdrawals, except for Libya, which perhaps used data-plane filters. Near universal compliance from regional ISPs, some resistance from International providers...

Interesting side-effect to see the DoS attacks.

## Reviewer #5

**Strengths:** The paper is useful in that it sheds light on a relatively new type of event (country-wide, censorship-induced outage). It also presents an interesting and practical idea (to use network telescopes to early-detect such outages).

**Weaknesses:** Ideally, this should have been a short paper (there is not enough technical content here to warrant 14 pages).

**Comments to Authors:** My only complaint is that this should have been a short rather than a long paper. The main contribution of the paper lies in the BGP and telescope measurements, and these can certainly fit in 6 pages. Section 4 can be significantly condensed: I believe that most people who would be interested in reading this paper already know about RouteViews and the UCSD telescope, and the way the authors used these sources is straightforward and similar to the way others have used them

before (in my opinion, it does not warrant two pages of explanation). Section 6 repeats the conclusions drawn in Section 5. Section 5 itself has a fair amount of redundancy and graphs that are larger than necessary (and Fig. 8, which can be completely omitted). In general, the paper gives the impression that it was written and formatted to use more space than necessary.

A less serious complaint is that, in my opinion, the paper overclaims. I think that the first and third contributions stated in the introduction are the same ("We document a rich view of the disruptions..." "We report previously unknown details of each event..."). Also, I think that the fourth contribution ("We sketch a general methodology...") is not supported by the paper: first, I don't think that the paper presents any novel methodology for collecting or analyzing measurements; second, even if it did, I do not think that \*sketching\* (as opposed to describing and evaluating) a methodology could count as an IMC-level contribution. However, this is a matter of opinion and presentation, so I did not hold it against the paper.

Minor questions:

- Section 5.1.2 says that "only 176 prefixes remain[ed] visible" after the start of the outage. Is there anything special about these prefixes? Is there any indication as to why they were spared (do they host any particularly important infrastructure)?

- The same section says that "several of the visible ... prefixes were reachable through ... (IntAS2) ... or ... (IntAS3) ...". Does this observation say anything about the outage? Is there any indication that the spared prefixes were spared \*because\* they were reachable through the particular ASes?

- Later, one reads that "at least three ... prefixes, amongst those ... not withdrawn ... were actually reachable." Is there any indication that the remaining non-withdrawn prefixes were \*not\* reachable? The text implies that the ad-hoc active measurements targeted only very few prefixes. Was there a particular reason for that? In general, a better description of this ad-hoc measurement process would help.

## Response from the Authors

We thank the anonymous reviewers for their insightful comments, which inspired improvements to the camera-ready version of the paper, both in terms of better clarity and additional content.

One of the reviewers observes "The work can have even bigger impact if the authors use the same analysis for other types of outages such as network failures, natural disasters, not just political censorships". This paper represents a first step in this direction, however, we focused on these specific events for two reasons: (i) we wanted to provide sufficient detail to demonstrate the power of our analysis methodology, including to reveal aspects of the observed phenomena not previously discussed; (ii) the prevalence of censorship-related Internet blackouts is growing, and there has been little scientific study of the different approaches and their effects. We recognize there are ethical issues, including revealing holes in censorship that could endanger people in the censored country, and we appreciate the reviewer support for our anonymization of operator names.

With regard to the paper being a better short than a long paper, we saw sufficient novel technical content to justify a full-length paper. Even if the tools and methods we used are not new, some

of them, e.g., the network telescope, have not been used to study macroscopic Internet outages or censorship. We combined several different methods and granularities to reveal new insight into censorship-induced network behavior. We highlighted their complementarities, and suggested a path toward automation of our methods. Background and analysis of the events, description of the data sources, and the findings we illustrated with the figures, just would not have fit into a small paper. Two reviewers

asked us to further elaborate on the reachability of IPv6 addresses. We changed the IPv6 paragraph in the discussion based on these comments, adding two possible causes. Our thoughts on the limitations of current geolocation databases, requested by one of the reviewers, were limited to the focus of the paper; we cite references that provide more specific studies of the topic. Finally, to better contextualize our study, we added information on the physical connections of these countries to the rest of the world.