



Security and Performance Implications of BGP Rerouting-Resistant Guard Selection Algorithms for Tor

Asya Mitseva, Marharyta Aleksandrova, Thomas Engel, Andriy Panchenko

► To cite this version:

Asya Mitseva, Marharyta Aleksandrova, Thomas Engel, Andriy Panchenko. Security and Performance Implications of BGP Rerouting-Resistant Guard Selection Algorithms for Tor. 35th IFIP International Conference on ICT Systems Security and Privacy Protection (SEC), Sep 2020, Maribor, Slovenia. pp.219-233, 10.1007/978-3-030-58201-2_15 . hal-03440868

HAL Id: hal-03440868

<https://inria.hal.science/hal-03440868>

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

Security and Performance Implications of BGP Rerouting-resistant Guard Selection Algorithms for Tor

Asya Mitseva¹, Marharyta Aleksandrova¹,
Thomas Engel¹, and Andriy Panchenko²

¹ University of Luxembourg, Esch-sur-Alzette, Luxembourg
{firstname.lastname}@uni.lu

² Brandenburg University of Technology, Cottbus, Germany
{firstname.lastname}@b-tu.de

Abstract. Tor is the most popular anonymization system with millions of daily users and, thus, an attractive target for attacks, e.g., by malicious autonomous systems (ASs) performing active routing attacks to become man in the middle and deanonymize users. It was shown that the number of such malicious ASs is significantly larger than previously expected due to the lack of security guarantees in the Border Gateway Protocol (BGP). In response, recent works suggest alternative Tor path selection methods preferring Tor nodes with higher resilience to active BGP attacks. In this work, we analyze the implications of such proposals. We show that Counter-RAPTOR and DPSelect are not as secure as thought before: for particular users they allow for leakage of user’s location. DPSelect is not as resilient as widely accepted as we show that it achieves only one third of its originally claimed resilience and, hence, does not protect users from routing attacks. We reveal the performance implications of both methods and identify scenarios where their usage leads to significant performance bottlenecks. Finally, we propose a new metric to quantify the user’s location leakage by path selection. Using this metric and performing large-scale analysis, we show to which extent a malicious middle can fingerprint the user’s location and what kind of confidence it can achieve. Our findings shed light on the implications of path selection methods on the users’ anonymity and the need for further research.

Keywords: BGP Routing Attacks · Tor · Onion Routing · Privacy.

1 Introduction

In the age of mass surveillance and censorship, users rely on anonymization techniques to exercise their right to freedom of expression and to freely access information. Currently, Tor [7] is the most popular low-latency anonymization network designed to hide users’ identities (i.e., IP addresses) from service providers and to prevent third parties from exposing the relationship between communicating partners on the Internet. To accomplish this goal, the user traffic sent via Tor is encrypted in multiple layers and forwarded through three Tor

nodes, known as *entry*, *middle*, and *exit*. Due to its popularity, it is an attractive target for adversaries aiming to compromise Tor users, e.g., by applying traffic analysis. An adversary who simultaneously observes the user traffic entering and exiting Tor can perform traffic correlation on packet sizes and timing and, thus, deanonymize user connections [15,23]. Website fingerprinting (WFP) is another type of traffic analysis, where the adversary aims to identify the website visited by a Tor user by observing patterns of data flows between the user and its entry node [14,22]. A common example of entities in the position to execute both types of attacks are autonomous systems (ASs), often called network-level adversaries, which lie on the path between the Tor user and its destination. Recent studies [23,25] have shown that natural Internet routing dynamics and active attacks against Border Gateway Protocol (BGP), the de-facto standard inter-domain routing protocol, dramatically increase the number of ASs that are in a position to compromise Tor traffic by applying traffic correlation or WFP.

In response, several works have focused on developing sophisticated Tor path selection methods for choosing entry and exit nodes that consider not only the nodes' capacity but also the presence of asymmetric routing, potentially colluding ASs, and the robustness of network paths against active routing attacks [2,21]. The most recent proposals are Counter-RAPTOR [24] and DPSelect [11]. Counter-RAPTOR is an alternative path selection method preferring entries with higher resilience to BGP attacks, as estimated based on the network location of the Tor user. The hardened protection against active routing attacks provided by this method, though, negatively influences the randomness of entry node selection and leaks information about the user's location. This allows an attacker to link a user to its AS [11]. Such a leak can be further exploited to deanonymize users. DPSelect aims to overcome this drawback by using *differential privacy*.

Hanley et al. [11] have recently explored the privacy loss of both Counter-RAPTOR and DPSelect with respect to the fingerprintability of user ASs by using notions of entropy. However, their evaluation is based on a small dataset of 95 ASs only. Hence, the results cannot be generalized to all Tor users. Moreover, no existing work has examined the potential threats of malicious middle nodes to deanonymize Tor users utilizing rerouting-resistant path selection. Recent study has shown that traffic analysis from middle nodes can be as effective as from entry and exit positions in case of Tor onion services, although middles neither directly know the user nor its destination [14]. In this work, we show to which extent a malicious middle can close this gap and successfully mount a deanonymization attack by localizing a user. Our contributions are as follows:

1. We show that Counter-RAPTOR and DPSelect are not as secure as previously thought: about 20% of users that rely on these methods select an entry node from their country with five times higher probability than vanilla Tor. This exposes user location and seriously endangers user's anonymity. Our method allows users to assess their vulnerability.
2. DPSelect is not as resilient as widely accepted. We show that it achieves only one third of its originally claimed resilience and, hence, does not protect users from active routing attacks.

3. We identify scenarios where the usage of both path selection methods leads to serious performance bottlenecks, as they prefer poorly-performing entries.
4. We propose a new metric, *confidence increase*, to quantify the user’s location leakage by Counter-RAPTOR and DPSelect. Using this metric, we show to which extent a malicious middle can fingerprint the location of a user and what kind of confidence it can achieve.

We perform a large-scale analysis and collect the most comprehensive set of ASs containing Tor users to analyze the security properties of Counter-RAPTOR and DPSelect. Our analysis allows to better understand the properties of these methods and warn vulnerable users about possible implications.

2 Background

Internet routing: The Internet comprises a large set of interconnected ASs identified by unique AS numbers (ASNs) [18]. Each AS possesses a set of delegated IP addresses that are aggregated into blocks (i.e., IP prefixes) and is responsible for forwarding traffic to and from them. Some ASs, called transit ASs, are also able to forward traffic whose source and destination are not in their IP prefixes. ASs set up dedicated links between each other, based on confidential business agreements, and distribute reachability data using BGP. Since BGP does not provide any security guarantees, ASs can manipulate the global routing by distributing bogus data [18]. An AS is capable of claiming that it originates an IP prefix not delegated to it, known as a *prefix hijack*, and attract a fraction of Internet traffic to it. As this AS does not possess any valid route to the victim AS, the redirected traffic will be dropped and affected users will experience connectivity problems. A more sophisticated attack, BGP *interception* [18], is an improved version of prefix hijacks, where the malicious AS has a valid route to the victim AS. It can not only redirect traffic through itself, but also forward it via a detour to the real destination without disturbing the connectivity. These BGP attacks have been often exploited for country-level censorship [26] and tracking users of anonymization networks such as Tor [23].

Tor [7] is the most popular anonymization network designed for low-latency applications, e.g., web browsing. The traffic between a Tor user and its destination is sent via a virtual tunnel (i.e., *circuit*), over three nodes, known as *onion relays* (ORs). Information about identities and status of the available ORs is periodically distributed to Tor users in the form of a *consensus* document. Tor users select ORs for circuits probabilistically according to the ORs’ bandwidth, availability, and exit policy to a given target. After negotiating a symmetric key between each OR in the circuit and the user, the user data is encrypted in multiple layers using these keys [7]. While forwarding the user data, each OR on the path removes (or adds, depending on the direction) one layer of encryption. Thus, none of the ORs in the circuit knows both the user and its destination at the same time. Each user also maintains a list of preselected entries, called *guard set*, to reduce the information leakage caused by the frequent selection of new entries for each new circuit. From its guard set, the user chooses a single entry

(i.e., *guard*) as the first hop for its circuits and continues using it for months [8]. Although the use of a guard reduces the probability of picking a malicious OR in a short time period, it does not prevent malicious ASs from being on the path between the Tor user and its guard and compromising user traffic.

Counter-RAPTOR: Sun et al. [24] proposed an enhancement to the original Tor path selection algorithm, called Counter-RAPTOR, aiming to decrease the probability of an AS actively putting itself on the path between a Tor user and its guard. According to the proposal, the user considers not only the available OR’s bandwidth $B(i)$ but also the OR’s resilience $R(i)$ to hijack attacks when choosing a guard. The OR’s resilience value indicates the fraction of ASs that will not succeed in hijacking the user traffic sent to the OR by falsely claiming to originate the IP prefix containing that OR. In other words, the probability of a guard i being selected is proportional to its weight $W(i)$:

$$W(i) = \alpha \cdot R(i) + (1 - \alpha) \cdot \bar{B}(i), \quad (1)$$

where $\bar{B}(i)$ is the OR’s bandwidth normalized in the range $[0, 1]$ and α is a configurable parameter to balance between the OR’s resilience to hijack attacks and its bandwidth. To limit a user-specific guard selection, a random sampling is applied to the ORs’ resilience values to produce a more uniform choice of guard. While not stated differently, in our work we use $\alpha = 0.5$ as recommended in [24].

DPSelect: Despite the use of random sampling to pick a guard from a set of ORs with high resilience values, Hanley et al. [11] showed that Counter-RAPTOR still leaks information about user locations. Consequently, the authors proposed DPSelect, which integrates a differential privacy metric into the weight function (1) of Counter-RAPTOR. This metric is intended to bound the difference between the largest probability of a user selecting a given guard and the least probability of another user choosing the same guard and, thus, prevents a statistical correlation between a guard selected by a user and the AS of that user. To ensure guard selection homogeneity among users, DPSelect relies on an exponential mechanism to compute the weight function $W(i)$ of a guard i :

$$W(i) = e^{\epsilon \cdot (\alpha \cdot R(i)^{x_1} + (1 - \alpha) \cdot B(i)^{x_2})}, \quad (2)$$

where ϵ defines how private the guard selection should be and x_1 and x_2 are optimization parameters aiming to preserve the main goal of the original Counter-RAPTOR approach with respect to high bandwidth and resilience values of the considered OR. Hanley et al. apply a Monte-Carlo sampling-based method with equally-weighted resilience and bandwidth values (i.e., $\alpha = 0.5$) to tune x_1 and x_2 and, so, achieve a reasonable trade-off between OR’s resilience and bandwidth.

3 Related Work

Threat of network-level attackers: The threat of an AS simultaneously observing both ends of Tor user connections was first examined by Feamster and Dingledine [10], who detected that up to 30% of randomly generated Tor circuits are vulnerable to an AS adversary. Due to the increased number of ASs carrying Tor traffic over the years, the natural intuition was that the likelihood of a single

AS being able to observe user traffic entering and exiting Tor reduced. In [9], this assumption was verified by using an updated model for Tor and showed that the risk of deanonymization by a single AS is not reduced. Another work [20] explored the threat to Tor users posed by Internet eXchange points (IXP) – a shared physical infrastructure in a single location connecting several ASs. The authors showed that an attacker, who is positioned at an IXP and observes the traffic passing through any AS co-located at that IXP, can correlate high-speed network flows even at low rates of sampling and compromise users’ anonymity.

Juen et al. [17] questioned the accuracy of the AS path inference methods used to evaluate the vulnerability of Tor to AS attackers. Wacek et al. [27] used traceroute data from the Center for Applied Internet Data Analysis (CAIDA) [6] to reconstruct the AS interconnectivity and showed that the same AS may still appear in both ends of 27.4% of randomly created Tor circuits. Johnson et al. [15] explored the amount of time needed by an attacker controlling a set of ASs or IXPs to compromise Tor circuits. Although the user’s security strongly depends on its location, they showed that an attacker possessing several ASs or IXPs has a much greater compromise speed, even against users in safer locations. Sun et al. [23] showed that traffic correlation attacks succeed even when an AS observes paths in different directions on both ends of a Tor circuit. In response, Nithyanand et al. [21] reevaluated the threat posed by these attacks and discovered that up to 40% of Tor circuits are vulnerable to traffic correlation by single ASs, 42% by colluding ASs, and 85% by state-level (i.e., the set of ASs located in a single country) attackers. In [23], the authors also showed that BGP hijack and interception attacks used to redirect Tor traffic can dramatically increase the likelihood of an AS eavesdropping on both ends of Tor connections. Tan et al. [25] extended this analysis and detected that more than 90% of the total bandwidth available in Tor is vulnerable to BGP hijack attacks.

Defenses against network-level attackers: The idea of avoiding a single AS that appears on both ends of a Tor circuit when choosing entry and exit was first proposed by Feamster and Dingledine [10] and developed by Edman and Syverson [9] by using a snapshot of the current AS topology. *LASTor* [1] is another alternative path selection method, which predicts the ASs through which the user traffic entering and exiting Tor is highly likely to be routed. However, Wacek et al. [27] showed that almost 25% of *LASTor* circuits remain vulnerable to network-level adversaries. To limit the effectiveness of hijack attacks used to redirect Tor traffic, Tan et al. [25] relied on periodical traceroute measurements to detect guards under active BGP attack and prevent users from selecting them. However, this method can neither detect short-lived BGP attacks nor protect already established circuits [24,23]. *DeNASA* [2] is another Tor path selection method that avoids a predefined set of large ASs often appearing on both sides of Tor circuits. *Astoria* [21] considers OR capacity, asymmetric routing, and potential colluding ASs during circuit creation. Contrary to Counter-RAPTOR and DPSelect, *DeNASA* and *Astoria* only focus on passive AS-level attackers. Both methods have also been shown to be user location-dependent and vulnerable to network-level attackers who can exploit user behavior over time to compromise anonymity [16,28]. Wails et al. [28] raised the further criticism that none of the

proposed location-aware approaches, including Counter-RAPTOR and DPSelect, consider user mobility, which dramatically reduces the anonymity provided by these methods over time. Wan et al. [29] showed that an attacker can exploit the location awareness of the methods and strategically launch ORs in locations that increase their likelihood of being selected as guards by target users.

4 Datasets

To thoroughly analyze the impact of the modified path selection schemes Counter-RAPTOR and DPSelect on privacy of the users, there is a need to use representative real-world data. To this end, we gathered information about (i) available guards and the ASs they are located in, (ii) the set of all possible user ASs, and (iii) existing AS relationships. To obtain Tor network data, we downloaded consensus from CollecTor³ for March 1, 2017 and extracted guards, whose IP addresses were mapped to ASs using Maxmind GeoIP database⁴. In total, we obtained 2,451 guards belonging to 475 unique ASs in 50 countries.

To acquire all possible user ASs, we collected all known ASs from CAIDA [3] during March 2017, comprising 57,015 unique ASNs, and filtered out transit and content hosting ASs (as these do not contain end-users). In total, we obtained 30,848 possible user ASs. As we are interested to know the estimated number of end-users in a given AS for our analysis, we also collected the number of IP addresses delegated to each of the ASs by using CAIDA AS Ranking dataset [5]. From our set of user ASs, we excluded 2,361 ASs for which we could not infer any data about delegated IP addresses and 2,606 ASs for which Counter-RAPTOR and DPSelect cannot compute a network path between these ASs and the collected guard ASs (breadth-first search of these methods may discard unprofitable AS relationship and yield no valid path). Finally, we obtained 25,881 possible user ASs distributed in 223 countries, where 81.4% of them were located in countries containing guard nodes too. We refer to this dataset as D .

For further analysis, we applied the method proposed in [27] to construct a reduced map of the Internet including latency measurements between hosts. We were able to extract latency values for a fully-connected graph of ASs containing 333 guard ASs and 7,052 user ASs from our initial sets of guard and user ASs. We refer to this dataset as D_{lat} . The guard ASs in D_{lat} cover 88.9% of all available guards in D and are located in 48 countries. The user ASs in D_{lat} contain 91% of all IP addresses delegated to the set of user ASs in D and are distributed in 187 countries. Moreover, 80.6% of the user ASs in D_{lat} are located in countries containing guards from D_{lat} . Table 1 summarizes the statistics for the sets of collected user and guard ASs.

Like [11,24], we used CAIDA AS Relationship dataset [4] to acquire data about existing commercial AS relationships and inferred network paths between ASs necessary to compute the resilience values between user and guard ASs.

DPSelect parameters for D and D_{lat} : DPSelect weight function (2) contains two parameters, x_1 and x_2 , that need to be tuned with regard to the set of

³ <https://metrics.torproject.org/collector.html> ⁴ <https://dev.maxmind.com/geoip/geoip2/geolite2/>

Table 1: Statistics for collected user and guard ASs.

Description	Number	Countries	Guards	Dataset
Total number of collected ASs	57,015	230	–	–
Total number of possible user ASs	25,881	223	2,451	D
Total number of guard ASs	475	50	2,451	
Number of user ASs with latency	7,052	187	2,180	D_{lat}
Number of guard ASs with latency	333	48	2,180	

possible user ASs and the current Internet topology. In [11], the authors showed that $x_1 = 2$ and $x_2 = 0.75$ are optimal for the set of top-93 Tor user ASs and $\alpha = 0.175$ achieves a reasonable trade-off between ORs resilience and bandwidth. As the number of user ASs in our dataset is larger by several orders of magnitude, we repeated the optimization procedure proposed in [11] before analyzing the security properties of DPSelect. However, we did not observe any appreciable impact of the newly obtained parameters on the average user resilience and bandwidth. Thus, we stuck to the optimal values as proposed in [11].

5 Vulnerabilities in Counter-RAPTOR and DPSelect

In this section, we analyze information leakage in Counter-RAPTOR and DPSelect and summarize our key findings. In particular, we show that DPSelect is not as secure concerning active routing attacks as widely accepted as it achieves only one third of its originally claimed resilience. Then, we show that for 20% of the users both methods select a guard from their country with five times higher probability than vanilla Tor and propose a method allowing users to assess their vulnerability. We also identify scenarios where the usage of both path selection methods leads to significant performance bottlenecks. Finally, we propose a new metric that allows users to assess their location leakage with respect to a malicious middle node and bound its confidence.

Comprehensive revision of DPSelect: The main goal of DPSelect is to prevent information leakage concerning user locations while achieving resilience to hijack attacks as by Counter-RAPTOR. Hanley et al. [11] evaluated the performance of DPSelect in terms of user resilience for the set of top-93 Tor user ASs and showed that DPSelect achieves very similar average user resilience to Counter-RAPTOR. Our analysis reveals that these results cannot be generalized for all Tor users. As shown in Figure 1, the user resilience achieved by DPSelect (densely dotted line) degrades significantly compared to Counter-RAPTOR (dash-dotted line). While for 57% of Tor users Counter-RAPTOR provides up to a 70% probability of being resilient to hijack attacks, this probability reduces by 10% for DPSelect. Compared to vanilla Tor, DPSelect is able to improve the average user resilience only by 12.5% while Counter-RAPTOR achieves an increase of up to 30.5%.

To identify the reason for the significant difference in user resilience obtained by Counter-RAPTOR and DPSelect, we revisited the implementations of both




Fig. 1: CDF of the average user resilience for different guard selection algorithms.

methods provided by the authors. To compute the weight function (2), DPSelect relies on a resilience probability $R_{prob}(i) = \frac{R(i)}{\sum_i R(i)}$ and a bandwidth probability $B_{prob}(i) = \frac{B(i)}{\sum_i B(i)}$ instead of a resilience value $R(i)$ and a normalized bandwidth $\bar{B}(i)$ used by Counter-RAPTOR. To see the impact of the adjusted values, we modified the original code of Counter-RAPTOR to use resilience and bandwidth probabilities. Figure 1 shows that DPSelect (densely dotted line) and the modified Counter-RAPTOR (loosely dotted line) produce very similar user resilience.

To sum up, we showed that DPSelect does not provide a user resilience as high as the original Counter-RAPTOR method. Hence, Counter-RAPTOR still remains the only alternative for keeping Tor users resilient against hijack attacks.

Information leaks in Counter-RAPTOR and DPSelect: To evaluate whether Tor users relying on Counter-RAPTOR and DPSelect increase their probability of selecting a guard from the same country as the user, denoted by p_C , we consider all guards and only those user ASs that are located in countries containing at least one guard from our dataset D . We obtained 21,064 user ASs comprising 81% of all user ASs in D . We computed how many times p_C increases compared to p_C for vanilla Tor. Figure 2 shows that for 80% of our users utilizing the location-aware methods, the probability of selecting a guard located in the same country as the user increases. For roughly 20% of users relying on Counter-RAPTOR ($\alpha = 0.5$), p_C increases by more than five times and for around 10% of DPSelect users, p_C increases by nine times. Moreover, the higher values of α in case of Counter-RAPTOR (i.e., the user is paying more attention to its resilience to hijack attacks than to the performance of its circuits) result in significantly higher probability of choosing a guard from the same country as the user.

We also examined information leaks in Counter-RAPTOR and DPSelect using the number of ASs between a Tor user and its guard and geographical distance between a Tor user and its guard (using GeoIP data). However, we did not observe any significant correlation when applying both methods.

In summary, Counter-RAPTOR and DPSelect users are more likely to choose a guard from the same country as the user, in contrast to vanilla Tor users. For 20% of the users, this probability is five times higher. This information leak can be further exploited to improve guard placement attacks, as proposed in [29].

Fingerprinting user locations: As shown above, Counter-RAPTOR and DPSelect usually leak information about a user by allowing an attacker to link a





Fig. 2: CDF of the ratio of p_C for location-aware methods to p_C for vanilla Tor.

user to its location. To quantify this information leak, we propose a new metric, *confidence increase*, and we use it to measure the increase of the attacker's confidence from a middle position⁵ about the location of Tor users. Our metric shows the ratio of the cumulative probability of users in the top- N most probable ASs selecting a guard to the cumulative fraction of IP addresses delegated to those ASs. The larger is the fraction of IP addresses delegated to the user ASs and the lower is the cumulative probability of those ASs, the smaller is the confidence increase achieved by the attacker for the user location.

Inspired by Gini index, confidence increase allows to account for small values of probabilities due to the large number of possible user ASs and to focus on inequality of distribution among the top most preferable ASs. A traditional metric, i.e., *decrease in entropy*, as compared to the baseline (when the probability is uniformly distributed among all user ASs), depends on the total number of available user ASs. As opposite, if the number of IP addresses between user ASs is uniformly distributed, the value of confidence increase will be constant. In case of non-uniform distribution of the number of IP addresses, confidence increase allows to additionally capture the information leak related to the association of high probability values with ASs that have a small number of potential users.

By using our dataset D , we first computed the Gini index for both location-aware methods and observed deviations from the uniform distribution of guards in case of vanilla Tor. Thus, we focus on the top of the inequality of the distributions for both methods, where our metric is more illustrative. We established $N = 25\%$ (inequality of distribution in the first quarter of top ASs) as a reasonable threshold for the efficiency of our metric. Figure 3a shows that the confidence increase for vanilla Tor (dashed gray line) always equals one as the probability of a user being located in a given AS is distributed uniformly among all user ASs and the cumulative probability of the first quarter of user ASs is always 0.25. As the cumulative fraction of IP addresses delegated to these user ASs changes depending on the chosen subset of user ASs, we take the mean of all combinations, which is equal to 0.25. Contrariwise, Counter-RAPTOR (dash-dotted gray line) and DPSelect (densely dotted gray line) leak a considerable amount of user-specific information. For more than 40% of guards selected by Counter-RAPTOR users and more than 20% of guards chosen by DPSelect users, the confidence increase of a middle connected to these guards is 100% higher than

⁵ This metric can be used for other scenarios as well, not only from a middle position.



(a) Influence of latency measures ($\theta = 5$). (b) Varying accuracy of latency measures.

Fig. 3: CDF of the confidence increase for different guard selection algorithms.

for vanilla Tor. Even worse, the confidence of the middle increases by 300% for 15% of guards connected to Counter-RAPTOR users and almost 10% of guards chosen by DPSelect users. Although the overall confidence increase for DPSelect is less than for Counter-RAPTOR, the protection provided by DPSelect is far smaller than originally claimed.

Next, we examine if the use of latency-based attacks proposed in [12] can further increase the attacker’s confidence about the location of Tor users. In 2007, Hopper et al. showed that an attacker, who has access to a web server, a network coordinate system, and an OR, can estimate the latency between a Tor user connected to the adversarial web server and its guard and localize the user. To do this, the authors apply the Tor circuit clogging attack suggested by Murdoch and Danezis [19] to detect the ORs utilized by the user in its circuit. Then, the attacker creates its own circuit via the same ORs to estimate the latency between the guard and the exit. The latency between the user and its exit is also easily measured as the user is connected to the adversarial web server. By subtracting both latency estimates, the attacker obtains the latency between the user and its guard. The adversary further utilizes the network coordinate system to compute a set of possible latencies between potential users and guards and localize the user based on the estimated latency. In our work, we assume even a weaker attacker model where the adversary controls only at least one middle node. For all user connections traversing an adversarial middle, the attacker knows the identities of the guards and exits utilized by these users. This allows an estimation of the latency between the guard and the middle and between the user and the middle and a computation of the latency between the user and its guard. The rest of the attack remains identical to [12].

For our evaluation, we used the dataset D_{lat} containing latency measures for each pair of ASs. We assume that each user AS contains at least one victim user that the attacker is attempting to compromise, and iterate through each user AS as a potential adversarial target. As the latency estimates measured by the attacker are not precise, we consider a reduced set of potential user ASs whose estimated latency to a guard is in the interval $r \in [lat_{meas} - \theta, lat_{meas} + \theta]$ where lat_{meas} is the real latency between the target user and its guard and θ is a configurable parameter indicating the inaccuracy of estimated latency measured by the attacker. Once the adversary measured the latency between




Fig. 4: Average throughput and download time for the first experiment.

all candidate user ASs and a guard, user ASs whose latencies are significantly different from the estimated latency of the target user (i.e., those outside the interval r) can be excluded from consideration. For the reduced number of user ASs, the attacker computes the confidence increase as described above. As shown in Figure 3a, for more than 40% of guards selected by Counter-RAPTOR users (dash-dotted black line) and more than 20% of guards chosen by DPSelect users (densely dotted black line), the confidence increase of a middle is nearly 200% higher when the attacker relies on latency estimates, compared to the confidence increase obtained for these methods without any latency measures, and almost 400% higher compared to vanilla Tor. Hence, taking into account latencies allows to further narrow down the possible location of the user.

Lastly, we examine the confidence increase of an adversarial middle when the accuracy of attacker’s latency measurements vary. We computed the confidence increase of the middle for different values of θ . As the change of the confidence increase for vanilla Tor was negligible for $\theta \in \{0, 5, 10, 25, 50\}$, for this method we present only $\theta = 0$ for simplicity. Figure 3b shows that, as expected, the attacker gains less information about the user location when the latency estimates are less precise. Still, even in the worst case scenario ($\theta = 50$) the confidence increase is doubled for nearly 50% of guards chosen by Counter-RAPTOR and DPSelect users compared to vanilla Tor. Hence, even imprecise latency measurements help to significantly boost the confidence increase in identifying a user location.

To sum up, we proposed a new metric for quantifying information leakage about user’s location when using Counter-RAPTOR and DPSelect. Using this metric, we showed that both location-aware methods strengthen the attacker’s ability to fingerprint user locations from a middle node. The impact becomes considerably higher when our approach is enhanced with the latency information.

Performance analysis: In [24,11], the authors examined the performance of Counter-RAPTOR and DPSelect and showed that both methods achieve very similar average bandwidth in guard selections for the top-93 Tor user ASs to vanilla Tor. We verify if these results can be generalized for all Tor users. To do this, we recompute the average bandwidth of selected guards for our dataset D . We observe that the average bandwidth of guards chosen by Counter-RAPTOR users reduces by 52% compared to the average bandwidth of guards selected by vanilla Tor users. This drop is by 30% for DPSelect. Thus, we can conclude that – when considering all users – both location-aware methods have significantly lower average bandwidth of selected guards than previously expected. In the rest, we elaborate on the reasons for this and describe our performance analysis.

A major shortcoming of the previous analysis in [24,11] is the fact that all guards in the simulated Tor network are highly-performing ORs (i.e., middles were the bottleneck for Tor circuits). Such analysis cannot capture the real impact of both methods on Tor performance, as the guards (in the case of saturated middles) do not have major influence on the performance of the circuits. The described situation in the previous analysis is a particular case and can change with the evolution of Tor. Thus, we revisit the performance of both methods by using Shadow [13] with the same number of simulated Tor users and ORs as in [24,11] but a slightly modified configuration of the Tor network as follows. First, the bandwidth of each middle and exit is higher than the maximum bandwidth available to each guard in order to avoid network bottlenecks created by middle or exit ORs. Second, the latency between each pair of users and ORs is equal in order to eliminate the impact of latency on the speed and quality of user connections. As Counter-RAPTOR and DPSelect require meaningful IP addresses for their operation, we assigned randomly chosen user IP addresses from our set of user ASs in D . As shown in Figure 4, while the download time for users relying on DPSelect and Counter-RAPTOR with $\alpha = 0.5$ to browse the web is similar to that of web users utilizing vanilla Tor, the download time for Counter-RAPTOR users paying more attention on their resilience to hijack attacks ($\alpha = 1$) increases substantially. In the case of bulk users (i.e., users downloading files of large size), the increase in the download time for users utilizing Counter-RAPTOR with $\alpha = 1$ is by nearly 10%. Moreover, the average sender and receiver throughput for all Tor nodes reduces significantly for DPSelect.

We modify the experiment presented above such that our simulated Tor network contains highly resilient low-bandwidth guards and highly-performing low-resilient guards. To this end, we assigned IP addresses to our users belonging to a single AS and distributed the guards in two other ASs, whose network paths with the user AS are low- and high-resiliently, respectively. The rest of the simulation configuration is the same as described above. As shown in Figure 5, we observe a significant increase of the download time for both, Counter-RAPTOR and DPSelect, users. While the download time for web users utilizing vanilla Tor does not exceed 0.3 seconds, only nearly 60% of Counter-RAPTOR users for $\alpha = 1$ are able to load a website within this time period. This drop is even worse for bulk users, whose download time increases by almost 20% for Counter-RAPTOR with $\alpha = 0.5$ and DPSelect and to almost 80% for Counter-




Fig. 5: Average throughput and download time for the second experiment.

RAPTOR with $\alpha = 1$. Like the previous experiment, the average sender and receiver throughput for all Tor nodes drops dramatically for DPSelect.

To sum up, we showed that Counter-RAPTOR and DPSelect negatively influence the Tor performance. We identified scenarios where the usage of both methods leads to significant bottlenecks, as users prefer poorly-performing guards.

6 Conclusion

We analyzed the susceptibility of the most recent location-aware Tor path selection methods Counter-RAPTOR and DPSelect with regard to malicious middle ORs. To do this, we collected comprehensive set of ASs containing at least one Tor user and available guards in Tor. We showed that both methods are not as secure as thought before: for some users they leak their location. Moreover, DPSelect is not as resilient as widely accepted. We showed that it achieves only one third of its originally claimed resilience. Hence, it does not protect users from routing attacks. We proposed a new metric to quantify the user’s location leakage and with its help performed a large-scale analysis to show to which extent a malicious middle can fingerprint the location of a user and what kind of confidence it can achieve. We also revealed the performance implications of both methods and identified scenarios where their usage leads to significant bottlenecks that were not originally anticipated. Our findings shed light on the implications of both location-aware methods on users’ anonymity and the need for further research.

References

1. Akhoondi, M., et al.: LASTor: A Low-Latency AS-Aware Tor Client. In: IEEE S&P (2012)

2. Barton, A., Wright, M.: DeNASA: Destination-Naive AS-Awareness in Anonymous Communications. In: PETS (2016)
3. CAIDA: AS Classification. <https://www.caida.org/data/as-classification/>
4. CAIDA: AS Relationships. <http://www.caida.org/data/as-relationships/>
5. CAIDA: ASRank. <https://asrank.caida.org/>
6. CAIDA: The IPv4 Routed /24 Topology Dataset. https://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml
7. Dingledine, R., et al.: Tor: The Second-generation Onion Router. In: USENIX Security (2004)
8. Dingledine, R., et al.: One Fast Guard for Life (or 9 months). In: HotPETs (2009)
9. Edman, M., Syverson, P.: AS-awareness in Tor Path Selection. In: ACM CCS (2009)
10. Feamster, N., Dingledine, R.: Location Diversity in Anonymity Networks. In: ACM WPES (2004)
11. Hanley, H., et al.: DPSelect: a Differential Privacy Based Guard Relay Selection Algorithm for Tor. In: PETS (2019)
12. Hopper, N., et al.: How Much Anonymity Does Network Latency Leak? In: ACM CCS (2007)
13. Jansen, R., Hopper, N.: Shadow: Running Tor in a Box for Accurate and Efficient Experimentation. In: NDSS (2012)
14. Jansen, R., et al.: Inside Job: Applying Traffic Analysis to Measure Tor from Within. In: NDSS (2018)
15. Johnson, A., et al.: Users Get Routed: Traffic Correlation on Tor by Realistic Adversaries. In: ACM CCS (2013)
16. Johnson, A., et al.: Avoiding The Man on the Wire: Improving Tor's Security with Trust-Aware Path Selection. In: NDSS (2017)
17. Juen, J., et al.: Defending Tor from Network Adversaries: A Case Study of Network Path Prediction. In: PETS (2015)
18. Mitseva, A., et al.: The state of affairs in BGP security: A survey of attacks and defenses. *Computer Communications* **124** (2018)
19. Murdoch, S.J., Danezis, G.: Low-cost Traffic Analysis of Tor. In: IEEE S&P (2005)
20. Murdoch, S.J., Zieliński, P.: Sampled Traffic Analysis by Internet-exchange-level Adversaries. In: PETS (2007)
21. Nithyanand, R., et al.: Measuring and Mitigating AS-level Adversaries Against Tor. In: NDSS (2016)
22. Panchenko, A., et al.: Website fingerprinting at internet scale. In: NDSS (2016)
23. Sun, Y., et al.: RAPTOR: Routing Attacks on Privacy in Tor. In: USENIX Security (2015)
24. Sun, Y., et al.: Counter-RAPTOR: Safeguarding Tor Against Active Routing Attacks. In: IEEE S&P (2017)
25. Tan, H., et al.: Data-plane Defenses against Routing Attacks on Tor. In: PETS (2016)
26. Tschantz, M.C., et al.: SoK: Towards Grounding Censorship Circumvention in Empiricism. In: IEEE S&P (2016)
27. Wacek, C., et al.: An Empirical Evaluation of Relay Selection in Tor. In: NDSS (2013)
28. Wails, R., et al.: Tempest: Temporal Dynamics in Anonymity Systems. In: PETS (2018)
29. Wan, G., et al.: Guard Placement Attacks on Path Selection Algorithms for Tor. In: PETS (2019)