Tor and circumvention: lessons learned (Abstract to go with invited talk)

Roger Dingledine

The Tor Project

Tor is a free-software anonymizing overlay network that helps people around the world use the Internet in safety. Tor's 2500 volunteer relays carry almost 10Gb/s of traffic for several hundred thousand users each day.

While many in the research community know Tor as the primary fielded system in the anonymous communications literature [2], Tor has also played a central role in recent research on *blocking resistance*. That is, even if an anonymity system provides great anonymity, a government censor can render it moot by simply blocking the relays. In recent years we streamlined Tor's network communications to look more like ordinary SSL, and we introduced "bridge relays" that are harder for an attacker to find and block than Tor's public relays [1].

Tor played a key role in several Middle Eastern countries in early 2011. In this talk I'll walk the audience through how Iran used its Nokia DPI boxes to filter SSL flows that used Tor's original Diffie-Hellman parameter p; the surge in Tor traffic when Egypt blocked Facebook and the flatline when they unplugged the net; the continued bad news for Libya's Internet; and an intriguing trend in Saudi Arabia. I'll also cover current trends in China and Tunisia (not pictured).

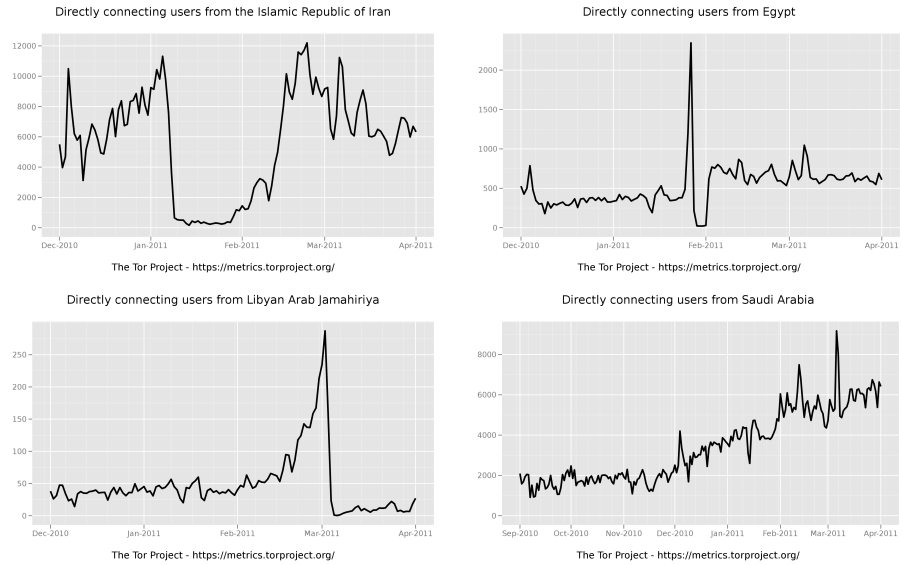

Fig. 1. Estimates of daily Tor clients connecting from each country

The data for these user graphs, along with historical Tor network data and ongoing performance statistics, are all available at https://metrics.torproject.org/. Our WECSR'10 paper [3] explains our aggregation techniques and why we think they're safe—we'd love for you to show us that we're wrong. Further, if you're working on Tor-related research, please talk to us (https://torproject.org/research) so we can explain what's available and help interpret your results.

Some open questions from the anonymity field.

Here are a few examples of open anonymity and blocking-resistance problems:

1) How effective is the traffic correlation attack really? Tor's threat model assumes that an adversary who can see a traffic flow into the Tor network and the corresponding flow out of the Tor network can correlate them with high probability and low false positives. Recent results from Steven Murdoch [4] show confirmation attacks even when both sides only see a small sample of traffic on each side. But how quick can the attack actually be in practice, using how little traffic? Are there effective padding schemes to make correlation less effective?

2) For various diversity metrics (like entropy), how has the diversity of the Tor network changed over time? How robust is it to change or attack?[1]

3) How can we automatically recognize blocking events—when Tor relays are censored at a firewall by destination address or by traffic flow characteristics?

4) Clients who are censored from the public Tor relays can use private addresses to "bridge" into the public Tor network. What strategies should we use to give out these addresses such that legitimate users get enough addresses but adversaries can't learn too many?

5) How can we make it hard for censors to recognize Tor traffic flows by content (e.g. distinguishing Tor's handshake from other expected protocols) and by traffic characteristics (packet size, volume, and timing)? We need *obfuscation* metrics to let us anticipate which protocols will blend in better with background traffic or otherwise defeat deep packet inspection (DPI) algorithms.

References

1. Roger Dingledine and Nick Mathewson. Design of a blocking-resistant anonymity system. Technical Report 2006-1, The Tor Project, November 2006.
2. Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proc 13th USENIX Security Symposium*, August 2004.
3. Karsten Loesing, Steven J. Murdoch, and Roger Dingledine. A case study on measuring statistical data in the Tor anonymity network. In Sven Dietrich, editor, *Proceedings of the Workshop on Ethics in Computer Security Research (WECSR 2010)*, LNCS 6054, Tenerife, Canary Islands, Spain, January 2010. Springer.
4. Steven J. Murdoch and Piotr Zieliński. Sampled traffic analysis by Internet-exchange-level adversaries. In Nikita Borisov and Philippe Golle, editors, *Proceedings of the Seventh Workshop on Privacy Enhancing Technologies (PET 2007)*, LNCS 4776, Ottawa, Canada, June 2007. Springer.

[1] https://blog.torproject.org/blog/research-problem-measuring-safety-tor-network