Tor: anonymity, anti-censorship, anti-surveillance (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 7000 volunteer relays carry over 48Gbit/s of traffic for around two million users each day.

In the research community, Tor is best known as the primary fielded system for anonymous communications [3]. Tor's anonymity comes from two components. The first is *distributed trust*: Tor clients build a path through three relays, such that no single relay learns about both the client and her destination. Thus the more relays there are, and the more diversity in terms of location and operators, the harder it is for an adversary of a given size to be in a position to match up clients to their destinations. The second component of Tor's anonymity is diversity of users and uses—ranging from ordinary citizens in Western countries to corporations, law enforcement and military, but also bloggers in oppressive regimes around the world. This diversity means that an adversary who learns that a given person is using Tor still doesn't learn *why* they are using Tor [1].

Aside from the anonymity side, Tor has also led 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 by address, or blocking the traffic flows using Deep Packet Inspection (DPI). Tor introduced *bridge relays* that are harder for an attacker to find and block than Tor's public relays [2], and also introduced *pluggable transports* that transform Tor traffic flows so they appear like more innocuous traffic [6].

Tor played a key role in several Middle Eastern countries starting in 2011. In this talk I'll cover how Iran used its Nokia DPI boxes to filter Tor flows, the surge in Tor traffic when Egypt blocked Facebook and the flatline when they unplugged the net, and stories from other countries around the world.

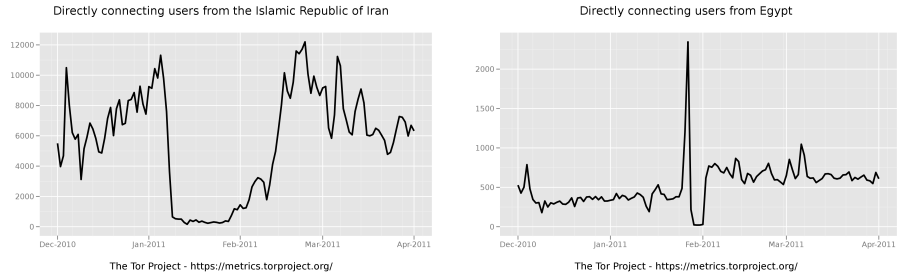

Fig. 1. Estimates of daily Tor clients connecting from each country

Tor and mass surveillance. The Snowden revelations over the past two years make it clear that Tor's anonymity side (traffic analysis resistance) and Tor's censorship circumvention side (DPI resistance) are more related than many people realize. An adversary who sees the traffic flow entering the Tor network, and the corresponding traffic flow exiting the Tor network, can use statistics to match them up [5]. But what are the realistic chances that a mass surveillance adversary can see both of these flows?

If they really did write down every packet on the Internet, Tor likely doesn't fare well. But surely long-term storage of every Internet packet is not practical. We could imagine instead that they write down flow summaries, or packet headers, or they have a rolling full packet capture of the last 72 hours of packets, meaning they have a limited amount of time to decide something is worth keeping before they're forced to discard it. But they can cut down the required storage space by starting out with a list of Tor relay IP addresses, and just decide that everything to or from them is worth keeping forever. It may even be practical to go out 2 hops or even 3 hops from the Tor relay IP addresses.

How can we make it hard for mass surveillance adversaries to recognize Tor traffic flows by address or content? Flashproxy [4] (originally designed as a pluggable transport for censorship circumvention) puts millions of transient IP addresses between Tor clients and the Tor network, giving a surveillance attacker an unmanageable number of addresses to track. Newer transports like meek (https://trac.torproject.org/projects/tor/wiki/doc/meek) which route traffic through Google, Akamai, or other sites could similarly stymie an attacker that tracks by address. This new *surveillance-resistance* field needs more attention.

References

1. Roger Dingledine and Nick Mathewson. Anonymity loves company: Usability and the network effect. In Ross Anderson, editor, *Proceedings of the Fifth Workshop on the Economics of Information Security (WEIS 2006)*, June 2006.
2. Roger Dingledine and Nick Mathewson. Design of a blocking-resistant anonymity system. Technical Report 2006-1, The Tor Project, November 2006.
3. Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proc 13th USENIX Security Symposium*, August 2004.
4. David Fifield, Nate Hardison, Jonathan Ellithorpe, Emily Stark, Roger Dingledine, Phil Porras, and Dan Boneh. Evading censorship with browser-based proxies. In *Proceedings of the 12th Privacy Enhancing Technologies Symposium (PETS 2012)*. Springer, July 2012.
5. 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.
6. The Tor Project. Tor pluggable transport specification. https://gitweb.torproject.org/torspec.git/tree/pt-spec.txt.