Enforced Community Standards For Research on Users of the Tor Anonymity Network*

Christopher Soghoian
Center for Applied Cybersecurity Research, Indiana University

Abstract

Security and privacy researchers are increasingly taking an interest in the Tor network, and have even performed studies that involved intercepting the network communications of Tor users. There are currently no generally agreed upon community norms for research on Tor users, and so unfortunately, several projects have engaged in problematic behavior – not because the researchers had malicious intent, but because they simply did not see the ethical or legal issues associated with their data gathering. This paper proposes a set of four bright-line rules for researchers conducting privacy invading research on the Tor network. The author hopes that it will spark a debate, and hopefully lead to responsible program committees taking some action to embrace these, or similar rules.

1 Introduction

Over the past few years, the Tor network has grown from an academic research project [4] to one of the most widely used privacy enhancing technologies, with several hundreds thousand of active users [9]. While little is known about the average Tor user, it is safe to assume that it is used by individuals seeking to protect their privacy, either denying their own ISP or government the ability to learn what they are doing online, or to stop websites from learning anything about their visitors. In order to achieve this degree of privacy protection, Tor's users pay a significant penalty, both in latency, as well as in general usability (as many popular plugins such as Flash must be disabled in order to prevent data leakage).

Just as privacy-seeking users have flocked to Tor, so too have researchers interested in learning more about its users and their use of the network. In some cases, these researchers specifically wish to observe Tor users in order to learn how it is being used. However, in others cases, the researchers simply seek to study general Internet behavior, and Tor is just a quick way to easily observe the traffic of thousands of Internet users – perhaps because major ISPs will not permit some kinds of traffic interception and network attacks on their customers, even in the name of research.

There are currently no widely accepted or publicized research community norms for studies on Tor and its users. As such, each team of researchers interested in studying the use of the network is left to determine what is right and wrong for themselves. While many researchers have gone out of their way to protect the privacy of Tor users as they collect data from the network [9], some are not getting it right, at least in the opinion of this author.

This paper will present two case studies in which researchers setup their own Tor servers, specifically in order to monitor the traffic that flows over the network. The paper will examine several ethical issues, and attempt to establish bright line rules for determining if the Tor network should be used to answer particular research questions. Finally, the paper will conclude by proposing that conference and workshop program committees play a strong role in establishing norms for this type of research, and enforcing these norms by rejecting papers that do not adhere to a few basic guidelines.

2 Prior Academic Studies on Tor Users

This section will summarize two academic research studies performed on the Tor anonymity network, one published in 2008, and one in 2010. These are not the only studies to involve the collection of data on the Tor network ([9] includes references to several others), but these papers are noteworthy in that they received strong post-publication criticism from the privacy community regarding the degree to which the researchers needlessly violated, or put at risk the privacy of Tor users. The purpose of this section is not to demonize the researchers, but to highlight the fact that the privacy community has failed to establish and enforce ethical norms for research studies that involve monitoring the Tor network.

2.1 Shining Light in Dark Places: Understanding the Tor Network

In 2008, McCoy *et al.* published the results of a study [11], which sought to determine the kind of traffic flowing over the Tor anonymity network [4]. In order to gather this data, the researchers setup a Tor *exit node* server on the University of Colorado's high-speed network, and added it to the publicly distributed list of Tor servers. During a four day period in December 2007, the researchers logged and stored the first 150 bytes of each network packet that went through their server their network. This revealed the kind of traffic that was crossing the Tor network, and the specific websites that users were accessing. In a second part of the study, the researchers ran an *entry node* to the network for fifteen days, which allowed them to determine the source IP address of a large number of Tor users. They used this to learn which countries use Tor more heavily than others.

Before starting their study, the researchers did not seek or obtanin a thorough evaluation of the legality of their activities. When later questioned by this author, one of the researchers stated that they "spoke informally with one lawyer, who told us that that area of the law is ill defined." Based on this, he said, the researchers felt that it was "unnecessary to follow up with other lawyers [14]." Similarly, the researchers did not seek the guidance and approval of their university's Institutional Review Board (IRB). "We were advised that it wasn't necessary," one of the researchers said, adding that the IRB review process is used "used more in medical and psychology research at our university," and was not generally consulted in computer science projects [14].

The researchers did not receive a warm welcome after presenting their work at the Privacy Enhancing Technologies Symposium. Several outspoken members of the academic privacy community were in the audience, as well as core developers of the Tor project, many of whom reacted harshly to the news that the researchers had monitored traffic on the network. As one example, when questioned by an audience member after the presentation, the researchers admitted that they had retained a copy of the logged Tor traffic, and further, that it was not held on an encrypted storage device. This disclosure was met with boos from the audience, even after the researchers stressed that the data was kept in a "secure" location [14].

Within days of the researchers' presentation, the University of Colorado announced that a post-review of the project had determined that the researchers did not violate university policies, specifically finding that:

> "Based on our assessment and understanding of the issues involved in your work, our opinion was that by any reasonable standard, the work in question was not classifiable as human subject research, nor did it involve the collection of personally identifying information. While the underlying issues are certainly interesting and complex, our opinion is that in this case, no rules were violated by your not having subjected your proposed work to prior [IRB] scrutiny. Our analysis was confined to this [IRB] issue [10]."

2.2 Private Information Disclosure from Web Searches

In 2010, Castelluccia *et al.* revealed a privacy flaw in several major search engines, in which an attaker can use a sniffed authentication cookie to reconstruct a users' search query history [3]. In addition to

demonstrating the flaw, the researchers also sought to determine the degree to which users are vulnerable, that is, how many users conduct web searches when "logged in" to a search engine, and how many of them have enabled Google's Web History feature. In order to determine this information, the researchers collected data via three different methods: First, network traces for the 500-600 daily users at their own research center were collected and analyzed. Second, the researchers established a Tor exit node server, and examined the network traffic exiting from it. Third, the researchers received opt-in consent from 10 users, whose Google session cookies the researchers sniffed, and then used to actively reconstruct the individuals' search history information.

During the one week period in which the researchers collected data from the Tor network, 1803 distinct Google users were observed, 46% of which were logged into their accounts. For each of these logged-in users, the researchers used the sniffed Google session cookies and attempted to access the users' first and last name; locations searched using Google Maps (along with the "default location", when available); blogs followed using Google Reader; full Web History (when accessible without re-entering credentials); finance portfolio; and bookmarks. In their paper, the researchers stress that their research application did not store any individual users' data – only aggregate statistical information was retained.

The researchers treated the three groups of users (the volunteers, co-workers at their research center and Tor users) quite differently. For example, the researchers did not actively attack the accounts of their colleagues, they merely passively analyzed the network traces, whereas users of the Tor network had their accounts actively attacked, and some of their data downloaded from Google's servers (although not retained). In their paper, the researchers do not reveal the reason for the restraint they showed in choosing to not actively attack their coleagues' accounts.

Similarly, the researchers did not actively probe the search history of either their colleagues' accounts or the Tor users, and restricted the use of this attack to just the 10 volunteers who had consented to assist with the study. The researchers describe the motivation for this difference, writing in their paper that "it would have been otherwise impossible to conduct our study on uninformed users without incurring legal and ethical issues." It is unclear from the content of the paper why the researchers found it ethically acceptible to actively attack Tor users' Google accounts, but not to download their search history.

When the researchers presented their paper at the Privacy Enhancing Technologies Symposium in 2010, they received a similar reaction from the audience as McCoy *et al.* had in 2008. The reaction of the audience is not terribly surprising, given that most of the people attending the conference spend their time working to protect users' privacy. What is surprising, and extremely relevant to the focus of this paper, is that the researchers presenting their paper in 2010 had not learned about the strong reaction from the community to the paper by McCoy *et al.* presented at the very same conference two years earlier.

2.3 Analyzing and Comparing the Two Studies

The academic privacy and security community can learn a few things by contrasting these two research papers. First, the published proceedings from the 2008 and 2010 Privacy Enhancing Technologies Symposia include the McCoy *et al.* and Castelluccia *et al.* papers, but nothing documenting the strong reactions from the audience. As such, any future researchers looking through previously published papers in this community may reasonably believe that such studies are appropriate, and blessed by the community.

Second, it is quite easy to differentiate between the McCoy and Castelluccia studies. The former specifically sought to learn more about users of the Tor network, whereas the latter simply used Tor users' network activity to assist in drawing broader conclusions about general Internet behavior. The fact that Castelluccia *et al.* performed only passive network monitoring on their own colleagues but actively attacked the accounts of Tor users likely indiciates that the researchers knew they were engaging in morally and ethically dubious behavior. If there were no problems with what they were doing, why would they not do it to their friends and colleagues, but were willing to do it to users who had specifically signaled a desire to protect their own privacy?

Third, neither group of researchers submitted their studies for Institutional Review Board approval –

McCoy *et al* did not believe they had to, while Castelluccia *et al.* did not have an IRB at their research institution.

Finally, Castelluccia *et al.* specifically designed their research tool to analyze individual users's data in-memory, and only retained aggregate statistical data. On the other hand, McCoy *et al.* retained individual users' browsing data, and performed statistical analysis of it after the fact. The former approach is clearly more privacy preserving, but the latter is more resistent to researcher-error. That is, had Castelluccia *et al.* made a mistake in their code, they would have had to collect new user data in order to analyze and aggregate it. By retaining individual users' data McCoy *et al.* were free to tweak their code as much as they wanted, as they could always re-run it against their previously colected data.

3 Towards A Community Standard

Program committees can and should play a major role in both establishing and enforcing community standards for research. Even if just one or two conferences establish and publicize such rules, it will send a clear signal to researchers and help them to take appropriate steps to protect user privay as they design their studies. Furthermore, since most of the Tor related research seems to be published at the Privacy Enhancing Technologies Symposium, it is likely that a strong set of community norms can be established via the decision of a single program comittee. In this section, I propose four bright-line rules for Tor related research – these are not exhausive, and it is still quite possible for researchers to meet these guidelines and still engage in irresponsible, privacy invading research. However, should researchers follow these rules, they should at least be able to avoid several privacy pitfalls present in earlier research studies.

Research should be focused on users of the Tor network Researchers seeking to gather Tor network usage data should be specifically focused on studying users of the Tor network, and should not be using Tor as as a convenient method of studying general Internet users' activity online. Researchers may be tempted to establish their own Tor exit node, as it is a very quick way of getting access to the Internet traffic of thousands of users. This may be a particularly attractive option for those researchers without close ties to a large Internet service provider, as well as for those researchers whose academic instutitions will not permit them to conduct the study on their colleagues and students. In spite of this temptation, it is simply not appropriate to violate the privacy of Tor users just because it is easier to do so than to get approval to monitor the network at one's own university. If the privacy of Tor users is to be intruded upon, it should at least be to answer questions specific to the Tor community, and not something that could be learned another way.

Minimize user data collection and retention Researchers should ensure that user data is examined in-memory only, and that the only data retained is aggregate in nature. The researchers should not put themselves in a position where they could be later compelled (by law enforcement agencies, for example) to disclose any identifable data either about specific users (such as originating IP addresses), or the specific web sites and web pages that Tor users visit.

Ensure that the research study is legal in the country where it is performed There are significant questions surrounding the legality of much network monitoring research, particularly when it is conducted in the United States, where communications privacy and interception law is exceedingly complex [12]. Computer scientists are simply not equipped to evaluate the legality of the research they perform, and as such, it is important that researchers seek the assistance of qualified legal experts as they design their studies. Program committees should require that the researchers identify the legal expert with whom they consulted, and should independently contact the named legal expert in order to verify that they do indeed believe that the researchers' study did not violate the law.

Research studies should be vetted by an IRB, if one exists While Institutional Review Boards exist at most research universities in the United States, they are far less common in many other countries. It is certainly true that there are legitimate concerns about the lack of technical expertise on many IRBs, however, these will lessen over time, as more and more computer scientists interact with IRBs. Furthermore, even if the IRB does not provide much in the way of useful technical oversight, the self-evaluation that the researchers have to perform as part of the review process (listing the kinds of possible harms that test subjects may face, and the steps they have taken to mitigate them) may be be useful.

4 Related Work

Loesing *at al.* presented two case studies in which data was gathered from the Tor network in a responsible, and privacy-preserving manner [9]. Drawing from these case studies, the researchers proposed three general guidelines for future Tor data collection: data minimalism, source aggregation and transparency. The researchers goal for the paper was to start a discussion, but they do do not call for enforcement of these rules.

Dittrich *et al.* proposed an ethical framework to guide and evaluate applied security research, motivated by a frustration among researchers, program committees, and professional organizations over the current state of affairs [6, 5]. Their goal too was to encourage a dialog, which would hopefully lead to some form of community consensus.

Allman examined the role that confererence program committees may play in guiding researchers towards ethical research methodologies [1]. Allman does not however propose a clear set of rules that program committees should adopt. Likewise, Landwehr has called on professional societies to develop ethical guidelines for their members who are facing these issues [8].

Allman, Garfinkel and Landwehr [1, 8, 7] all suggested that Institutional Review Boards may play a positive role, but have voiced concerns about the degree to which IRB members lack enough computer security skills and an awareness of the existing values of the computer science community to to effectively judge the risks involved in such research.

Sicker *et al.* [12] outlined several areas of potential legal liability for researchers engaging in network monitoring research. The authors strongly encourage the broader network monitoring community to establish community norms, but do not suggest what these norms should be. On a similar note, Burstein and Soghoian each offered specific recommendations to security researchers engaging in cybersecurity and phishing research in order to avoid specific legal pitfalls [2, 13].

5 Conclusion

In this paper, I have proposed four easy, bright-line rules that can be used to evaluate and guide researchers seeking to engage in studies involving the Tor anonymity network. The community has lagely failed, thus far, to establish and enforce any standards for this type of research. Both of the problematic research projects summarized in this paper were published at a highly ranked peer reviewed conference. This creates two incentive problems: the researchers who conducted the earlier studies pay no real long-term price for recklessly violating the privacy of Tor users, and future researchers who read through the published conference proceedings may be reasonably lead to believe that the methods employed in these studies are legitimate and blessed by the community.

The community must promptly agree upon, establish and enforce a set of easy to understand guidelines for acceptible Tor research (those presented in this paper, those created by Loesing *at al.*, or a different set). Future Tor related research projects that violate these guidelines should be rejected from the Privacy Enhancing Technologies Symposium, as well as other top-tier privacy and security conferences.

References

[1] Mark Allman. What ought a program committee to do? In *Proceedings of the conference on Organizing Workshops, Conferences, and Symposia for Computer Systems*, pages 9:1–9:5, Berkeley, CA, USA, 2008. USENIX Association.

[2] Aaron J. Burstein. Conducting cybersecurity research legally and ethically. In *Proceedings of the 1st Usenix Workshop on Large-Scale Exploits and Emergent Threats*, pages 8:1–8:8, Berkeley, CA, USA, 2008. USENIX Association.

[3] Claude Castelluccia, Emiliano De Cristofaro, and Daniele Perito. Private information disclosure from web searches. In Mikhail J. Atallah and Nicholas J. Hopper, editors, *Privacy Enhancing Technologies*, volume 6205 of *Lecture Notes in Computer Science*, pages 38–55. Springer, 2010.

[4] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: the second-generation onion router. In *Proceedings of the 13th conference on USENIX Security Symposium - Volume 13*, SSYM'04, pages 21–21, Berkeley, CA, USA, 2004. USENIX Association.

[5] David Dittrich, Michael Bailey, and Sven Dietrich. Have we Crossed the Line? The Growing Ethical Debate in Modern Computer Security Research. In *(Poster at) Proceedings of the 16th ACM Conference on Computer and Communication Security (CCS '09)*, Chicago, Illinois, USA, November 2009.

[6] David Dittrich, Michael Bailey, and Sven Dietrich. Towards community standards for ethical behavior in computer security research. Technical Report 2009-01, Stevens Institute of Technology, Hoboken, NJ, USA, April 2009.

[7] Simson L. Garfinkel. Irbs and security research: myths, facts and mission creep. In *Proceedings of the 1st Conference on Usability, Psychology, and Security*, pages 13:1–13:5, Berkeley, CA, USA, 2008. USENIX Association.

[8] Carl E. Landwehr. Drawing the line. *IEEE Security and Privacy*, 8:3–4, 2010.

[9] Karsten Loesing, Steven Murdoch, and Roger Dingledine. A case study on measuring statistical data in the tor anonymity network. In *Financial Cryptography and Data Security*, volume 6054 of *Lecture Notes in Computer Science*, pages 203–215. Springer Berlin, 2010.

[10] Damon McCoy, Kevin Bauer, Dirk Grunwald, Tadayoshi Kohno, and Douglas Sicker. Response to tor study, July 25 2008. `www.verisign.com/static/039933.pdf`.

[11] Damon Mccoy, Kevin Bauer, Dirk Grunwald, Tadayoshi Kohno, and Douglas Sicker. Shining light in dark places: Understanding the tor network. In *Proceedings of the 8th international symposium on Privacy Enhancing Technologies*, PETS '08, pages 63–76, Berlin, Heidelberg, 2008. Springer-Verlag.

[12] Douglas C. Sicker, Paul Ohm, and Dirk Grunwald. Legal issues surrounding monitoring during network research. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, IMC '07, pages 141–148, New York, NY, USA, 2007. ACM.

[13] Christopher Soghoian. Conducting cybersecurity research legally and ethically. In *Proceedings of eCrime Researchers Summit*, 2008.

[14] Christopher Soghoian. Researchers could face legal risks for network snooping. *Surveilance State*, July 24 2008. `news.cnet.com/8301-13739_3-9997273-46.html`.