Anonymity on QuickSand: Using BGP to Compromise Tor

Laurent Vanbever, Oscar Li, Jennifer Rexford, and Prateek Mittal
Princeton University

ABSTRACT

Anonymity systems like Tor are known to be vulnerable to malicious relay nodes. Yet, we argue that the Autonomous Systems (ASes) delivering the traffic also present a serious threat, due to their powerful eavesdropping capabilities. An AS (or set of colluding ASes) that lies between the client and the first relay, and between the last relay and the destination, can perform timing analysis to compromise user anonymity. In this paper, we show that AS-level adversaries are much more powerful than previously thought. First, routine BGP routing changes can significantly increase the number of ASes that can analyze a user's traffic successfully. Second, ASes can actively manipulate BGP to put themselves on the paths to and from relay nodes, to gain visibility into user traffic. Third, an AS can perform timing analysis even when it sees only one direction of the traffic; as such, asymmetric routing ironically *increases* the fraction of ASes able to analyze a user's traffic. We present a preliminary evaluation of our attacks using measurements of BGP and Tor. Our findings motivate the design of approaches for anonymous communication that are resilient to AS-level adversaries.

1. INTRODUCTION

Due to increased surveillance of online communications, *anonymity* systems have become a key privacy-enhancing technology. Anonymity systems enable users to communicate privately by hiding their identities from the recipient or third parties on the Internet. For example, the Tor network [14] is a popular anonymity system that serves millions of users every day, and is comprised of over 5000 volunteer relays [1]. Tor is used today by journalists, whistle-blowers, political dissidents, military, intelligence agencies, law enforcement, and businesses, as well as ordinary citizens [2].

Although Tor traffic is encrypted, an adversary that controls both the first and last relays in the path can deanonymize a user by correlating the timing and size of the packets. However, we argue that the large Autonomous Systems (ASes) that carry the traffic are a more serious threat. An AS, like an Internet Service Provider (ISP), can easily eavesdrop on a portion of all links, and observe any unencrypted information, packet headers, packet timing, and packet size. Recent news revelations by Edward Snowden have shown that ASes pose a realistic threat: the NSA's Marina program stores vast amounts of metadata for up to a year [10], while the GCHQ's Tempora program buffers data for three days and metadata for 30 days [23]. In particular, Tor has been specifically targeted by such adversaries in collusion with ASes [4, 6, 7].

Tor's vulnerability to AS-level adversaries has received relatively little attention [15, 17, 21, 26]. Prior work focuses on (a) inferring the *static* AS-level paths to compute the chance of a single AS observing traffic between the client and the first relay, and the last relay and the destination, or (b) on selecting paths to minimize the risk of timing-analysis attacks. But this is just the tip of the iceberg. Interdomain (BGP) routes *change* over time, placing even more ASes in a good position to perform timing analysis. These routing changes happen naturally due to equipment failures and changes in routing policies. The problem grows more serious if we consider ASes that actively manipulate BGP [11, 33] to gain strategic visibility into remote communications of Tor relays. For example, an active adversary could launch a prefix hijack attack to take ownership of a Tor relay's IP prefix, or launch a prefix interception attack to become an intermediate AS for traffic destined to a Tor relay.

Interception attacks are especially dangerous, because the traffic continues to flow between the communicating hosts. BGP interceptions have become increasing common in recent years [5]. In one high-profile example, China Telecom intercepted traffic for tens of thousands of IP prefixes all over the world for around 18 minutes [3]. During this time, China Telecom would have seen packets destined to any Tor relay nodes in these address blocks. Naturally, China Telecom can always see the traffic its own customers exchange with Tor guard nodes. By putting this information together, an AS has sufficient data for an accurate timing analysis on that traffic. Of course, it is not easy to know whether BGP interceptions are intentional or accidental; the more important point is that interceptions substantially increase an AS's ability to determine what sites Tor users are accessing.

While BGP churn and BGP attacks are well known in the networking community, their impact on the security of anonymity systems like Tor is not well understood. We show that BGP routing changes, whether incidental or intentional, can have devastating consequences for user anonymity. In addi-

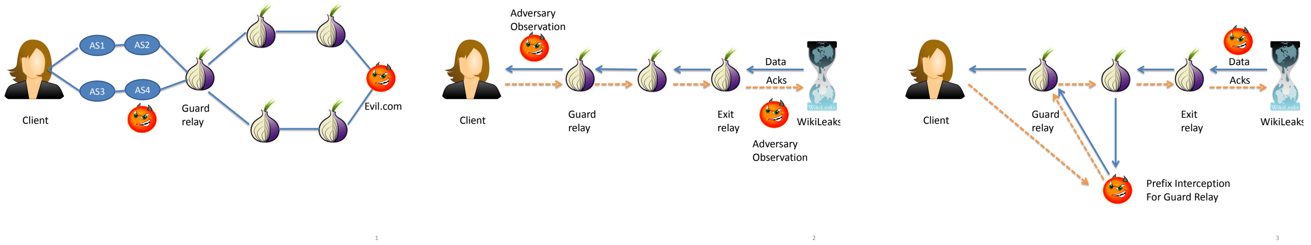

(a) Path changes degrade anonymity (b) Asymmetric traffic analysis attack (c) Active BGP manipulations to intercept traffic

Figure 1: AS-level Traffic Analysis. We show that (a) natural routing changes and (c) malicious BGP manipulations increase AS surveillance capabilities, enabling it to deanonymize Tor clients using timing/traffic analysis and (b) it suffices for an adversary to observe communication at both ends in *any* direction.

tion, we argue that the adversarial AS (or set of ASes) need only see one direction of the traffic at each end of the communication. For example, an adversary could correlate data packets from the client to the first relay with TCP acknowledgments from the server to the last relay. As such, these attacks are effective even under asymmetric routing. In fact, asymmetric routing only *increases* the security risk, by increasing the number of ASes that lie on some path (either forward or reverse) at each end of the communication.

We quantify the threat of AS-level adversaries based on a preliminary analysis from real-world Tor and BGP data, and also propose countermeasures for our traffic-analysis attacks. Overall, our work motivates the design of new approaches for anonymous communication, that account for the powerful capabilities of AS-level adversaries.

2. TOR BACKGROUND

The focus of this work is on low-latency anonymity systems such as the Tor network [14]. Low-latency designs are suitable for interactive communications on the Internet, as they do not inject any timing delays, but are also vulnerable to timing analysis attacks. The Tor network is a popular deployed system for low-latency anonymous communication that serves millions of clients a day, and carries 8 GBps traffic. As of July 2014, the Tor network comprises more than 5000 volunteer relays all over the world [1]. Tor clients first download information about Tor relays (called network consensus) from directory servers. Then Tor clients select three relays for anonymously forwarding users' traffic to the destination (source routing). Layered encryption is used to ensure that each relay learns the identity of only the previous hop and the next hop in the communications, and no single relay can link the client to the destination. Moreover, to load balance the network, clients select relays with a probability that is proportional to their network capacity.

Threat model and conventional attacks: *End-to-end Timing Analysis*: It is well known that if an attacker observes encrypted traffic from a client to the first relay as well as from the final relay to the destination (or traffic from the destination to the final relay and from the first relay to the client), then it can leverage correlation between packet timing and sizes to infer the identities of clients and destinations. Typical security analysis of Tor mostly considers the threat of end-to-end timing analysis due to malicious relays. Note that this attack requires the adversary to insert a large number of malicious relays in the Tor network, and has some fundamental limitations discussed below. In this work, we focus on the threat posed by AS-level adversaries. In particular, an AS-level adversary can launch passive attacks, and is also capable of certain types of active attacks.

Long-term anonymity: When users' communicate with recipients over multiple time instances, then there is a potential for compromise of anonymity at every communication instance [27, 29]. Therefore, the anonymity protection received by users degrades over time. Prior work considered this threat from the perspective of malicious relays: to defend against such long-term attacks, Tor clients choose their first hop relay from a small set of three relays (called *guards*). The set of three guard relays per client is kept fixed for about a month. [1] Without the use of guard relays, the probability of user deanonymization approaches 1 over time. With the use of guard relays, if the chosen guards are honest, then the user cannot be deanonymized for the lifetime of guards. Some of our attacks rely on the observation that even if the set of guard relays for a client stay the same across communication instances, the set of ASes on the paths between the client and the guard relays does change!

3. AS-LEVEL TRAFFIC ANALYSIS

In this section, we show how ASes can exploit natural BGP dynamics, or even launch active attacks, to compromise the anonymity of Tor users. We then discuss how seeing just one direction of the traffic for each segment (between the sender and the guard, and between the last relay and the destination) is sufficient for the adversary.

3.1 Exploiting Natural Temporal Dynamics

[1] The Tor Project is considering increasing the duration of the time period to 9 months [13].

When communicating with recipients multiple times, a user's traffic is susceptible to adversarial analysis at each communication instance. Thus, anonymity can degrade over time. Tor's use of guard relays defends against this threat with respect to adversarial relays, but *not* against AS-level adversaries. The underlying Internet paths between a client and guard relay vary over time due to changes in the physical topology (e.g., failures, recoveries, and the rollout of new routers and links) and AS-level routing policies (e.g., traffic engineering and new business relationships). These changes give a malicious AS surveillance power that increases over time! For example, AS 4 in Figure 1(a) does not lie on the original path from the client to the guard, but a BGP routing change can put AS 4 on the path for some period of time.

Let us suppose that the probability of any AS being malicious is f, and that the set of malicious ASes collude. Also, let us suppose that there are n AS-level paths between a client and a particular guard relay comprising x distinct ASes. Then, over time, the adversary's probability of observing the client's communication with the guard approaches $1-(1-f)^x$, i.e., the probability that at least one of out of the x ASes is malicious. Observe that this probability increases *exponentially* with the number of ASes (x)! Moreover, the Tor network uses multiple guard relays for improved availability (set to three guard relays in the current implementation). The average probability of an adversary observing communications between a client and any of the l guard relays is computed as $1-(1-f)^{l \cdot x}$. Thus, we can see that the impact of temporal dynamics (which increases the value of x) is further amplified due to the use of multiple guard relays.

The severity of the problem depends on the frequency of routing changes—and the diversity of ASes on these paths. Our preliminary measurement study in Section 4 shows that routing changes do, in fact, give ASes much greater power. In addition, the convergence process—where BGP explores multiple options before settling on a new stable path—allows even more farflung ASes to get a (temporary) look at the client's traffic. BGP convergence, while notoriously slow, is probably fast enough to prevent these ASes from performing a successful traffic-analysis attack. Still, these ASes can learn about a client's use of the Tor network (and a particular guard)—information that can be combined with other data to implicate the client. As an example, the suspect in the recent bomb scandal at Harvard university was implicated purely due to the use of the Tor network from the Harvard campus [12]. While in this case, FBI had direct visibility into the suspect's communications, route convergence enables remote (off-path) attackers to draw similar inferences.

3.2 Manipulating Interdomain Routing

Internet routing is vulnerable to well-known attacks which enable a malicious router or AS to manipulate routing by advertising incorrect BGP control messages. For example, an AS could launch a prefix hijack attack [31] by advertising a particular IP prefix as its own, in which case, a certain fraction of Internet traffic destined to destinations in that prefix would be captured by the AS. AS-level adversaries can exploit these vulnerabilities in several ways to compromise user anonymity:

Traffic analysis via prefix hijack: In order to deanonymize the user associated with a target connection (say an observed connection to the WikiLeaks website), the adversary can first use existing attacks on Tor that can infer what guard relay the connection uses [19,24,25,27]. Next, the adversary can learn the identify of the client by launching a prefix-hijack attack against the prefix corresponding to the discovered guard relay. The attack allows a malicious AS to see the traffic destined to the guard relay. The attack essentially blackholes all traffic destined to the guard relay, so the client's connection only remains active for a limited amount of time, after which it will be dropped. For the duration of the connection, the malicious AS can learn the set of clients associated with the guard relay (anonymity set), by inspecting the IP headers. As we previously discussed in the Harvard example, this reduced anonymity set can already be incriminating for the user.

Nevertheless, prefix hijack only enables a limited form of traffic analysis (discussed above), since the connection is eventually dropped. We note that a malicious AS cannot perform a man in the middle attack pretending to be the guard since the Tor software is shipped with cryptographic keys of trusted directory authorities; these trusted authorities digitally sign the Tor network consensus containing cryptographic keys of guards

Traffic analysis via prefix interception: To perform exact deanonymization of the user via end-to-end traffic analysis, malicious ASes could launch a variant of the prefix hijacking attack, known as a prefix interception attack [11]. A prefix interception attack allows the malicious AS to become an intermediate AS in the path towards the guard relay, i.e., after interception, the traffic is routed back to the actual destination. Such an interception attack allows the connection to be kept alive, giving an additional opportunity for fine grained traffic analysis. For example, if the flow of traffic is from the user towards the destination website (say, a file upload to WikiLeaks), then the adversary can correlate users' traffic to the guard with the target flow at the destination, and fully deanonymize the user. In case the flow of traffic is towards the client (file download from WikiLeaks), then correlation can be performed using the asymmetric traffic analysis mechanism discussed next. The latter scenario is illustrated in Figure 1(c).

In addition to "classical" prefix hijacking and prefix interception, Renesys [34] recently shed light on a man-in-the-middle attack using BGP communities. BGP communities are labels that can be attached to BGP routes in order, most of the time, to change the default behavior of the BGP decision process. In particular, some ISPs use well-known labels in order to limit the propagation of the routes coming from their clients to a subset of their neighbors. Using communities, an attacker can therefore limit the propagation of a hijacked prefix to a few ASes, making the attack extremely hard to detect.

Our above attacks enable an AS-level adversary to deanonymize user identity corresponding to a monitored target connection. Similarly, ASes that act as the Tor client's own ISP already see the client's traffic to the guard, so they only need to intercept traffic from the exit relay to the destination. Furthermore, our attacks can be extended to perform general surveillance of the Tor network by intercepting traffic at *both* guard and exit relays. Since Tor clients select relays with a probability proportional to their bandwidth, high bandwidth relays observe a significant fraction of Tor traffic. Thus, an adversary could intercept traffic towards high bandwidth guard relays and exit relays (last hop), and perform traffic correlation to break user anonymity in Tor.

3.3 Asymmetric Traffic Analysis

Next, we present a novel traffic-analysis attack that AS-level adversaries can use to compromise user anonymity. In particular, our attack can be used in conjunction with the previously discussed interception attacks to increase adversaries' surveillance capabilities.

Let us suppose that a Web server is sending a large file to a client. Recall that conventional end-to-end timing analysis considers a scenario where a malicious AS observes traffic from the Web server to the last relay, as well as from the first guard relay to the client.[2] However, Internet paths are often asymmetric; thus, the path between the Web server and the last relay may differ from the path between the last relay and the server. This observation has interesting consequences for traffic analysis. Given the asymmetric nature of Internet paths, we can view the conventional end-to-end attack scenario as a setting in which the adversary is able to observe traffic at both ends of the anonymous path, and in *the same direction as the flow of traffic*.

In our new traffic-analysis attack, the adversary may observe traffic at both ends of the anonymity path, but *in opposite directions to each other*. For example, our attack is applicable to the scenario where an adversary observes traffic from the (a) last relay to the Web server, and first relay to the client (illustrated in Figure 1(b)), or (b) the Web server to the last relay and client to the first relay. Note that, in scenario (b), both paths have destinations that are Tor relay nodes, meaning that an adversary can easily attack a large number of users simply by launching an interception attack on destination prefixes that include Tor relay nodes.

We call such an attack an *asymmetric traffic analysis*. In this new setting, the adversary might not be able to observe data traffic at one end of the anonymity circuit, but it can still observe TCP acknowledgement traffic. In most deployed anonymity systems, SSL/TLS encryption is used, which is initialized at the session layer, and leaves the TCP header unencrypted. Our attack inspects TCP headers to infer the number of bytes being acknowledged using the TCP sequence number field. Our traffic-analysis attack considers the number of bytes seen in data packets at one end, the number of bytes acknowledged by TCP at the other end, and analyzes correlation between these fields over time. Note that a new correlation analysis is required here since TCP acknowledgements are cumulative, and there is not a one-to-one correspondence between packets seen at both ends of the communication.

In a more extreme variant of the attack, an adversary observes only the acknowledgment traffic at both ends of the connection. In this case, our attack correlates the number of acknowledged bytes at both ends of the path over time. Our preliminary evaluation in Section 4 shows the feasibility of such asymmetric traffic analysis.

4. PRELIMINARY RESULTS

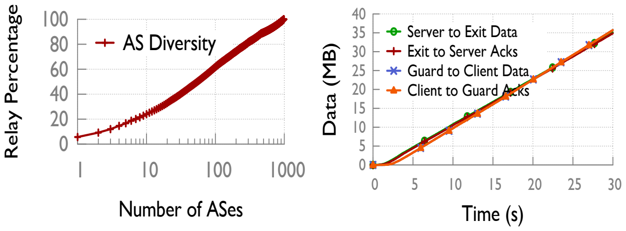

Figure 2: AS-level Traffic Analysis: Tor relays are concentrated in a handful of ASes, presenting attractive targets for routing attack. Just 3 ASes host 10% of Tor relays, while 7 ASes host 20% of Tor relays (left). ASes can deanonymize clients by observing traffic at both communication ends in *any* direction. The data sent from server to exit is nearly identical to the amount of data acknowledged by the client to the guard across time (right).

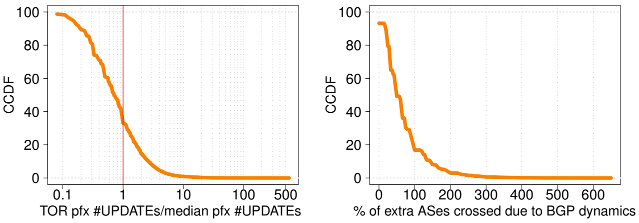

t

Figure 3: Tor prefixes seeing at least one UPDATE over the month (99% of them) often end up being less stable than BGP prefixes, with 40% of them receiving more UPDATEs than the median amount (left). This relative instability increases surveillance capability for these prefixes as 50% of the paths changes cause at least 50% more ASes (with respect to the shortest AS-PATH) to see Tor traffic for more than 5 minutes (right).

[2] Alternatively, if the traffic is flowing from the client to the server, then end-to-end timing analysis considers a scenario where the adversary observes traffic from the client to the first relay and from the last relay to the server.

In this section, we show that BGP temporal dynamics significantly increases AS-level surveillance capabilities and should therefore be considered during Tor relay selection. We also show the feasibility of asymmetric traffic analysis.

Methodology and datasets. We collected all the BGP updates received by 4 RIPE collectors over more than 100 eBGP sessions between May 1 and May 31, 2014. To ensure meaningful results, we removed any artificial updates caused by BGP session resets using the method developed in [30]. We also collected information (IP address, flags and bandwidth) about 4586 Tor relays listed by the Tor project [1]. 1918 of these relays were listed as "guards". For each Tor relay, we identified the most specific BGP prefix that contains it and ended up with 2361 BGP prefixes, announced by 985 distinct ASes. We refer to those prefixes as *Tor prefixes*. The distribution of the number of relays per prefix is skewed, with a median number of relay per prefix of 1, a 75-th percentile of 2, and maximum of 55 relays per prefix.

Tor relays are concentrated in a handful of ASes. Figure 2(left) depicts the lack of AS diversity among Tor relays. A point (x, y) on the curve means that x number of ASes are hosting $y\%$ of Tor relays. We found that all Tor relays are found in 2361 BGP prefixes announced by only 985 distinct ASes. Alarmingly, 3 of these ASes host 10% of Tor relays (Hetzner Online AG, OVH SAS, Abovenet Communications), while 7 ASes host 20% of Tor relays. Thus, we can see that a few ASes have significant visibility into Tor communications; these ASes also present an attractive target for active BGP attacks (by other AS-level adversaries).

Tor prefixes seeing UPDATEs over the month (99% of them) tend to be relatively unstable. We started by computing the relative chattiness of Tor prefixes with respect to the rest of the BGP prefixes. To do so, we computed the number of BGP UPDATEs received by each prefix on each session. Figure 3 (left) plots the number of UPDATEs seen by Tor prefixes on a session divided by the median number of monthly UPDATEs seen by a BGP prefix on the same session, excluding Tor prefixes that did not receive any UPDATE over the month (less than 1% of them). Results are presented as Complementary Cumulative Distribution Functions (CCDFs). A point (x, y) of the curve means that $y\%$ of Tor prefixes are seeing more than $x\%$ times the median amount of UPDATEs on the corresponding session.

Close to 40% of the Tor prefixes seeing at least one UPDATE also saw more UPDATEs than a normal BGP prefix since the ratio was strictly greater than one. Also, 5% of Tor prefixes received four times more UPDATEs than the median case and one Tor prefix received close to 600 times more UPDATEs than the median case. We also observed that unstable Tor prefixes (ratio greater than 10) were unstable on multiple sessions (on average, 6) which tends to indicate that the root cause of the UPDATE is located relatively close to the originator AS.

BGP temporal dynamics significantly increase AS-level surveillance capabilities. As a second step, we computed how many more ASes were seeing traffic directed to a Tor prefix as a result of BGP temporal dynamics. As baseline, we considered the case of a static BGP in which only the ASes belonging to the shortest AS-PATH are seeing traffic. Also, to be fair, we did not consider an AS if it was crossed for less than 5 minutes as it is anyway unlikely that an attack can be performed on such a short timescale. The right part of Figure 3 describes the results as a CCDF. In 50% of the cases, the number of ASes seeing Tor traffic increases by more than 50% with respect to a static BGP. In close to 20% of the cases, the number of ASes seeing Tor traffic doubles. These results tend to indicate that the AS-level surveillance capabilities is actually much worse that what has been shown in previous studies [15, 21].

Asymmetric traffic analysis is feasible. We performed a wide-area experiment over the live Tor network using a Tor client and an Apache webserver. The Tor client and server were chosen in different geographical locations. We used the torsocks program at the client to tunnel wget requests over Tor, and downloaded a large file from the webserver. We used tcpdump to collect data at the server and the client, and show the number of MBs sent or acknowledged (computed by inspecting TCP headers) at various segments of the path in Figure 2 (right). We can see that data sent or acknowledged at all 4 segments is nearly identical across time. Thus, it suffices for an AS-level adversary to observe traffic at both ends of the communication in *any* direction.

5. COUNTERMEASURES

Limiting impact of BGP dynamics: To minimize opportunities for AS-level traffic analysis, the Tor network can monitor the path dynamics between the clients and the guard relays, and between the exit relays and the destinations. Information about path dynamics can be obtained using data-plane (e.g., traceroute) or control-plane (e.g., BGP feed) tools. For instance, each relay could publish the list of any ASes it used to reach each destination prefix in the last month. This information can be distributed to all Tor clients as part of the Tor network consensus data. Tor clients can use this data in relay selection, perhaps in combination with their own traceroute measurements of the forward path to each guard relay. For example, Tor clients should select relays such that the same AS does not appear in both the first and the last segments, after *taking path dynamics into account*.

Detecting and reacting to routing manipulations: We can extend the data-plane and control-plane based monitoring framework discussed above, to perform *real-time* monitoring of prefixes corresponding to the Tor relays. The monitoring system can leverage state-of-art techniques for detecting prefix hijacks and interception attacks [11,22,28,31–33]. For anonymity systems, false positives are much more acceptable than false negatives, so we can afford to be aggressive in classifying anomalies as attacks, rather than risk compromising user anonymity. So, if the monitoring system has a suspicion that a relay might be under attack (or even that a relay is now using a path containing a new AS), this information can be broadcast through the Tor network in real time, so clients can avoid selecting this relay.

Favoring relays with shorter AS-PATHs: Monitoring is particularly effective at detecting subprefix attacks, where an adversary advertises a more-specific prefix for the victim relay, as all ASes would eventually see the bogus announcement. However, the adversary could use stealthier attacks, such as advertising an existing prefix or using the BGP community attack [34]. These attacks tend to affect only ASes that have relatively long paths to the legitimate destination AS, since other ASes will tend to favor the (shorter) route to the real destination. Thus, Tor clients can mitigate such routing manipulations by preferring guard relays with shorter AS-PATHs. Still, the client should balance this strategy with the need to limit the number of guard relays, to protect against conventional attacks on long-term anonymity.

Mitigating asymmetric traffic analysis: Using IP-layer encryption (e.g., IPsec) rather than SSL/TLS would thwart our asymmetric traffic-analysis attack, by hiding the TCP sequence numbers from the adversarial ASes. However, using IPsec would come at a significant cost; because IPsec is not widely used, it makes Tor traffic much easier to identify, and limits its applicability for important applications such as censorship-resilient communications.

6. RELATED WORK

AS-level adversaries: Most security analysis of anonymity systems focuses on the threat of end-to-end timing analysis by malicious or compromised relays/proxies. However, a similar timing analysis attack can also be performed by AS-level adversaries, though existing literature on this front is more limited. Feamster and Dingledine [17], and later Edman and Syverson [15] explored this aspect, and considered the probability of a single AS being on the path between a client and the first relay as well as on the path between the last relay and the destination, using the AS-level path simulator of Gao et al [18]. Recently, Johnson et al. [21] analyzed the impact of such attacks using user-understandable metrics for anonymity, and Akhoondi et al. [8] considered path selection algorithms that minimize opportunities for AS-level end-to-end traffic analysis. Finally, Murdoch et al. [26] considered the analogous analysis with respect to Internet exchange level adversaries, which are also in a position to observe significant fraction of Internet traffic. We build upon prior work in this domain, and show increased surveillance capabilities of AS-level adversaries, due to BGP path changes, active routing attacks, and asymmetric traffic analysis.

Tor Traffic Analysis: There is an exciting thread of research that aims to investigate the traffic analysis attacks on anonymity systems such as Tor. For example, Murdoch and Danezis [25] showed how a remote adversary could congest a Tor relay by sending traffic, and observe the impact of the impact of congestion on other flows to infer the relay's membership in an anonymity channel (also known as a circuit). Evans et al. [16] and Jansen et al. [20] show how an adversary could exploit protocol level details to cause similar relay congestion (and even shutdown) with minimal resources. The work of Mittal et al. [24] and Hopper et al. [19] further studies the impact of leveraging network level characteristics such as circuit throughput and latency to make probabilistic inferences about Tor relays and clients that are part of a target anonymity circuit. Most of the above attacks are only able to provide probabilistic information about Tor relays, and do not fully de-anonymize the actual clients. In contrast, we show that a remote adversary can fully de-anonymize Tor clients by actively manipulating inter-domain routing.

Other work: There has been a lot of work on prefix hijack attacks [28,31–33], and prefix interception attacks [11], but we are the first to analyze the implications of these attacks on privacy technologies such as anonymous communication. Recent work by Arnbak and Goldberg [9] discusses surveillance possibilities by AS-level adversaries from a legal perspective, but does not focus on anonymity systems.

7. CONCLUSION

The security of privacy technologies like Tor depends on how the underlying Internet infrastructure delivers traffic. In this paper, we show that normal BGP routing changes greatly increase the likelihood that an AS (or set of colluding ASes) can perform traffic-analysis attacks, and ASes can easily manipulate BGP to gain even wider visibility into user traffic. In fact, the adversary need only lie on one direction of each path, between the client and guard and between the last relay and the server. Our initial experiments illustrate that these vulnerabilities can be easily exploited in practice.

Improvements in BGP security can go a long way toward addressing the most serious concerns. However, deployment of BGP security solutions—and particularly techniques that prevent interception attacks—has proven challenging. We hope that the concerns we raise about compromises of user anonymity help build much-needed momentum for improving BGP security in the long term, and real-time detection of BGP anomalies in the short term. In our future work, we plan to conduct a more extensive measurement study, including an analysis of recent BGP interception attacks to quantify the potential risks to the anonymity of Tor users.

8. REFERENCES

[1] Tor metrics portal. https://metrics.torproject.org. Accessed, July 2014.

[2] Who uses Tor? https://www.torproject.org/about/torusers.html.en. Accessed, July 2014.

[3] Chinas 18-minute mystery, Nov. 2010. Renesys Blog Posting, http://www.renesys.com/2010/11/chinas-18-minute-mystery/.

[4] How the NSA attacks Tor/Firefox users with QUANTUM and FOXACID. https://www.schneier.com/blog/archives/2013/10/how_the_nsa_att.html, Oct. 2013.

[5] The new threat: Targeted Internet traffic misdirection, Nov. 2013. Renesys Blog Posting, `http://www.renesys.com/2013/11/mitm-internet-hijacking/`.

[6] Peeling back the layers of Tor with EgotisticalGiraffe. `http://www.theguardian.com/world/interactive/2013/oct/04/egotistical-giraffe-nsa-tor-document`, Oct. 2013.

[7] Tor stinks. `http://www.theguardian.com/world/interactive/2013/oct/04/tor-stinks-nsa-presentation-document`, Oct. 2013.

[8] M. Akhoondi, C. Yu, and H. V. Madhyastha. Lastor: A low-latency AS-aware Tor client. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy*, SP '12, pages 476–490, Washington, DC, USA, 2012. IEEE Computer Society.

[9] A. Arnbak and S. Goldberg. Loopholes for circumventing the constitution: Warrantless bulk surveillance on americans by collecting network traffic abroad. In *HotPETs*, 2014. Available at http://ssrn.com/abstract=2460462.

[10] J. Ball. NSA stores metadata of millions of web users for up to a year, secret files show. `http://www.theguardian.com/world/2013/sep/30/nsa-americans-metadata-year-documents`, Sep. 2013.

[11] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the Internet. In *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, SIGCOMM '07, pages 265–276, New York, NY, USA, 2007. ACM.

[12] R. Brandom. FBI agents tracked harvard bomb threats despite Tor. http://www.theverge.com/2013/12/18/5224130/fbi-agents-tracked-harvard-bomb-threats-across-tor. Accessed, July 2014.

[13] R. Dingledine, N. Hopper, G. Kadianakis, and N. Mathewson. One fast guard for life (or 9 months). In *HotPETs*, 2014.

[14] R. Dingledine, N. Mathewson, and P. 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.

[15] M. Edman and P. Syverson. AS-awareness in Tor path selection. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, CCS '09, pages 380–389, New York, NY, USA, 2009. ACM.

[16] N. S. Evans, R. Dingledine, and C. Grothoff. A practical congestion attack on Tor using long paths. In *Proceedings of the 18th Conference on USENIX Security Symposium*, SSYM'09, pages 33–50, Berkeley, CA, USA, 2009. USENIX Association.

[17] N. Feamster and R. Dingledine. Location diversity in anonymity networks. In *Proceedings of the 2004 ACM Workshop on Privacy in the Electronic Society*, WPES '04, pages 66–76, New York, NY, USA, 2004. ACM.

[18] L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Trans. Netw.*, 9(6):733–745, Dec. 2001.

[19] N. Hopper, E. Y. Vasserman, and E. Chan-TIN. How much anonymity does network latency leak? *ACM Trans. Inf. Syst. Secur.*, 13(2):13:1–13:28, Mar. 2010.

[20] R. Jansen, F. Tschorsch, A. Johnson, and B. Scheuermann. The sniper attack: Anonymously deanonymizing and disabling the Tor network. In *Proceedings of the 21st Annual Network and Distributed System Security Symposium (NDSS '14)*. Internet Society, 2014.

[21] A. Johnson, C. Wacek, R. Jansen, M. Sherr, and P. Syverson. Users get routed: Traffic correlation on Tor by realistic adversaries. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security*, CCS '13, pages 337–348, New York, NY, USA, 2013. ACM.

[22] J. Karlin, S. Forrest, and J. Rexford. Autonomous security for autonomous systems. *Computer Networks*, Oct. 2008.

[23] E. MacAskill, J. Borger, N. Hopkins, N. Davies, and J. Ball. GCHQ taps fibre-optic cables for secret access to world's communications. `http://www.theguardian.com/uk/2013/jun/21/gchq-cables-secret-world-communications-nsa`, June 2013.

[24] P. Mittal, A. Khurshid, J. Juen, M. Caesar, and N. Borisov. Stealthy traffic analysis of low-latency anonymous communication using throughput fingerprinting. In *Proceedings of the 18th ACM Conference on Computer and Communications Security*, CCS '11, pages 215–226, New York, NY, USA, 2011. ACM.

[25] S. J. Murdoch and G. Danezis. Low-cost traffic analysis of Tor. In *Proceedings of the 2005 IEEE Symposium on Security and Privacy*, SP '05, pages 183–195, Washington, DC, USA, 2005. IEEE Computer Society.

[26] S. J. Murdoch and P. Zieliński. Sampled traffic analysis by Internet-exchange-level adversaries. In *Proceedings of the 7th International Conference on Privacy Enhancing Technologies*, PET'07, pages 167–183, Berlin, Heidelberg, 2007. Springer-Verlag.

[27] L. Overlier and P. Syverson. Locating hidden servers. In *Security and Privacy, 2006 IEEE Symposium on*, pages 15 pp.–114, May 2006.

[28] X. Shi, Y. Xiang, Z. Wang, X. Yin, and J. Wu. Detecting prefix hijackings in the Internet with Argus. In *Proceedings of the 2012 ACM Conference on Internet Measurement Conference*, IMC '12, pages 15–28, New York, NY, USA, 2012. ACM.

[29] M. Wright, M. Adler, B. N. Levine, and C. Shields. Defending anonymous communications against passive logging attacks. In *Proceedings of the 2003 IEEE Symposium on Security and Privacy*, SP '03, pages 28–, Washington, DC, USA, 2003. IEEE Computer Society.

[30] B. Zhang, V. Kambhampati, M. Lad, D. Massey, and L. Zhang. Identifying BGP routing table transfer. *SIGCOMM 2005 MineNet Workshop*, August 2005.

[31] Z. Zhang, Y. Zhang, Y. C. Hu, and Z. M. Mao. Practical defenses against BGP prefix hijacking. In *Proceedings of the 2007 ACM CoNEXT Conference*, CoNEXT '07, pages 3:1–3:12, New York, NY, USA, 2007. ACM.

[32] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush. iSPY: Detecting IP prefix hijacking on my own. *IEEE/ACM Trans. Netw.*, 18(6):1815–1828, Dec. 2010.

[33] C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis. A light-weight distributed scheme for detecting IP prefix hijacks in real-time. In *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, SIGCOMM '07, pages 277–288, New York, NY, USA, 2007. ACM.

[34] E. Zmijewski. The end of undetected BGP route hijacking. http://www.renesys.com/wp-content/uploads/2014/05/Linx851.pdf, May 2014.