Who Carries Tor? Measuring Bandwidth-Weighted Transit Concentration

Isabelle Wang
Smith College
Northampton, Massachusetts, USA
ilwang@smith.edu

Shinyoung Cho
Smith College
Northampton, Massachusetts, USA
scho@smith.edu

Abstract

The Tor network remains vulnerable to the underlying inter-domain routing infrastructure. A small set of Autonomous Systems (ASes) act as critical intermediaries and may gain visibility into substantial fractions of inbound and outbound Tor traffic, creating potential chokepoints that could facilitate traffic correlation attacks.

To quantify this risk, we adapt the AS Hegemony metric and extend it by introducing a Symbolic ASN to represent the Tor network as a logical destination and incorporating consensus bandwidth weighting to reflect actual traffic patterns. Using this methodology, we conduct a longitudinal study spanning 2015–2025 and a country-level analysis of Tor's AS-level transit dependency. Our findings reveal persistent concentration in a small number of transit ASes, a dramatic surge of SURF (AS1103) to over 50% exit hegemony in 2024, and significant sensitivity to the distribution of relay capacity across ASes at the country level. We further propose a diversity-aware relay placement heuristic to guide future mitigation efforts.

Keywords

Tor networks, AS dependencies, transit concentration

1 Introduction

The security and privacy of the Tor network [3] rest on the assumption of a decentralized infrastructure. However, this logical decentralization is often undermined by the physical reality of inter-domain routing. The Internet is organized into Autonomous Systems (ASes)—independently operated networks managed by ISPs, universities, or companies—that exchange routing information using the Border Gateway Protocol (BGP) [14]. In practice, a small, persistent set of ASes act as critical intermediaries for a disproportionate share of global traffic, including Tor. These ASes gain significant visibility into both inbound and outbound Tor traffic, creating systemic security, privacy, and resilience risks.

This concentration has direct security consequences. AS-level centralization enables traffic correlation attacks, where an adversary observing traffic at both ends of a Tor circuit can deanonymize users with high statistical confidence [4, 9, 18]. Beyond active adversaries, even a trustworthy AS poses operational risk: an AS carrying 30% of Tor's bandwidth-weighted traffic means that a single outage or BGP misconfiguration could render that share of the network unreachable. Concentration of relays within a few dominant ASes further exposes the network to large-scale censorship and targeted disruption.

Despite the growth of Tor in relay count and total capacity, this AS-level concentration has persisted. Current relay deployment is largely organic and volunteer-driven—a testament to the community's commitment to privacy—yet this naturally leads to geographic and topological clusters, as volunteers tend to deploy in well-connected networks that happen to be highly hegemonic. These findings suggest that the distribution of relay capacity across ASes can substantially influence transit concentration and motivate the development of tools and guidelines that account for AS-level routing diversity.

In this paper, we quantify this risk by adapting the AS Hegemony metric [5, 8]—a measure of how dependent paths are on a given AS—to the Tor network. We extend it with two contributions: a Symbolic ASN representing the Tor network as a logical destination, and consensus bandwidth weighting to reflect actual path selection probabilities. We apply this methodology to conduct a decade-long study (2015–2025) and a country-level analysis [7] of Tor's AS-level transit dependency. Our findings reveal persistent concentration, dramatic short-term shifts driven by relay deployment, and structural risks from cascading upstream dependencies. We conclude by proposing a diversity-aware relay placement heuristic as a tool for exploring potential reductions in transit concentration.

2 Methodology

To quantify Tor's dependence on transit providers, we adapt AS Hegemony, a robust alternative to standard centrality measures [5]. AS Hegemony measures the extent to which observed routing paths depend on a given AS. We further introduce two Tor-specific extensions: a Symbolic ASN representing the Tor network as a logical destination and consensus bandwidth weighting to reflect path selection probabilities.

2.1 From Destination-Based to Service-Specific Centrality

Betweenness Centrality (BC) measures the fraction of paths traversing a node. In BGP settings, however, BC is highly sensitive to viewpoint placement, often inflating the scores of ASes near collectors. AS Hegemony ($\mathcal{H}$) mitigates this bias using a trimmed mean, discarding the top and bottom α fraction of BC values across viewpoints to reduce sampling bias in partial BGP views [5].

We extend AS Hegemony to a service-specific setting by introducing a Symbolic ASN (AS_{Tor}), a logical destination representing the Tor network. We append AS_{Tor} to BGP paths terminating in ASes hosting Tor relays, thereby capturing paths toward Tor's

Free and Open Communications on the Internet 2026(2), 37–41

entry or exit surface. This symbolic destination enables the computation of *Local Hegemony* with respect to the Tor network, formally, $\mathcal{H}(v, AS_{Tor})$, defined as the fraction of bandwidth-weighted paths to AS_{Tor} that transit AS v, aggregated across all viewpoints. The resulting metric quantifies transit dependence on the Tor network as a service rather than on individual relay-hosting ASes.

2.2 Bandwidth-Weighted Betweenness Centrality

Tor clients select relays proportionally to their consensus weight (B_w), a bandwidth-derived score assigned by directory authorities. As a result, high-capacity relays carry a disproportionate share of traffic. To reflect this behavior, we introduce B_w as a path weight in the BC calculation. We retain the IP-space weighting of $\sigma_{jw}(v)$ from prior work [5], preventing ASes that fragment address space into many small prefixes from receiving inflated scores. Together, these extensions account for both Tor's bandwidth-driven relay selection and BGP deaggregation bias.

$$BC_{Tor}(v) = \frac{1}{n \cdot \sum_{w \in V_{Tor}} \sigma_w \cdot B_w} \sum_{j \in V} \sum_{w \in V_{Tor}} \sigma_{jw}(v) \cdot B_w \quad (1)$$

where V_{Tor} is the set of ASes hosting Tor relays, and $\sigma_{jw}(v)$ is the IP-space-weighted reachability of AS v from viewpoint j through AS w [5]. σ_w is the total address space of AS w, used to normalize the denominator. B_w is the aggregate consensus bandwidth of relays hosted in AS w [11], and n is the number of BGP observation points.

By weighting each path by B_w, a transit AS's hegemony score reflects not only how frequently it appears on routes toward Tor relays, but how much Tor traffic those relays carry: transit ASes on paths toward high-bandwidth relays contribute proportionally more than those serving low-capacity nodes.

2.3 Calculation of Tor Hegemony

We refer to our adaptation of Local Hegemony for the Tor network as Tor Hegemony. The final Tor Hegemony score is calculated by applying the trimmed mean aggregation to these weighted values:

$$\mathcal{H}_{Tor}(v, \alpha) = \frac{1}{n - (2\lfloor \alpha n \rfloor)} \sum_{j=\lfloor \alpha n \rfloor + 1}^{n - \lfloor \alpha n \rfloor} BC_{Tor}^{(j)} \quad (2)$$

Following the methodology established in previous research [5], we arrange the $BC_{Tor}^{(j)}$ values from S viewpoints in ascending order and discard the top and bottom α ratio to eliminate BGP sampling bias. The resulting score $\mathcal{H}_{Tor}(v)$ ranges from 0 to 1, representing the fraction of Tor's bandwidth-weighted traffic that transits AS v.

2.4 Calculation of Country-Based Tor Hegemony

Tor Hegemony aggregates viewpoints from across the Internet, reflecting Tor's aggregate transit dependencies. Country-based Tor Hegemony restricts computation to viewpoints within a given country, revealing the transit ASes most central to that country's connectivity to Tor. The country-based metric enables comparisons of transit concentration across countries and helps identify country-specific routing dependencies and potential points of exposure.

We group viewpoints by their geolocated country rather than treating all viewpoints as a single global set, computing a separate hegemony score for each country using only its viewpoints [7]. Each viewpoint's IP address is geolocated using MaxMind [10]; viewpoints that cannot be geolocated are excluded. The hegemony score is computed using Equation (2) with n set to the per-country viewpoint count rather than the total global viewpoint count.

3 Our Measurement

We measure AS-level transit dependencies for both entry and exit relay surfaces of the Tor network, analyzing paths directed toward Tor rather than originating from it. Our dataset spans January 2015 to December 2025, combining BGP routing data with Tor consensus snapshots collected on matching dates.

BGP Data. We collect BGP routing tables from four route collectors: rrc00 and rrc10 (RIPE RIS [15]), and route-views2 and route-views.linx (RouteViews [19]). We analyze paths directed toward the Tor network—from clients to entry relays and from destination services to exit relays—and collect separate snapshots for each direction. To capture a ten-year longitudinal view, we take a single snapshot on January 1st of each year from 2015 to 2024. For 2025, we collect a snapshot on the first of each month to track recent dynamics at finer granularity.

Tor Consensus Data. We collect Tor consensus data [11] on the same dates as the BGP snapshots. From each consensus, we extract relay IP addresses along with their entry or exit classification and consensus bandwidth weights, which serve as the B_w values in our hegemony computation.

IP-to-ASN Mapping. We map relay IP addresses to ASNs using the CAIDA prefix-to-AS dataset [2], then aggregate per-relay bandwidth to the AS level to obtain B_w for each relay-hosting AS.

Country Geolocation. For country-based Tor Hegemony, we geolocate BGP viewpoint IPs using MaxMind [10]. All viewpoint IPs were successfully geolocated. Since this analysis targets current routing structure rather than long-term trends, geolocation is applied only to the 2025 monthly snapshots.

4 Findings

We analyze transit AS Hegemony across two dimensions: longitudinal trends over 2015–2025 period and country-level concentration in 2025. In both cases, we find that hegemony is shaped by a combination of backbone infrastructure shifts and relay placement.

4.1 Tor Hegemony (2015-2025)

Entry Relays. Figure 1(a) shows that entry relay transit dependency is primarily shaped by changes in the Tier 1 provider landscape, including corporate mergers and routing restructuring, rather than Tor relay placement. AS3356 (Level 3, later CenturyLink/Lumen) dominated entry hegemony from 2015 to 2017, peaking at 0.15 in 2015, indicating that nearly one in six bandwidth-weighted paths toward entry relays transited a single AS.

AS3356's hegemony declined sharply after 2017, coinciding with its acquisition by CenturyLink. During 2020–2022, Telia (AS1299) became the dominant transit AS, reaching a peak hegemony of approximately 0.13 before declining after its 2022 restructuring as Arelion. Since 2023, no single AS has exceeded 0.08, indicating

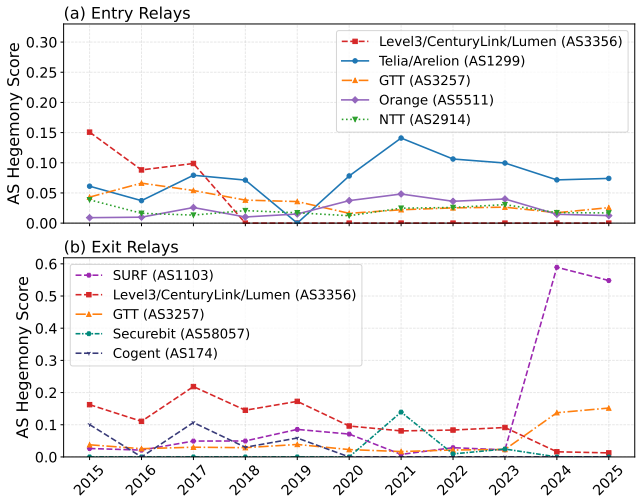

Figure 1: Top-5 transit AS Hegemony toward Tor entry (a) and exit (b) relays, measured annually from 2015 to 2025.

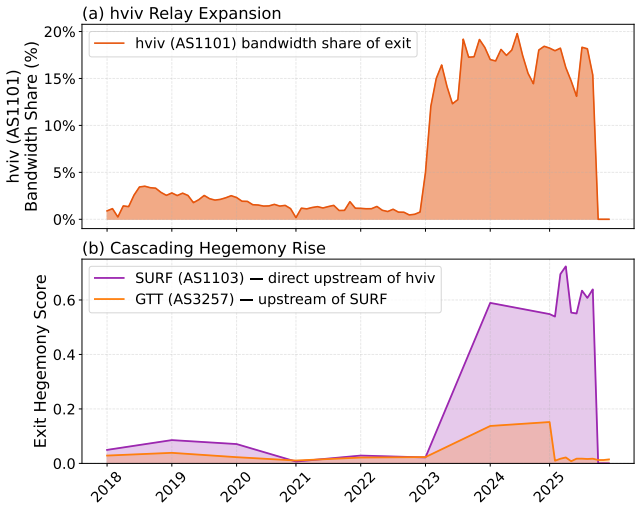

Figure 2: (a) Bandwidth share of exit relays operated by hviv (AS1101), and (b) cascading exit hegemony rise in SURF (AS1103), the direct upstream of hviv, and GTT (AS3257), upstream of SURF. The rapid relay expansion beginning in mid-2023 drove AS1103's hegemony to 0.59 by 2024.

a more distributed transit structure. Overall, these shifts suggest that changes in backbone connectivity and routing policies—largely outside the control of Tor relay operators—can substantially reshape Tor's transit dependencies.

Exit Relays. Figure 1(b) shows that exit relay transit dependency can be strongly influenced by relay deployment, in addition to backbone infrastructure. The most striking trend is the rise of SURF (AS1103) from near zero to 0.59 in 2024, meaning that nearly three in five bandwidth-weighted paths toward exit relays transited a single AS. As shown in Figure 2, this increase coincides with the rapid expansion of exit relays operated by Hart voor Internetvrijheid (hviv), a Dutch Internet freedom NGO, which grew from fewer than 20 relays to over 300 high-bandwidth relays hosted in AS1101 between 2023 and 2024. This concentration also propagated upstream, raising the hegemony of SURF's primary transit provider, GTT (AS3257), to 0.15. Unlike the entry-relay trends, these results suggest that concentrating relay capacity within a single AS can substantially reshape exit-side transit dependencies.

4.2 Country-Level Hegemony (2025)

We analyze Singapore and Malaysia as illustrative case studies, as both are classified as Partly Free in the Freedom House Internet Freedom Index [6].

Singapore. Figure 3 shows the monthly hegemony scores for exit relays geolocated in Singapore throughout 2025. From January through November, Viewqwest Pte Ltd (AS38001, SG) and NewMedia Express Pte Ltd (AS18106, SG) each maintain a consistent hegemony score of approximately 0.33, indicating that roughly one-third of paths to Singaporean exit relays transit each AS.

In December 2025, Viewqwest (AS38001) rises to 0.50 while NewMedia Express (AS18106) drops to 0.00, and SG.GS (AS24482) disappears from the top rankings entirely. This transition is driven primarily by changes in relay placement and consensus bandwidth rather than backbone infrastructure. AS18106's score collapses because the relays hosted by Shinjiru Technology (AS45839, MY)

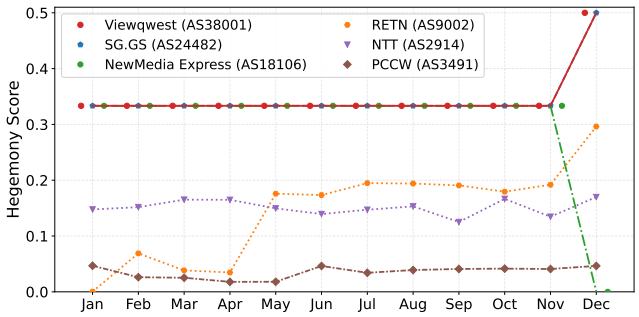

Figure 3: Top 6 exit AS Hegemony scores for Singapore-geolocated Tor relays by month in 2025.

experienced a substantial reduction in aggregate consensus bandwidth. At the same time, relays hosted in Orange Espagne (AS12479, ES) increased their aggregate consensus bandwidth to over 800,000 KB/s and expanded from four to five relays, shifting a substantial share of Singaporean exit traffic through Viewqwest (AS38001) as its upstream transit provider. These results illustrate that changes in relay capacity in one AS can propagate through the routing hierarchy, substantially reshaping country-level transit dependencies in geographically distant networks.

Malaysia. Figure 4 shows the monthly hegemony scores for entry relays geolocated in Malaysia throughout 2025. From August to September, the hegemony scores of the three dominant transit ASes drop simultaneously: PCCW Global (AS3491, HK) falls from 0.907 to 0.459, NTT (AS2914, US) from 0.278 to 0.142, and Liberty Global (AS6830, EU) from 0.091 to 0.045.

These three ASes serve as transit providers for multiple Tor-hosting ASes in Malaysia, including British Telecommunications

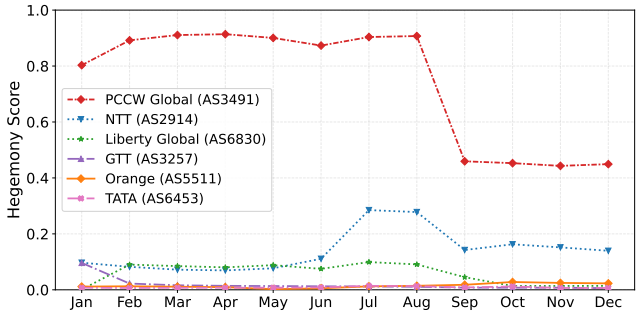

Figure 4: Top 6 entry AS Hegemony scores for Malaysia-geolocated Tor relays by month in 2025.

(AS6871, UK), FranTech Solutions (AS53667, US), and velia.net (AS30083, DE). The simultaneous decline in transit hegemony follows reductions in relay capacity across these transit ASes. FranTech Solutions (AS53667) experiences a sharp decline in hegemony between July and August as its relay count drops from 175 to 81.

British Telecommunications (AS6871) similarly declines between May and July as its relay count falls from two to one and its aggregate consensus bandwidth halves. Although velia.net (AS30083) maintains a single relay throughout the period, two reductions in consensus bandwidth coincide with further decreases in transit hegemony. Together, these results show that country-level transit dependencies can be highly sensitive to coordinated changes in relay capacity across multiple Tor-hosting ASes.

5 Toward Diversity-Aware Relay Deployment

Our findings demonstrate that transit hegemony is sensitive to localized changes in relay placement: bandwidth growth in a single AS in Spain can shift country-level transit concentration in Singapore. This sensitivity suggests that strategic relay placement can serve as a powerful tool for decentralization. We therefore propose a placement strategy that, given a deployment budget B and a set of candidate ASes, minimizes the maximum transit hegemony score.

We model this as a flow-based optimization problem over the Tor network graph $G = (\mathbb{V}, \mathbb{E})$. Here, "flow" represents the bandwidth-weighted probability of a path transiting a given AS. The objective is to minimize $\max_{v \in \mathbb{V}} F_v$ (the peak hegemony across all transit ASes) subject to budget and capacity constraints. Because optimal placement is computationally intensive, we propose a greedy heuristic: iteratively identifying the transit AS with the highest hegemony and deploying a relay in a candidate origin AS that provides the greatest marginal reduction in that peak.

While a full evaluation is deferred to future work, the dramatic surge in SURF (AS1103) hegemony suggests that even modest diversity-aware deployments could reduce reliance on dominant transit providers. Relay deployment, however, is only one factor influencing transit concentration. Because Tor's bandwidth-weighted relay selection favors high-capacity relays, it may also affect the distribution of traffic across transit providers. Understanding the tradeoff between performance, load balancing, and transit diversity remains an important direction for future work.

6 Related Work

Prior work on AS-level threats to Tor focuses primarily on traffic correlation attacks [4, 12, 18] and AS-aware path selection [1, 13, 16, 17]. These works highlight the importance of AS diversity in reducing exposure to routing-level adversaries. Our work is more closely related to AS Hegemony measurement: Fontugne et al. [5] introduce AS Hegemony as a betweenness-centrality metric for transit dependence, and IHR [8] provides ongoing longitudinal hegemony monitoring. Huffaker et al. [7] extend hegemony to country-level rankings, which informs our country-based analysis. To our knowledge, ours is the first study to apply bandwidth-weighted AS Hegemony specifically to the Tor network and track its evolution over a decade.

7 Discussion

Our analysis relies on BGP-derived paths from a finite set of observation points. These paths may differ from actual data-plane routes, omit routing asymmetry, and may not reflect the geographic or network distribution of Tor clients. Consequently, the reported hegemony scores should be interpreted as estimates of transit concentration rather than exact user-visible exposure.

In addition, our bandwidth-weighted metric reflects Tor's relay-selection probabilities at a given consensus but does not model guard persistence, whereby clients retain entry guards for extended periods. Consequently, it characterizes structural routing concentration rather than the temporal evolution of individual clients' entry-side exposure.

At the methodology level, we append a symbolic Tor ASN to paths originating in the Tor network, which can inflate the scores of origin ASes. While we filter known origin ASes, relay-hosting ASes that also provide transit and ASes serving as BGP viewpoints may still receive elevated hegemony scores. Future work could better isolate pure transit contributions by excluding path endpoints and separating relay-hosting and transit roles.

8 Conclusion

Our analysis reveals that transit AS Hegemony in Tor is driven by a complex interplay of backbone infrastructure shifts and relay placement. Because bandwidth weights propagate from Tor-hosting ASes through their transit paths, localized changes in relay capacity—such as the December shift in Singapore exit relays—can cascade into measurable hegemony changes at the transit level. This effect is most visible in the surge of SURF (AS1103), which peaked at over 50% exit hegemony in 2024, creating a significant structural dependency.

We present the first bandwidth-weighted AS Hegemony analysis of the Tor network spanning 2015 to 2025. Our findings show persistent concentration in a small number of transit ASes, country-level sensitivity to relay placement, and structural risks from cascading upstream dependencies. To help mitigate these risks, we propose a diversity-aware relay placement heuristic that minimizes transit AS hegemony. More broadly, Tor Hegemony provides a quantitative framework for evaluating how relay placement and Internet routing jointly shape transit concentration, enabling more informed deployment strategies for improving AS-level routing diversity.

Acknowledgments

This work was initiated with support from the Open Technology Fund. We thank the anonymous reviewers for their constructive feedback, which substantially improved this paper. We are grateful to Romain Fontugne for his generous guidance on the original AS Hegemony codebase. We also thank the Smith College undergraduate students who contributed to this project: Tasha Adler, Elaine Demetrion, and Sophie Crane.

References

[1] Masoud Akhoondi, Curtis Yu, and Harsha V Madhyastha. 2012. LASTor: A low-latency AS-aware Tor client. In *2012 IEEE Symposium on Security and Privacy*. IEEE, 476–490.

[2] CAIDA. 2026. Routeviews Prefix-to-AS Mappings Dataset (pfx2as). https://www.caida.org/catalog/datasets/routeviews-prefix2as/. Accessed: 2026.

[3] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2004. Tor: The Second-Generation Onion Router. In *Proceedings of the 13th USENIX Security Symposium*. USENIX Association, 303–320.

[4] Nick Feamster and Roger Dingledine. 2004. Location diversity in anonymity networks. In *Proceedings of the 2004 ACM workshop on Privacy in the electronic society*. 66–76.

[5] Romain Fontugne, Anant Shah, and Emile Aben. 2018. The Thin Bridges of AS Connectivity: Measuring dependency using AS hegemony. In *Proceedings of the Passive and Active Measurement Conference (PAM)*.

[6] Freedom House. 2026. Freedom on the Net. https://freedomhouse.org/report/freedom-net. Accessed: 2026.

[7] Bradley Huffaker, Romain Fontugne, Alexander Marder, and kc claffy. 2023. On the Importance of Being an AS: An Approach to Country-Level AS Rankings. In *Proceedings of the 2023 ACM on Internet Measurement Conference*. 52–65.

[8] IIJ Research Lab. 2026. Internet Health Report. https://ihr.iijlab.net/. Accessed: 2026.

[9] Aaron Johnson, Chris Wacek, Rob Jansen, Micah Sherr, and Paul Syverson. 2013. Users Get Routed: Traffic Correlation on Tor by Realistic Adversaries. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security* (Berlin, Germany) *(CCS '13)*. ACM, New York, NY, USA, 337–348. https://doi.org/10.1145/2508859.2516651

[10] MaxMind. 2026. GeoLite2 Free Geolocation Data. https://dev.maxmind.com/geoip/geolite2-free-geolocation-data. Accessed: 2026.

[11] Tor Metrics. 2026. Tor Project: Anonymity Online. https://metrics.torproject.org. Accessed: 2026.

[12] Milad Nasr, Alireza Bahramali, and Amir Houmansadr. 2018. Deepcorr: Strong flow correlation attacks on Tor using deep learning. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. 1962–1976.

[13] Rishab Nithyanand, Rachee Singh, Shinyoung Cho, and Phillipa Gill. 2016. Holding all the ASes: Identifying and Circumventing the Pitfalls of AS-aware Tor Client Design. *arXiv preprint arXiv:1605.03596* (2016).

[14] Y. Rekhter, T. Li, and S. Hares. 2006. *A Border Gateway Protocol 4 (BGP-4)*. RFC 4271. RFC Editor. https://www.rfc-editor.org/rfc/rfc4271

[15] RIPE NCC. 2026. RIPE Routing Information Service (RIS). https://ris.ripe.net. Accessed: 2026.

[16] Florentin Rochet, Ryan Wails, Aaron Johnson, Prateek Mittal, and Olivier Pereira. 2020. CLAPS: Client-Location-Aware Path Selection in Tor. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. 17–34.

[17] Oleksii Starov, Rishab Nithyanand, Adva Zair, Phillipa Gill, and Michael Schapira. 2016. Measuring and mitigating AS-level adversaries against Tor. In *Network and Distributed System Security (NDSS)*.

[18] Yixin Sun, Anne Edmundson, Laurent Vanbever, Oscar Li, Jennifer Rexford, Mung Chiang, and Prateek Mittal. 2015. RAPTOR: Routing Attacks on Privacy in Tor. In *Proceedings of the 24th USENIX Security Symposium (USENIX Security)*. 271–286.

[19] University of Oregon Route Views Project. 2026. Route Views Project. http://www.routeviews.org. Accessed: 2026.