



Actively Probing Routes for Tor AS-Level Adversaries with RIPE Atlas

Wilfried Mayer, Georg Merzdovnik, Edgar Weippl

► To cite this version:

Wilfried Mayer, Georg Merzdovnik, Edgar Weippl. Actively Probing Routes for Tor AS-Level Adversaries with RIPE Atlas. 35th IFIP International Conference on ICT Systems Security and Privacy Protection (SEC), Sep 2020, Maribor, Slovenia. pp.234-247, 10.1007/978-3-030-58201-2_16 . hal-03440850

HAL Id: hal-03440850

<https://inria.hal.science/hal-03440850>

Submitted on 22 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Actively Probing Routes for Tor AS-level Adversaries with RIPE Atlas

Wilfried Mayer¹, Georg Merzdovnik¹, and Edgar Weippl²

¹ SBA Research, Vienna, Austria
{wmayer,gmerzdovnik}@sba-research.org

² University of Vienna, Austria
edgar.weippl@univie.ac.at

Abstract. Tor provides anonymity to millions of users around the globe, which has made it a valuable target for malicious actors. As a low-latency anonymity system, it is vulnerable to traffic correlation attacks from strong passive adversaries, such as large autonomous systems. Estimations of the risk posed by such attackers as well as the evaluation of defense strategies are mostly based on simulations and data retrieved from BGP updates. However, this might only provide an incomplete view of the network and thereby influence the results of such analyses. It has already been acknowledged in previous studies that direct path measurements, e.g. with traceroute, could provide valuable information. But in the past, such measurements were thought to be impossible, because they require the placement of measurement nodes in the same ASes as the respective Tor network nodes. With the rise of new technologies and methodologies, this assumption needs to be re-evaluated.

In this paper we present a novel methodology to utilize the RIPE Atlas framework, a network of more than 10,000 probes worldwide, to actively perform traceroute commands from and to Tor guard and exit relays to clients and destinations. Based on multiple global scans our results validate previous results and show the large influence on Tor posed by a limited set of ASes. These are in a strong position to carry out effective correlation attacks on Tor traffic. With this work, we provide an additional source of information that can be used together with BGP route information to increase the accuracy of future models and simulations of Tor and ultimately improve anonymity on the Internet.

Keywords: Tor · RIPE Atlas · Traceroute Measurements.

1 Introduction

Tor is the most notable anonymity network, used by 2 to 3 million people on a daily basis and advertising up to 400 Gbit/s of bandwidth by utilizing around 6,500 voluntarily operated Tor relays. It provides anonymity by routing traffic via three different Tor nodes. As a low latency network, due to its design, it is not capable of guaranteeing anonymity in the case of a global passive observer. This form of attacker is explicitly excluded from the threat model, assuming that such

a global passive observer does not exist. Although not global, powerful observers exist, potentially threatening the anonymity of Tor users. However, their capabilities are not exactly clear. One reason for this is the theoretical assumption, that the underlying Internet hierarchy is flat and evenly distributed. Trivially, this is not the case, as the Internet is shaped in different tiers and various entities with different levels of control, e.g., Internet Exchange Points (IXP) with a high level of control, and small ISPs with a low level of control. Also, the Tor network does not utilize the Internet in an evenly distributed manner, as the location of Tor relays is depending on various parameters, e.g., economical (the price of bandwidth) or political (censorship, prosecution) reasons. Prior work [9, 11, 19] has shown that traffic through the Tor network only takes a limited set of routes on the Internet, making the threat of a powerful passive observer far more likely. They point out that few AS-level entities provide a high proportion of the Tor bandwidth, thus making them powerful entities for traffic correlation attacks. These studies rely on BGP updates and a prediction of routes taken. While they also describe that a traceroute-based approach could potentially yield better results, they also argue that it would not be feasible for measuring AS-level adversaries in the Tor network, because of the need to have measurement nodes placed on the same ASes as Tor nodes and destinations. However, with the introduction of the RIPE Atlas network [21] this assumption can no longer be taken for granted. The RIPE Atlas network is a global measurement network, which can be used by researchers to measure Internet connectivity and reachability. It has already been used for several studies [4] concerning network routing [10] as well as censorship measurements [3].

Our work presents a novel method of measuring the routes that traffic takes from and to the Tor network by utilizing active network probing, in contrast to estimations via BGP updates. We do this by utilizing probes that are placed in autonomous systems (AS) also in use by Tor relays, Tor users or Tor connection recipients. For this purpose, we utilize the RIPE Atlas network, which consists of more than 10,000 globally distributed probes connected to many different autonomous systems. To measure the routes a packet takes from and to the Tor network, we execute traceroute commands on these probes and collect information on the ASes observed on the respective paths. With this method, we gather data to create better predictions of powerful adversaries existing on the Internet and thus to improve the anonymity of Tor users. More specifically, the contributions of this paper are as follows:

Active Measurements of AS Interconnections: Using *traceroute* based measurements we estimate the capabilities of AS-level adversaries and show the influence of only a few ASes on a large amount of traffic.

Open-Source Active Measurement Tool: To improve the evaluation of future attacks and defenses against Tor, we provide an open-source framework to perform active measurement to acquire routing information for Tor nodes.

2 Related Work

Tor, described originally by Dingledine et al. in 2004 [8], grew to the most important anonymity system online nowadays. As a low-latency overlay network, it is inherently vulnerable to passive attacks by global observers, which is already described in the original specification. Instead, they work with a threat model that includes attackers that can only observe fractions of the network traffic.

Feamster and Dingledine [11] provided the first analysis of location diversity in the Tor network for independently operated autonomous systems based on BGP routing tables. They analyzed the probability of an entry path to the network and an exit path from the network will cross through the same AS. Their analysis shows that previous methods of choosing paths/nodes based on IP prefixes are not sufficient to guarantee a diverse set of ASes, since in about 10% to 30% of the time both the entry and exit path to the mix network will cross through the same AS. A refinement of this approach by Edman and Syverson in 2009 [9] shows that the previous study even underestimated the potential threat. A study of Tor security properties against traffic correlation attacks was presented by Johnson et al. [16]. Their results show that, depending on location, a user’s chance of compromise can be at 95% within 3 months of monitoring against a single AS. One mitigation they propose is to carefully select which entry and exit nodes to use. Wacek et al. [25] built a graph of the Tor network to capture the networks AS boundaries. Using this graph they provide an evaluation of a set of proposed relay selection methods and quantify their respective anonymity properties. Their results show that bandwidth is an important property for the performance of such algorithms, and should not be neglected.

The importance of location diversity in the Tor network has been shown by several attacks proposed in recent years. Vanbever et al. [24] provide a study of the capabilities of AS level adversaries. Sun et al. [23] describe a set of advanced routing attacks on Tor, named *Raptor*. They also describe the feasibility of asymmetric AS-level attacks by observing not only data traffic from exit relay to the server but also TCP acknowledgment traffic on other routes which increases the capabilities of AS-level adversaries. In 2016, Nithyanand et al. [19] also use data on the Internet’s topology [13] in a combination with AS-topology simulations [12] to estimate the threat posed by adversaries to Tor users. While previous attempts at the correlation of traffic [15, 17] had very limited performance or required a large amount of captured traffic or time, *DeepCorr* [18], developed by Nasr et al. greatly improves the feasibility of such attacks. By leveraging emerging learning mechanisms they manage to achieve drastically higher performance compared to existing state-of-the-art systems.

To mitigate the threat posed by an AS to be able to monitor Tor users, various kinds of protection mechanisms have been proposed [2]. Nithyanand et al. proposed *Astoria* [19], an AS-aware Tor client. While similar in functionality to *LASTor* [1], it provides improved protection with concern to threat models and attacker capabilities. Sun et al. [22] presented a measurement study on the security of Tor against BGP hijacking attacks and presented a new relay selection mechanism to mitigate such attacks on Tor. In contrast to previous approaches

DeNASA from Barton et al. [5] provides a mechanism for AS-aware path selection independently of the destination. Additionally, they propose another system for the creation of efficient and anonymous Tor circuits [6]. Hanley et al. [14] proposed an extension to the work presented by Sun et al. [22] to increase the provided privacy and anonymity guarantees. Wan et al. [26] showed that several attacks against a set of the proposed protections are still possible, but they also proposed simple solutions, which allow mitigating the threat posed by their developed methods.

3 Active Acquisition of Routing Information

In the following section, we describe a novel method to measure strong AS-level observers, which are in a good position to conduct correlation attacks. As an overlay network, Tor depends on the underlying structure of the Internet. While often a flat hierarchy is assumed, it is clear that this is not the case. We can model the structure of the Internet by looking at autonomous systems identified by a unique AS number (ASN). One AS can be seen as an administrative entity that is responsible for a defined routing policy. Some AS are large and include a lot of Tor users, destinations or relays, others do not contain users and destinations but are used for routing Tor traffic through the Internet and others are not important for Tor routing at all. Thus, some entities can observe more traffic than others. With our measurements, we find a way to quantify which entities are in a stronger position. Figure 1 illustrates the basic idea of a standard traffic correlation attack, where one adversary (AS2) is placed on the incoming route to Tor as well as on the outgoing route to the destination. Sun et al. [23] showed that it is also possible to correlate reverse-path traffic. Other work already quantified strong adversaries with the help of BGP route updates. In contrast, we develop a method that utilizes the RIPE Atlas framework to actively acquire routing information.

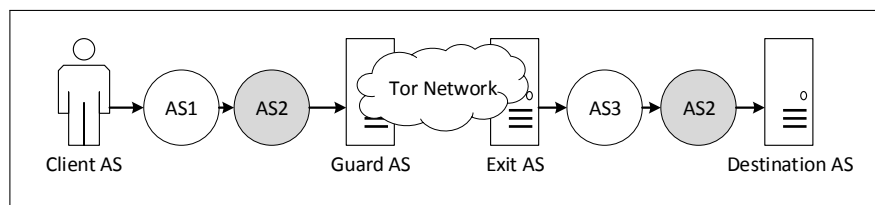


Fig. 1. AS2 in a possible position for a traffic correlation attack

3.1 Relay AS Diversity

As shown in Table 1 the Tor network currently consists of 6,509 relays (January 5th, 2020). Only relays with the *Guard* flag (stable and reliable relays after a

ramp-up phase [7]), are used as entry relay. Only relays configured to allow exiting traffic are potential exit relays in a Tor circuit. Because of the more stringent requirements, the number of guard and exit relays (with guard/exit probability > 0) is a lot smaller than 6,509. This also affects the AS diversity, which is the number of different ASes these relays are placed in. Current numbers are shown in Table 1. Tor relays are chosen based on their flags and consensus weight. In Figure 3 we show the AS diversity relation to guard and exit probability. We see that a small number of AS has a large share of exit (a) and guard (b) probability. Eight ASes have more than 50% exit probability and 48 ASes have more than 90%. We also see that only four AS have more than 50% guard probability and 122 have more than 90%. So although all Tor relays are distributed over more than 1,100 ASes, the majority of entry and exit routing endpoints are placed in a few ASes.

Table 1. Tor Relay overview

	Relays	Diff. AS	BW (Gbit/s)
All Relays	6,509	1,104	418.07
Exit Relays	1,000	275	112.90
Guard Relays	2,415	470	254.61

3.2 The RIPE Atlas Framework

The RIPE Atlas framework is a highly distributed measurement network consisting of more than 10,000 available probes, deployed in over 3,500 different ASes. It allows us to execute various low-level commands, e.g., ping or traceroute, on these probes and further process the results. We will utilize this to execute traceroute commands from RIPE Atlas probes that are deployed in the same ASes as Tor guard or exit relays as well as clients and popular destinations. Figure 2 illustrates the global distribution of RIPE Atlas probes and the global distribution of Tor relays. We see that countries with a higher number of Tor relays also run more RIPE Atlas probes.

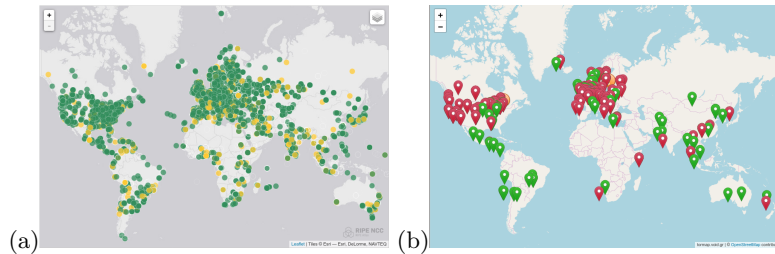


Fig. 2. (a) Worldwide RIPE Atlas Coverage³(b) Visualization of Tor Relays⁴

Figure 3 also shows the cumulated guard and exit probability for autonomous systems that contain RIPE Atlas probes. From 275 ASes that contain exit relays, only 112 also contain a probe (419 relays out of 1,000). Still, that makes approx. 41% of the total exit probability (35% with only 17 ASes). This differs from the cumulated guard probability. From 470 ASes that contain 2,415 relays, 238 ASes (with 1,848 relays) also include a RIPE Atlas probe, which represent guard relays with a sum of 83% guard probability (80% with 98 ASes). Especially for exit relays, these numbers could be drastically increased if only a few, exit-focused ASes would also host RIPE Atlas probes. Table 2 identifies ASes, that are currently not hosting any RIPE probes. By adding only 5 probes we could measure ASes with 76% exit probability in total and 10 probes would gain up to 87% probability in total.

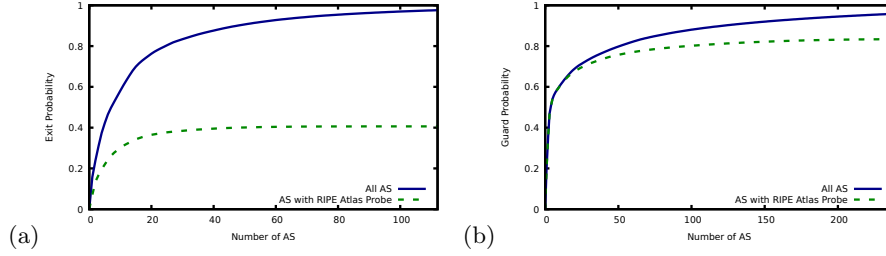


Fig. 3. Accumulated percentage of (a) exit, and (b) guard probability with the number of autonomous systems

Table 2. AS with Tor relays currently not hosting a RIPE Atlas probe

AS Name	Relays	Gbit/s BW	P_{exit}	P_{guard}
200052 FERAL	54	17.01	.158	.004
208323 APPLIEDPRIVACY	16	7.28	.082	.001
53667 FRANTECH	94	8.78	.048	.011
8972 HOSTEUROPE	23	2.60	.000	.010
63949 LINODE-AP	162	3.71	.001	.008

3.3 Active *traceroute* Probing with RIPE Atlas

As illustrated in Figure 4 we perform *traceroute* measurements to identify routes taken for four different directions: (1) all client ASes to all guard ASes, (2) exit ASes with probes installed to the destination ASes, (3) destination ASes to all exit ASes, and (4) guard ASes with probes installed to the client ASes. With

³ <https://atlas.ripe.net/>

⁴ <https://tormap.void.gr/>

these measurements, we do not cover all possible routes since not all ASes have probes installed. In the different directions we measure (1) 1s00% (2) ~40% (3) 100% (4) ~83% in terms of route probability.

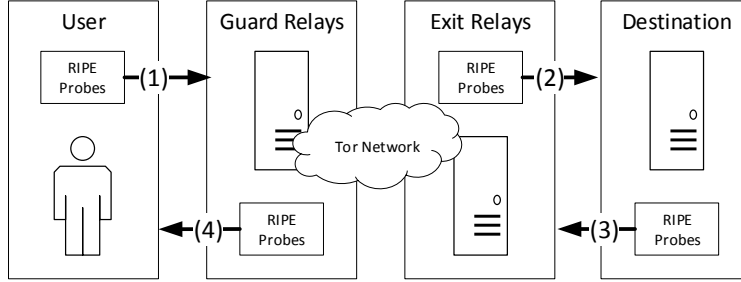


Fig. 4. Four different directions of active RIPE Atlas *traceroute* scans

In detail, this process works as follows:

1. Create the following sets:
 - i. AS_{client} ... ASes of the clients
 - ii. AS_{guard} ... all ASes with guard relays
 - iii. $AS_{guard+probe}$... all ASes with guard relays and RIPE atlas probes
 - iv. AS_{exit} ... all ASes with exit relays
 - v. $AS_{exit+probe}$... all ASes with exit relays and RIPE atlas probes
 - vi. $AS_{destination}$... ASes of the destinations
2. Generate ICMP traceroute measurement definitions for the following directions:

(1) AS_{client}	$\xrightarrow{\text{traceroute}}$	AS_{guard}
(2) $AS_{exit+probe}$	$\xrightarrow{\text{traceroute}}$	$AS_{destination}$
(3) $AS_{destination}$	$\xrightarrow{\text{traceroute}}$	AS_{exit}
(4) $AS_{guard+probe}$	$\xrightarrow{\text{traceroute}}$	AS_{client}
3. Execute the *traceroute* with the RIPE Atlas measurement API. ("protocol": "ICMP", "response_timeout": 20000, "packets": 1). Every RIPE Atlas measurement is charged with credits, obtained by hosting RIPE Atlas probes. For the current deployment that estimates to $20 \cdot 1230 = 24600$ credits for one client and one destination.
4. Process all results and look up the corresponding AS from the *ip2asn* database.
5. For every *traceroute*, mark all included ASes with the probability of that path being chosen, i.e., the corresponding guard/ exit probability.
6. Combine the values for the directions 1 and 4 for the entry side, and 2 and 3 for the exit side, s.t., if an AS appears on either the forward or the reverse path it is assigned with the probability of that path being chosen. For multiple destinations, all traceroutes are combined.
7. Point out the top ASes, that appear on entry and exit side by looking at $P_{guard} \cap P_{exit}$.

3.4 Origin and Destination AS

The sets of guard and exit relays can be derived by combining the Tor consensus with the RIPE Atlas probe overview. However, a client set and a destination set has to be chosen to conduct a measurement. A single client and a single destination are easily scannable, but it doesn't give us a full picture. However, executing traceroutes for all possible client and destination ASes is not feasible. Thus, we have to choose client and destination AS sets for our measurements. In 2008, Edman and Syverson [9] captured traffic from Tor relays to determine top autonomous systems. We are choosing a different approach using popular destinations and large client ASes. For the client set, we choose different countries and pick the 10 ASes containing the most RIPE Atlas probes. Then, we pick one probe thereof. E.g., Germany has 1,485 probes installed, in 343 different ASes, and we pick the ASes with most probes installed⁵. For the US we do the same⁶. For the most common destinations, we derive a list of top destinations from the Tranco [20] top sites list⁷. We take the 100 most popular domains, resolve the domain, and match the corresponding ASes. From the 100 top sites in 44 different ASes, ten ASes also have a RIPE Atlas probe installed⁸ and will be used as our destination ASes.

3.5 Data Sources

To facilitate reproducibility and encourage openness, all used data files are publicly available at the project website⁹. In particular, our work relies on following data sources:

1. The Tor consensus that contains all Tor relays with their IP address, associated flags (particularly "Guard" and "Exit"), advertised bandwidth and guard and exit probability. We collect this information via the Tor network status protocol *onionoo*¹⁰.
2. Statistical data about the *RIPE Atlas probes*¹¹. We use different data (e.g., id, number and AS of the probes) to find all probes connected to the same ASes as guard and exit relays.
3. Freely accessible *ip2asn*¹² databases to match IP addresses with the corresponding AS number.
4. Active RIPE Atlas *traceroute* results¹³. The measurements used for this paper are accessible at the projects website.

⁵ Client ASes Germany: 3320, 6830, 31334, 8881, 3209, 6805, 553, 680, 8422, 9145

⁶ Client ASes USA: 7922, 701, 7018, 209, 20115, 22773, 5650, 20001, 10796, 11427

⁷ Available at <https://tranco-list.eu/list/YL6G>

⁸ Destination ASes: 3, 15169, 4837, 24940, 36351, 14618, 16509, 14907, 3356, 794

⁹ Project website: <https://github.com/sbaresearch/ripe-tor>

¹⁰ onionoo: <https://metrics.torproject.org/onionoo.html>

¹¹ probes: <https://atlas.ripe.net/probes/>

¹² ip2asn: <https://iptoasn.com/>

¹³ measurements: <https://atlas.ripe.net/measurements>

4 Evaluation

In the following section, we evaluate our traceroute scans and show results. We start with an evaluation of a basic scan. Then, we present a larger measurement with multiple clients and destinations. We assess both directions on the guard and the exit side separately and also look at the combined results.

4.1 Measurement with a Single Client and a Single Destination

As an illustration of the capabilities of our methodology, we evaluate the results of measurements with one fixed client AS and one fixed destination AS. Therefore, we choose the AS of our research center as $AS_{client} = \{AS1764\}$, and the AS of one mirror of the `torproject.org` website as $AS_{destination} = \{AS24940\}$. We choose RIPE Atlas probes deployed in these ASes (id: 26895, 50609). We then execute 1,240 traceroute commands as defined in Section 3.3. Thereof, 269 only contain the client and destination AS, while 971 contain additional ASes on the path. In Table 3 we show various results. As expected, the client and destination AS (Hetzner, Nextlayer) are found on all traceroutes. ASes with a high guard or exit probability (Feral, Applied Privacy, OVH) also have a great share, although they are not intermediary and only found on the single traceroute to/from their AS. Large transit ASes, that appear on many routes are more interesting. In our measurement, we identified AS6939, AS47147, AS1200, and AS174 to be in a powerful position, as they appear on many routes and gain probability of up to 18%. Table 4 shows that for this single measurement only few ASes have a probability higher than 1% to appear on both sides.

Table 3. Results for a single client and single destination

AS Name	Dir.	P	P_{relays}	P_{routes}	Routes
24940 HETZNER-AS	exit	.988	.004	.984	269
200052 FERAL	exit	.161	.161	-	1
6939 HURRICANE	exit	.158	.001	.157	20
47147 AS-ANX	exit	.116	-	.116	4
1200 AMS-IX1	exit	.068	-	.068	17
1764 NEXTLAYER	guard	.992	-	.992	454
24940 HETZNER-AS	guard	.202	.202	-	1
16276 OVH	guard	.152	.152	-	1
1200 AMS-IX1	guard	.180	-	.180	55
174 COGENT-174	guard	.095	.007	.088	87

As described in Section 3.2, not all ASes have RIPE Atlas probes installed, $AS_{exit+probes} \rightarrow AS_{destination}$ only represents around 38% of total exit probability and $AS_{guard+probes} \rightarrow AS_{client}$ only represents around 83% of total guard probability. This means real values are estimated to be even higher.

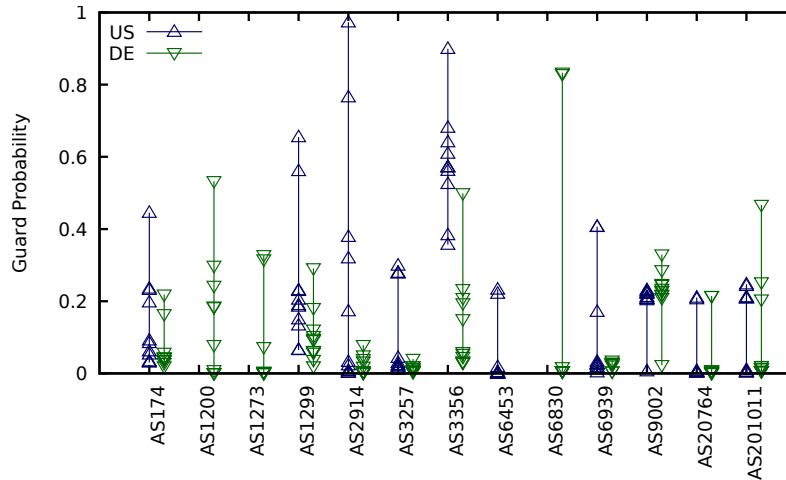
Table 4. Combined results for a single destination and a single client

AS Name	P_{guard}	P_{exit}	$P_{combined}$
24940 HETZNER-AS	.202	.988	.199
1200 AMS-IX1	.180	.068	.012
16276 OVH	.152	.065	.010

4.2 Measurements with Multiple Clients and Multiple Destinations

We conducted the scans on the guard side and exit side separately, and afterward combined the results. We conducted 15,160 successful traceroutes for the 10 entry side ASes originating in the US and Germany. On the destination side, we gathered 4,270 successful traceroute results. The scans were performed around 31.12.2019.

Client to Guard Relays We found ASes that have a high probability to appear on the route to/from guard relays. Figure 7 shows the probability of different ASes to be on a route to/from a guard relay in the US and Germany. The different data points represent the different originating ASes, and the line the min and max values. We identify ASes in good positions for both countries. AS3356 (LEVEL 3) will be traversed with a high probability for all originating ASes and has also a high probability for the set of German probes. We identify AS1200 and AS1273 and AS6830 only in German ASes. AS1299 (TELIANET), AS2914 (NTT-COMMUNICATIONS-2), AS174 (COGENT) and AS9002 (RETN) are strong for both client sets.

**Fig. 5.** Summarized guard probability for ASes that appear on a route from the originating client AS to all guard ASes

Destination Results For the destination set, we are using probes of ten different autonomous systems, derived from the Tranco Top pages list, as explained in Section 3.4. We identified all ASes that were located on the routes for every destination AS. We then combined these values to represent the possibility of a client connecting to all destinations. Table 5 shows all ASes that have a probability over 20%. Figure 6 additionally shows the data points for every single destination AS. We excluded all destination ASes, because they appear with certainty, and excluded ASes that only appear because exit relays are hosted (AS200052, AS208323 - Applied Privacy).

Table 5. Results on the exit side, with a summarized exit probability over 20%

AS Name	P
6939 HURRICANE	0.808
6461 ZAYO-6461	0.510
174 COGENT-174	0.415
1299 TELIANET	0.377
1200 AMS-IX1	0.370
2914 NTT-COMMUN	0.362
10578 GIGAPOP-NE	0.359
3257 GTT-BACKBO	0.290

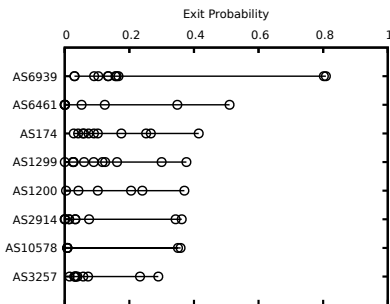


Fig. 6. Summarized probability with single data points representing the different destination ASes

Combined Results Combining all results, we identify ASes that have a high probability to be on the guard side as well as on the exit side. We investigate combinations of single client ASes with all ASes on the exit side, because users connect from one client AS to different destinations. In Figure 7, we can identify strong ASes for our measurement setup. AS3356 (LEVEL3) has a combined value of up to 67.1% ($P_{guard} = .681 \cdot P_{exit} = .985$) for the client AS7018. Other notable ASes are AS6939, AS1299 and AS2914 with combined values $> 20\%$.

5 Discussion

We presented a methodology to utilize the RIPE Atlas network to gather valuable routing data from and to Tor relays. Related work already quantified AS-level adversaries' capabilities for traffic correlation attacks. Thus, our work will not provide any surprising insights. However, our methods can be used to refine existing models with timely and actively gathered routing data. While the result set of this paper is rather limited with only 16,500 executed traceroute commands and a small number of probes utilized, the methodology is highly scalable. For future work, we plan to scale up the number of measurements performed in various ways. First, we want to enlarge the measured client and destination sets.

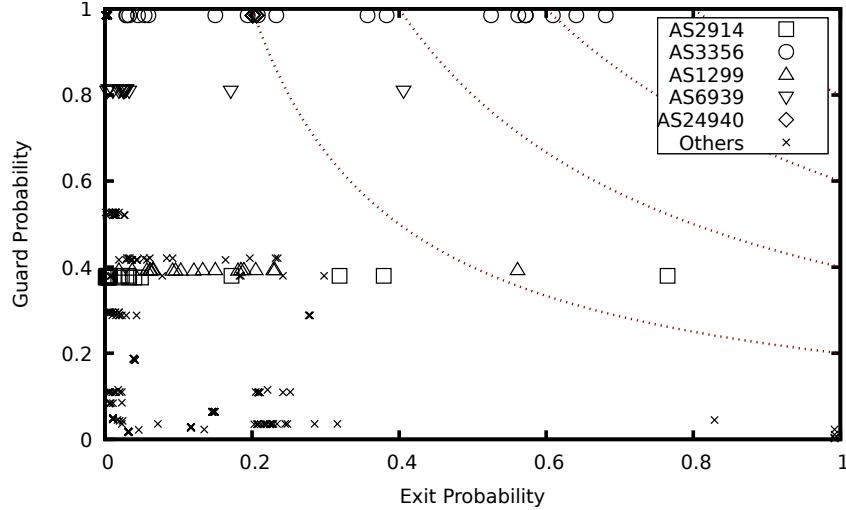


Fig. 7. Combined probability of ASes appearing on the client and destination path

Second, we think about reoccurring scans in contrast to *oneoff* measurements conducted in this paper. Last, a more fine-grained measurement, using probes in the same IP subnets as the relays could improve the results. We publish our source code openly available as free software. This enables other entities, such as large relay operators, to also perform measurements. All measurement results gathered with RIPE Atlas are also openly available and could include valuable results for the Tor network. We argue that large relay operators should deploy RIPE Atlas probes in their networks, not only to further improve our future results but also to enable other measurements. Only a few more probes would increase the coverage significantly. In Section 3.2 we identified the largest relay operators (AS-wise) without RIPE probes. The evaluation illustrates the possibilities of our methodology. However, it is limited in various ways. We currently do not consider various factors that are important to accurately quantify the threat of AS-level observers. This includes user behavior, Tor circuit creation algorithms, and others. Hence, a combination of our data acquisition method with other simulations is necessary to correctly quantify the traffic correlation threat.

Finally, we argue for increased AS diversity in the Tor network. Even with simple measurements, we see that the distribution of Tor relays is skewed. We hope that our measurements can improve the informed decision how this diversity should be achieved.

6 Conclusion

To address Tor traffic correlation attacks through ASes we presented a novel way to analyze the network routes taken by traffic from and to the Tor network.

While previous research relied on the analysis of BGP routing information and simulations, we proposed a new method to utilize the RIPE Atlas framework to measure network routes. We implemented a measurement framework that utilizes the RIPE Atlas probes to perform traceroute commands between clients, servers and Tor endpoints to collect information on the ASes involved in traffic routing. Next, we utilized the collected information to create a model of paths to locate and quantify strong observers.

By leveraging this methodology we were able to identify a small set of ASes which have a great influence on the total amount of Tor bandwidth. This shows that the collected information is a valuable additional data source when analyzing attacks and defenses based on AS topology.

Acknowledgments

We want to thank David Schmidt for his preliminary work on this topic. This research was funded by the Austrian Science Fund (FWF): P30637-N31, the Josef Ressel Center (JRC) project TARGET and the Austrian Research Promotion Agency (FFG) through project AutoHoney(I)IoT. The competence center SBA Research (SBA-K1) is funded within the framework of COMET – Competence Centers for Excellent Technologies by BMVIT, BMDW, and the federal state of Vienna, managed by the FFG.

References

1. Masoud Akhoondi, Curtis Yu, and Harsha V Madhyastha. LASTor: A Low-Latency AS-Aware Tor Client. In *Symposium on Security and Privacy*. IEEE, 2012.
2. Mashael Alsabah and Ian Goldberg. Performance and Security Improvements for Tor: A Survey. *ACM Computing Surveys (CSUR)*, 2016.
3. Collin Anderson, Philipp Winter, et al. Global Network Interference Detection Over the RIPE Atlas Network. In *USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2014.
4. Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwälder. Lessons learned from using the ripe atlas platform for measurement research. *ACM SIGCOMM Computer Communication Review*, 2015.
5. Armon Barton and Matthew Wright. DeNASA: Destination-Naive AS-Awareness in Anonymous Communications. *Proceedings on Privacy Enhancing Technologies*, 2016.
6. Armon Barton, Matthew Wright, Jiang Ming, and Mohsen Imani. Towards Predicting Efficient and Anonymous Tor Circuits. In *USENIX Security Symposium*, 2018.
7. Roger Dingledine. The lifecycle of a new relay, 2013.
8. Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The Second-Generation Onion Router. In *USENIX Security Symposium*, 2004.
9. Matthew Edman and Paul Syverson. AS-awareness in Tor path selection. In *Conference on Computer and Communications Security*. ACM, 2009.

10. Rod  rick Fanou, Pierre Francois, and Emile Aben. On The Diversity of Interdomain Routing in Africa. In *Passive and Active Network Measurement Conference*, 2015.
11. Nick Feamster and Roger Dingledine. Location Diversity in Anonymity Networks. In *Workshop on Privacy in the Electronic Society*. ACM, 2004.
12. Phillipa Gill, Michael Schapira, and Sharon Goldberg. Modeling on quicksand: Dealing with the scarcity of ground truth in interdomain routing data. *ACM SIGCOMM Computer Communication Review*, 2012.
13. Vasileios Giotsas, Matthew Luckie, Bradley Huffaker, and kc claffy. Inferring Complex AS Relationships. In *Internet Measurement Conference*. ACM, 2014.
14. Hans Hanley, Yixin Sun, Sameer Wagh, and Prateek Mittal. DPSelect: A Differential Privacy Based Guard Relay Selection Algorithm for Tor. *Proceedings on Privacy Enhancing Technologies*, 2019.
15. Nicholas Hopper, Eugene Y Vasserman, and Eric Chan-Tin. How Much Anonymity does Network Latency Leak? *ACM Transactions on Information and System Security (TISSEC)*, 2010.
16. Aaron Johnson, Chris Wacek, Rob Jansen, Micah Sherr, and Paul Syverson. Users get routed: Traffic correlation on Tor by realistic adversaries. In *Conference on Computer and Communications Security*. ACM, 2013.
17. Prateek Mittal, Ahmed Khurshid, Joshua Juen, Matthew Caesar, and Nikita Borisov. Stealthy Traffic Analysis of Low-Latency Anonymous Communication Using Throughput Fingerprinting. In *Conference on Computer and Communications Security*. ACM, 2011.
18. Milad Nasr, Alireza Bahramali, and Amir Houmansadr. DeepCorr: Strong Flow Correlation Attacks on Tor Using Deep Learning. In *Conference on Computer and Communications Security*. ACM, 2018.
19. Rishab Nithyanand, Oleksii Starov, Adva Zair, Phillipa Gill, and Michael Schapira. Measuring and Mitigating AS-level Adversaries Against Tor. In *Network and Distributed System Security Symposium (NDSS)*, 2016.
20. Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczynski, and Wouter Joosen. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *Network and Distributed System Security Symposium*, 2019.
21. RN Staff. RIPE Atlas: A global internet measurement network. *Internet Protocol Journal*, 18(3), 2015.
22. Yixin Sun, Anne Edmundson, Nick Feamster, Mung Chiang, and Prateek Mittal. Counter-RAPTOR: Safeguarding Tor Against Active Routing Attacks. In *Symposium on Security and Privacy*. IEEE, 2017.
23. Yixin Sun, Anne Edmundson, Laurent Vanbever, Oscar Li, Jennifer Rexford, Mung Chiang, and Prateek Mittal. RAPTOR: Routing Attacks on Privacy in Tor. In *USENIX Security Symposium*, 2015.
24. Laurent Vanbever, Oscar Li, Jennifer Rexford, and Prateek Mittal. Anonymity on QuickSand: Using BGP to Compromise Tor. In *Proceedings of the 13th ACM Workshop on Hot Topics in Networks*. ACM, 2014.
25. Chris Wacek, Henry Tan, Kevin S Bauer, and Micah Sherr. An Empirical Evaluation of Relay Selection in Tor. In *Network and Distributed System Security Symposium*, 2013.
26. Gerry Wan, Aaron Johnson, Ryan Wails, Sameer Wagh, and Prateek Mittal. Guard Placement Attacks on Path Selection Algorithms for Tor. *Proceedings on Privacy Enhancing Technologies*, 2019.