Chris Ritzo and Jordan McCarthy

OTI Explores Frameworks for Ethically Responsible and Scalable Network Interference Measurement

OTI fundamentally believes that everyone has the right to access an Internet that is open and secure. Over the last few years, however, the proliferation of wide-scale network interference and surveillance has exploded. As a result, there has been growing interest in exploring how to detect and measure network interference, and understanding the various forms that this can take. While some interference activity is blatant and obvious, much of it is far more subtle, and by design nearly impossible for anyone but an an IT expert to detect, much less protest or mitigate.

In 2014, with support in part from the Knight Foundation, Open Technology Fund, and TIDES Foundation, OTI began researching network interference measurement to explore how it could support our larger research and policy objectives. This research consisted of experimenting with censorship measurement tools and analysis techniques, as well as convening researchers and tool developers already studying and experimenting in the space.

As a result of this research and convenings, some core themes emerged as areas of focus and improvement for the field:

1. **Community building**
   Given the nascent nature of the field, better communication is crucial to building a strong community. Communication facilitates establishing and enforcing norms around vetting methodologies, sanitizing and securing data, and other facets of research which are not currently formalized.

2. **Vetting methodologies**
   We need to cultivate greater interdisciplinary expertise. It is important to ensure that all new methodologies are vetted by people with sufficient interdisciplinary expertise to allow them to make meaningful determinations about what is responsible, ethical, and possible, and what's not.

3. **Informed consent**
   We need flexible guidelines and structures for determining when informed consent must be obtained, and how to do it. For example, when a study involves highly-technical interventions, adjustments must be made to facilitate the participation of people who have almost no prior experience with the subject and may not to understand the technical language.

4. **Data standards and sharing**
   Common standards, structures, and processes for collecting and analyzing data will allow for better sharing within the community and beyond, allowing for external validation of studies and comparable results across tools and processes. This should include standard practices for what personal metadata can be safely gathered and what needs to be scrubbed, limited, or not collected at all.

5. **Information and systems risk**
   Developing standards for study design that take into account a realistic and iterative assessment of risk to every participant in the study (human or otherwise) is critical. These standards need to incorporate the input of major stakeholders, including the provider(s) of infrastructure used for the study, institutional review boards (IRBs), and study subjects themselves.

Responsible and Ethical Best Practices for Network Interference Measurement

In 2014 and 2015, OTI and M-Lab staff convened or sponsored multiple gatherings of researchers and tool developers working in the field of interference measurement. The team held four convenings with researchers and tool developers, one workshop on interference measurement, co-led with the Oxford Internet Institute, and sponsored a pre-conference workshop at 2015 ACM SIGCOMM. At these meetings, we facilitated discussion, documented the state of the field, and sought to increase awareness of the need for further research and analysis of interference measurement tools and research.

In short, the systematic, rigorous study of network interference and digital censorship is a very new discipline – but one that is absolutely essential to the protection of civil liberties as people increasingly transition more of their lives to digitally-mediated spaces.

The people who were invited to these convenings are laying the foundation of a new field of research and development that will serve the critical function of protecting freedom of speech and access to information by detecting violations, and subsequently exposing them. OTI felt it was essential that this nascent community have regular opportunities to meet — to share best practices, explore new research questions, forge new collaborations, and examine and iteratively redefine itself to meet emerging use cases. Additionally, network interference research often conflates political speech with "scientific data" in ways that are tricky, and new to traditional technical fields. It is for this reason that interrogating and continually refining collectively held research and analysis methods is so important — it ensures that researchers and tool developers have shared approaches that can result in meaningful and diverse implementations that complement one another.

About half of the researchers attending these gatherings were already running experiments on the M-Lab platform, and the other half were not. These meetings provided an opportunity to discuss the base assumptions and optimal analytical methods of each experiment running on the M-Lab platform, interrogate and discuss current instruments for measuring interference or censorship, and consider broadly the use cases that improved or new measurement tests would address. The output of these workshop convenings was documented on a wiki as a part of a larger outreach campaign engaging additional researchers globally. It is important to OTI and M-Lab that these periodic gatherings continue, as an important venue to engage with the research community and encourage growth and support rigorous research practices.

Areas for Improvement Identified for the Field of Interference Measurement

An important outcome of the researcher and tool developer convenings discussed above was a set of "core areas for improvement" for practitioners in the field of interference measurement. These emerged over multiple meetings and were documented by OTI. The five core areas are

listed below and represent areas of focus that OTI, M-Lab, and others could pursue in the future, given continued interest and funding.

1. **Community building**
   Given the nascent nature of the field, better communication is crucial to building a strong community. Communication facilitates establishing and enforcing norms around vetting methodologies, sanitizing and securing data, and other facets of research which are not currently formalized.

2. **Vetting methodologies**
   We need to cultivate greater interdisciplinary expertise. It is important to ensure that all new methodologies are vetted by people with sufficient interdisciplinary expertise to allow them to make meaningful determinations about what is responsible, ethical, and possible, and what's not.

3. **Informed consent**
   We need flexible guidelines and structures for determining when informed consent must be obtained, and how to do it. For example, when a study involves highly-technical interventions, adjustments must be made to facilitate the participation of people who have almost no prior experience with the subject and may not to understand the technical language.

4. **Data standards and sharing**
   Common standards, structures, and processes for collecting and analyzing data will allow for better sharing within the community and beyond, allowing for external validation of studies and comparable results across tools and processes. This should include standard practices for what personal metadata can be safely gathered and what needs to be scrubbed, limited, or not collected at all.

5. **Information and systems risk**
   Developing standards for study design that take into account a realistic and iterative assessment of risk to every participant in the study (human or otherwise) is critical. These standards need to incorporate the input of major stakeholders, including the provider(s) of infrastructure used for the study, institutional review boards (IRBs), and study subjects themselves.

In summary, OTI's work in the area of interference measurement aligns with our work on broadband quality of access and surveillance, and the work described here represents our ongoing support of this nascent field. OTI continues to seek support for continuing work with researchers and tool developers working to identify the potential indicators of online censorship, while being pragmatic about the capability of any software tool or test claiming to detect those indicators. Our utmost concern in this area is acting ethically and responsibly by framing the discussion around user privacy and security and advocating for tools and research approaches that consider the implications of studying this issue on the individuals and groups who may be at risk for heightened online scrutiny.

Exploring the Detection of Network Interference Ethically

As a founding member of Measurement Lab (M-Lab), OTI develops and maintains the global platform that is the only source of open measurement data on network traffic. Tests hosted on the platform are initiated by individuals and the data are available via open source licensing. M-Lab hosts tests that measure everything from basic network data such as speed, quality, and routing to various forms of network interference. For example Glasnost and Neubot, two of the tests on the platform, both attempt to detect whether an ISP is performing application-specific traffic shaping.

In January 2015, OTI explored how to deploy a test to the platform that would focus on detecting technical means of network interference, such as HTTP header manipulation. This type of test compares the headers of an HTML file hosted on M-Lab servers when downloaded via standard means against the same file downloaded through a non-standard service.

For any test to de deployed on M-Lab, the developers of the test must answer questions about how the test works, what type of data is collected and published, as well as security questions related to what the test needs access to, or who needs access to the test. M-Lab's charter explicitly requires tests to be initiated by a user and not from our servers and everything, including the data, the tool, and the process, must be open source.

- *Standard Questions:*
    - How does the test work? Provide a brief technical summary.
    - What data must be collected on the M-Lab server and published?
    - What data does the test gather from the client?
    - What data does it gather on the server?
    - What is logged on the client?
    - Who has access to client's logs?

In order to hosts tests that look into more sensitive data such as interference, censorship or manipulation, the M-Lab team determined that additional precautions were needed to determine if the test would adhere to the M-Lab principles and not put users at risk.

- *Concern: M-Lab principles are to provide a public open data set to the M-Lab repository for public use.*
    - Are there privacy issues raised by the data?
    - Is any Personally Identifiable Information (PII) collected or stored?
    - Does any combination of the data collected, effectively constitute PII?
    - Could the data be paired with an external dataset and become PII?
    - In the test data files, are any client-specific metadata saved?
    - Is there any risk that submitted test results could be traced back to or used to identify individual clients?
    - Does the client's IP address get included in report files?
    - Are there any other kinds of potentially personally identifiable information saved in the test results?

- *Concern: M-Lab's standard mandate for data collection is that only IP address and active test data should be collected (i.e. a record of how the network reacted to a synthetic stream of data).*
    - How is anonymization of data done? Is it robust enough?
    - Where does the anonymization take place?
- *Additional Questions:*
    - On test submission, what validation is performed to confirm that the test is complete, or that it was submitted by the initiating client?
    - Can a malicious actor submit results without ever having run a test?
    - How are incoming reports validated/sanitized before the data is written to disk?
    - Are both the backend and the collection endpoint reasonably well-protected against malicious code that might be sent to them in requests/reports?

M-Lab is highly concerned with maintaining the anonymity of users, as well as their safety, which is sometimes at odds with what is technically possible. As a global platform, M-Lab must operate within constraints that help maintain the integrity of its measurement platform, being utilized by multiple national regulators as well as by non-governmental organizations working on Internet related issues. Measuring network interference is often conflated politically with "detecting censorship". In reality, network congestion or stability can exhibit similar behavior to some forms of censorship, so the primary distinction is whether or not poor network performance is the result of a political actor. As a result of understanding the role of interference tests on the M-Lab platform OTI and the M-Lab team focused on bringing researchers and tool developers together to discuss interference measurement generally and the ethical considerations of such work. As set out in the project goals, these gatherings provided the starting point to establish shared resources and to encourage the community working in this area to codify best practices and identify areas of improvement.