Errata for Characterizing the Nature and Dynamics of Tor Exit Blocking

Rachee Singh[1], Rishab Nithyanand[2], Sadia Afroz[3,4]
Paul Pearce[3], Michael Carl Tschantz[4], Phillipa Gill[1], Vern Paxson[3,4]

[1]*University of Massachusetts – Amherst,* [2]*Iowa State University,*
[3]*University of California – Berkeley,* [4]*International Computer Science Institute*

Abstract

In our USENIX Security 2017 paper, we reported a method for distinguishing between Tor relays being blacklisted in a proactive and a reactive manner. However, we have since concluded that this method can draw false conclusions.

1 Introduction

The IP addresses of Tor relays often end up on blacklists. This could be caused by the relay actually being involved in some abusive behavior that got reported to a blacklist, which reacted to the report. Alternatively, it could be caused by the blacklist proactively listing every relay, or at least exit relay, in the Tor consensus.

In our prior work, Singh et al. [4], which appeared at USENIX Security 2017, we reported a method for distinguishing between these two cases. This errata explains how this method can draw false conclusions. In particular, that method made assumptions about the speeds of various methods of blacklisting that may not hold.

Singh et al. provide numerous results. (Henceforth, we will use the third person to refer to our prior selves.) We believe the errors described herein are mostly limited to Section 5.2 of Singh et al., although these errors affect some comments about the nature of blacklists made in Sections 5.3, 5.4, and 5.6 as well. We know of no major issues in the other parts of the paper.

We start with background and a summary of the problematic section of Singh et al. Readers with that work fresh in their mind can skip ahead to Section 4, where we discuss the problems. We end with some thoughts about the difficulties of producing a more accurate analysis.

2 Background

To provide anonymity, Tor obscures the connection between a website visitor and the website's server by routing it through *relays*, that function similarly to proxies. An *exit relay* serves as the last Tor relay in the route from the visitor to the website. By being the relay visible to the website, they appear to be the cause of any abuse that website receives via Tor from the visitor. Thus, the exit relays may face complaints, IP blacklisting, take-down notices, and even visits from law enforcement. Given this, only a relatively small number of relay operators allow their relay to serve as an exit. To enable the creation of routes from website visitors to websites, the *Tor consensus* publicly announces all the relays and which are willing to serve as exits.

If the IP address of a Tor exit relay ends up on a blacklist, any Tor user exiting through that relay will be blocked by websites using that blacklist. The paucity of exit relays combined with them being shared across all Tor users makes fate sharing and such blacklisting a serious concern. Unfortunately, some people visit websites using Tor to engage in the exact sorts of behavior that gets one blacklisted. As a result, Tor exit relays often end up blacklisted and Tor users, blocked. Indeed, prior work has found many websites unreachable by Tor users through many of the exit relays [2, 4].

The Tor community has taken steps to reduce abusive Tor use, not only to be good netizens, but also to reduce the number complaints exit relay operators receive and the amount of blacklisting and blocking of Tor users. For example, many Tor relay operators block port 25 to reduce their use for sending spam [5]. While any reduction in abuse is good, to maximize the reduction in blacklisting, requires some understanding how blacklisting happens.

3 Summary of Singh et al.'s Analysis

Herein, we summarize Section 5.2 of Singh et al. Singh et al. distinguishes between two forms of Tor relay blacklisting [4]:

- *Reactive* blacklisting happens when an IP address is added to a blacklist due to the blacklist maintainer receiving a report of abuse coming from the IP address.
- *Proactive* blacklisting happens when a blacklist maintainer adds an IP address in response to seeing it listed as a Tor relay, for example, in the Tor consensus.

Figure 1 shows the functioning of each.

Singh et al. developed and used a method of classifying blacklists into those that are reactive and those that are proactive. They classify four lists as proactive based upon the lists documentation. They then attempted determine whether the remaining blacklists were proactive or reactive by measuring the speed at which each list added Tor relays. If a blacklist added many relays very quickly, their method classified it as proactive. Otherwise, the method labeled it as reactive.

They measure the speed of black listing by looking at the difference between a Tor relay entering the consensus and appearing in the blacklist (§5.2). The Tor consensus is publicly available and achieved, allowing determining the time at which a relay entered the consensus. To determine the time at which a relay was blacklisted, they used a data set that logged the status of numerous IP threat intelligence feeds, including those found in Facebook's ThreatExchange [1] in particular (§5.1).

Figure 2 shows CDFs for the time lags between a relay appearing in the consensus and being blacklisted for two different threat feeds from Singh et al.'s paper (fig. 2). They classify CDF 2(a) as fast and the threat intelligence feed as proactive since most relays were quickly added, as visible from the sharp raise and extreme slop of the CDF. They classify the CDF 2(b) as reactive due to its the CDF's more gradual raise. This classification looks not just at the typical time lag, which is short, but at the tail where the slop flattens out but is not zero. They reasoned that such a tail would not exist for a proactive list since the list operator can easily quickly add all new relays simply by checking the Tor consensus regularly. They conclude since the list is unlikely to be a reactive, it is probably reactive.

4 Errors

Ian Goldberg questioned Singh et al.'s classification method (during the Q&A following the paper's presentation). He pointed out that one of the blacklists classified as reactive due to slowly adding Tor relays was blacklisting both Tor exit relays and non-exit relays at the same rate (see Singh et al.'s Figure 2(b) – reproduced herein as our Figure 2(b)). Since we would expect abuse complaints, the hallmark of reactive blacklisting, to be more common for exit relays than non-exit relays, shouldn't we expect that reactive blacklists would add exits more quickly than non-exits?

While it turns out that Tor abuse can come from non-exit relays [3], it continues to seem unlikely, to us, that exit and non-exit relays would be blacklisted at the same rate for a reactive list. Looking into how to answer Goldberg's question, we found two errors with our work.

Error 1. The match up between the slopes seems close to perfect as to improbable for any data set with as much noise as we would expect for this one. With this in mind, we re-computed the CDF getting somewhat different looking results, shown in Figure 2(c). We have not been able to track down the source of this difference. Maybe a bug in the old analysis code resulted in the same data being used for both CDFs shown in the old figure and something like floating point errors adding the tiny amount of jitter between them.

Our new results shows a similar drawn-out slope for the tail of each type of relay. This time, the slopes for the two types of relays are not so similar as to suggest an impossibility. However, they are still similar enough to that Goldberg's question remains salient.

Error 2. Our main concern is our realization that Singh et al.'s analysis implicitly assumes that because proactive blacklists lists can quickly add any new Tor relay, they will do so. This ignores the possibility of a slow-moving proactive blacklist. The operator of a proactive blacklist may be slow at its job and only rarely download the Tor consensus.

In fact, a slow proactive list can produce drawn-out tails and moderate slopes that Singh et al. used to label lists as reactive. For example, consider a simplified case where relays join the consensus at uniformly random times such that each hour the probability of seeing a new relay in the hourly consensus is p. Suppose a slow proactive blacklist adds relays from the consensus every h hours. In expectation, it will find $p*h$ new relays. The time lag for these new relays are equally likely to be any number between 0 and h. Thus, the CDF will, in expectation, have a shape $y = 1/h*x$ for $x \leq h$ and then flattening off at $y = 1$ for $x > h$. This simple model produces a line that approximates the CDF for exits with h at a bit over 9000, suggesting a blacklist update about once a year.

Similarly, Singh et al.'s analysis implicitly assumes that all reactive blacklisting will be slow. If Tor is experiencing high levels of abuse and a reactive blacklist quickly receives and processes complaints about every relay, it may add relays shortly after they join the consensus.

Thus, we conclude that the time lag for blacklisting is of limited value in distinguishing between proactive and

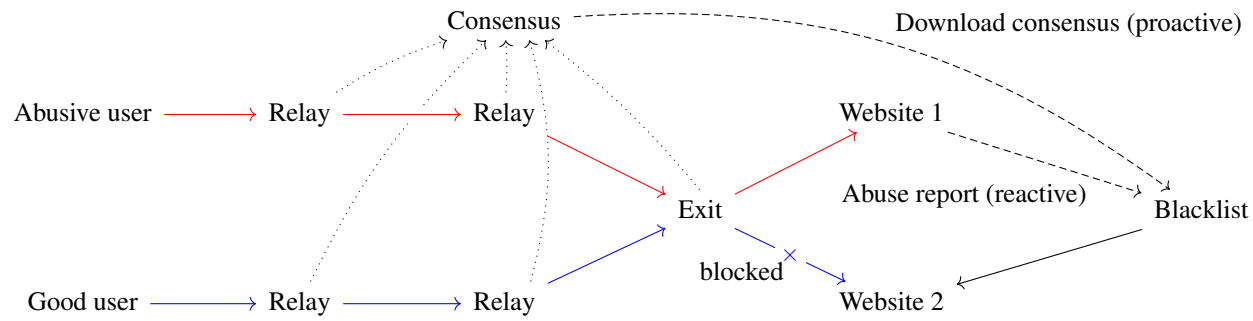

Figure 1: The Tor Network and Associated Players. The good user finds that Website 2 will not accept his connection via the Tor exit relay since it is on the blacklist used by the website. This could either be because the blacklist includes all exit relays from the Tor consensus (proactive) or because a different website reported earlier abuse transiting the exit. These two possibilities are shown as dashed arrows.

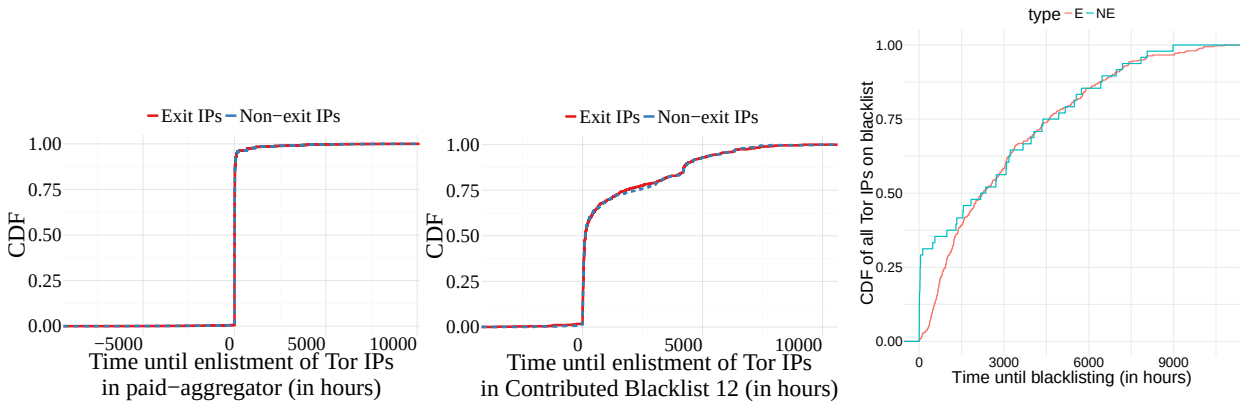

(a) Labeled proactive by Singh et al.

(b) Labeled reactive by Singh et al.

(c) Our recomputation of 2(b)

Figure 2: Reprints of two graphs from by Singh et al. [4] and a recomputation of one. They each give the time (in hours) between first seeing the IP address in the consensus until the blacklist adds the IP address. Negative values indicate the IP address being blacklisted before appearing in the consensus.

reactive blacklisting. Very fast blacklisting does seem to suggest proactive listing since it is hard to believe that reactive blacklists could receive and process abuse reports beyond some speed. Very slow blacklisting seems to suggest reactive listing since we would expect some reasonable rate of list updates. However, it is difficult to put numbers to these thresholds, and everything in between these extremes could be explained with either listing class.

5 Future Work

We took a second look at the our data set in hopes finding a better way of distinguishing between different types of blacklisting. So far, we have not found one. Rather, we have come away with the conclusion that many blacklists appear to be neither proactive nor reactive in a straightforward manner.

Perhaps, some use a bit of both methods. For example, consider a blacklist that only adds Tor relays when it gets complaints about one (like reactive ones), but instead of just adding the one exit relay, adds all exits (like proactive ones), or even adds all relays (exit and non-exit). Such blacklisting may result from an maintainer who only gets around to re-downloading the consensus upon receiving complaints.

Acknowledgments

We thank Ian Goldberg for asking us about our results, prompting us to look into these issues. The authors would like to thank Facebook Threat Exchange for providing threat intelligence data. We acknowledge funding support from the Open Technology Fund and NSF grants CNS-1237265, CNS-1518918, CNS-1406041, CNS-1350720, CNS-1518845, CNS-1422566. The opinions in this paper are those of the authors and do not necessarily reflect the opinions of a sponsor or the United States Government. All errors in our work are our own.

References

[1] Facebook Threat Exchange. `https://developers.facebook.com/products/threat-exchange`.

[2] KHATTAK, S., FIFIELD, D., AFROZ, S., JAVED, M., SUNDARESAN, S., PAXSON, V., MURDOCH, S. J., AND MCCOY, D. Do you see what I see?: Differential treatment of anonymous users. In *Network and Distributed System Security Symposium 2016* (2016), IETF.

[3] MASERGY. [tor-relays] abuse email for non-exit relay (masergy). Archived email to the Tor Relays mailing list: `https://www.mail-archive.com/tor-relays@lists.torproject.org/msg12097.html`, Sept. 2017. See also the linked to replies.

[4] SINGH, R., NITHYANAND, R., AFROZ, S., PEARCE, P., TSCHANTZ, M. C., GILL, P., AND PAXSON, V. Characterizing the nature and dynamics of Tor exit blocking. In *USENIX Security* (Aug. 2017).

[5] THE TOR PROJECT. What about spammers? Abuse FAQ. `https://www.torproject.org/docs/faq-abuse.html.en#WhatAboutSpammers`.