Invisible Internet Project and Spatial Restrictions: A Systemic Vulnerability

KEMI AKANBI, School of Information Technology, University of Cincinnati, USA
TAREK KRAYEM, College of Sciences and Technologies, University of Bordeaux, France
SIDDIQUE ABUBAKR MUNTAKA, School of Information Technology, University of Cincinnati, USA
JACQUES BOU ABDO, School of Information Technology, University of Cincinnati, USA

The Invisible Internet Project (I2P) is a peer-to-peer overlay network that assumes complete abstraction of the Internet's spatial distribution. I2P's freedom from spatial restrictions is a prerequisite to all the anonymity guarantees it offers. All I2P literature has implicitly considered this assumption to be true, but if it turned otherwise, I2P's anonymity guarantees will be significantly degraded, and much of the previous literature must be reconsidered. In this paper, we test this assumption and demonstrate that a router's view of the I2P network depends on its geographic location. Using measurements collected from multiple geographic locations, we demonstrate that a router's view of the network is not uniform but instead strongly shaped by location. Statistical tests (Chi-square and Kolmogorov–Smirnov) confirm that peer distributions differ significantly across locations, challenging the assumption of spatial independence. Furthermore, latency analysis reveals that routers in regions with weaker infrastructure may experience higher delays, introducing measurable performance disparities. In other words, our observation suggest that this assumption is false.

CCS Concepts: • **Networks**; • **Network properties**; • **Network privacy and anonymity**;

Additional Key Words and Phrases: Invisible Internet Project, Darkweb, Empirical Evaluation, Network measurement

ACM Reference Format:
Kemi Akanbi, Tarek Krayem, Siddique Abubakr Muntaka, and Jacques Bou Abdo. 2025. Invisible Internet Project and Spatial Restrictions: A Systemic Vulnerability. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation emai (Conference acronym 'XX)*. ACM, New York, NY, USA, 21 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

The World Wide Web, commonly referred to as the web, is a complex system [30] that comprises three primary components: the surface, the deep web, and the dark web [31]. The surface web is easily accessible and indexed by search engines such as Yahoo and Google. In contrast, the deep web, also known as the invisible web, contains approximately 95% of the web content [26]. Within the deep web lies the dark web [30]. The primary goal of the dark Web is to ensure anonymity to circumvent censorship and protect the privacy of users [39]. Anonymous communication in the dark web is facilitated by a decentralized network [24] and has become a significant aspect of the contemporary digital environment.

Authors' Contact Information: Kemi Akanbi, akanbikr@mail.uc.edu, School of Information Technology, University of Cincinnati, Cincinnati, Ohio, USA; Tarek Krayem, tarek.krayem@etu.u-bordeaux.fr, College of Sciences and Technologies, University of Bordeaux, Bordeaux, Nouvelle-Aquitaine, France; Siddique Abubakr Muntaka, muntaksr@mail.uc.edu, School of Information Technology, University of Cincinnati, Cincinnati, Ohio, USA; Jacques Bou Abdo, bouabdjs@ucmail.uc.edu, School of Information Technology, University of Cincinnati, Cincinnati, Ohio, USA.

Manuscript submitted to ACM

Anonymity networks, with their layered encryption and unique routing methods, protect a user's communications and identity, making it difficult to trace who is communicating with whom [39]. In addition to hiding the sender and receiver's identities, anonymity on the dark net also masks the connection between them [24]. Several anonymity networks exist, including Invisible Internet Project (I2P), The Onion Router (Tor) which are currently the most popular choices for anonymous communication[16, 24]. Tor has a very wide userbase, with approximately 2 million daily users [3], but the number of routers supporting the traffic is around 9 thousand [2]. This makes Tor a small anonymity network in comparison with I2P, which has a routerbase of more than 50 thousand [18, 28]. Additionally, Tor was designed with reduced anonymity guarantees as a compromise for better performance [7]. This makes I2P the anonymity network of choice for high-importance traffic and the focus of this study.

The Invisible Internet Project (I2P) is a fully decentralized, peer-to-peer overlay network [15, 16], that operates under the assumption that the underlying Internet's spatial distribution has been abstracted away. This principle of spatial independence implies that a user's network location does not affect their participation or anonymity. To date, most I2P literature has treated this assumption as valid. However, if spatial dependencies do exist, they could alter the strength of anonymity assurances. To better understand I2P's current state, structure, and performance, a further empirical study is required. In this work, we examine the network through the lens of geolocation to assess whether spatial biases are present. Identifying such constraints is critical, as they may introduce uneven anonymity protections, leaving peers in less connected or underrepresented regions at greater risk

The remainder of the paper is organized as follows: Section 2 introduces Related works. Section 3 describes the Network description. Section 4 presents a comparison of views at a continental level. Section 5: measures the Network Latency. Section 6: Present the results of the experiments and provide the discussion and synthesis of the findings. Section 7 summarizes the conclusion and future directions.

2 Related Works

In this section, we review the literature on peer-to-peer overlay networks, focusing on their topology, classification, and applications. Then, conclude by describing the Invisible Internet Project (I2P), an anonymous overlay network.

2.1 P2P Overlay Network

A peer-to-peer (P2P) network is a decentralized network architecture where all connected devices, called peers, directly exchange information and services with one another and can act as both clients (consuming services) and servers (providing services), without relying on a central server [6]. In this architecture, users participate via software agents running on their devices. This allows peers to communicate directly with their known neighbors to share requests and responses [6], forming the basis for many modern distributed systems. A P2P overlay network is a virtual layer built on top of the Internet's physical infrastructure as shown in Figure 1. Overlays decouple logical communication from physical topology, enabling peers to interact directly through logical links rather than strictly relying on physical connections [21]. This abstraction enables the underlying internet infrastructure to support overlays, facilitating advanced functions such as distributed search, efficient routing, and content delivery [33]. Overlay networks are broadly classified into two types: structured and unstructured. In structured overlay, strict mathematical rules are adopted for their topology, often using Distributed Hash Tables (DHTs), such as Chord or Kademlia, to map content to specific nodes deterministically, making structured overlays ideal for applications that require exact matches and fast searches due to their efficient lookup and predictable routing [21]. Unstructured overlays organize peers randomly without a rigid topology. Content discovery relies on flooding queries or random walks, which increases redundancy but makes

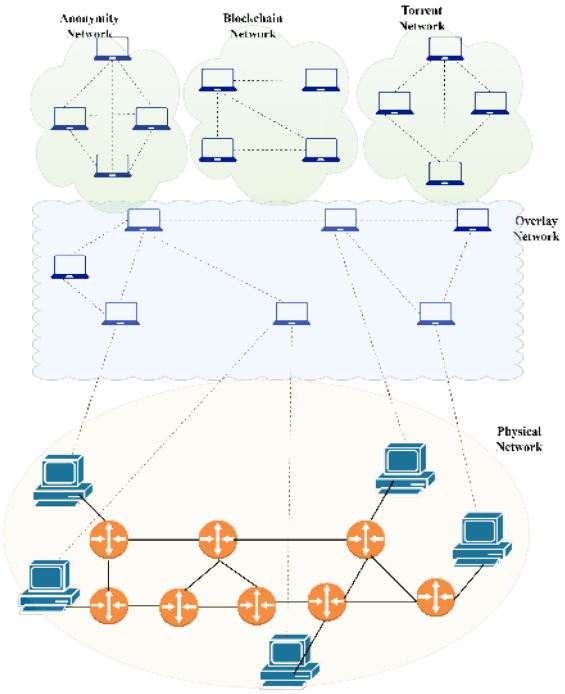

Fig. 1. P2P Overlay Network

the system easier to build and more robust against churn. While unstructured systems excel at finding popular content, they are inefficient for rare content searches and face scalability limitations due to high lookup costs [21, 43]. Han explains that Hybrid overlays are overlays that combine the strengths of both models, where the super-peers take on more responsibility for indexing and routing. At the same time, ordinary peers connect to them for discovery and communication. This approach offers the search efficiency of structured systems while retaining the flexibility and fault tolerance of unstructured designs [12].

2.2 P2P Overlay Network Application

P2P overlays support a wide range of applications, including file sharing, blockchain, and anonymity networks.

2.2.1 File sharing: File Sharing remains the most prominent application. Protocols like BitTorrent, Gnutella, KaZaA have enabled users to exchange large volumes of digital content, including movies, software, and books [9, 27, 43]. Gnutella was among the first large-scale decentralized protocols, allowing peers to search and share content without the need for central servers, supporting a network that could adapt to the fluidity introduced as users joined or left [9]. Building on these ideas, BitTorrent emerged as the most widely adopted P2P application, revolutionizing content distribution by breaking files into smaller pieces that could be downloaded simultaneously from multiple peers, dramatically boosting both scalability and download speed[9, 43]. Later, KaZaA introduced a hybrid model that combined the strengths of structured and unstructured overlays. By relying on a two-tier hierarchy of ordinary nodes

and super-nodes, the super-nodes carried more of the network's workload, enabling KaZaA to achieve better efficiency while still maintaining the flexibility of decentralized design [9, 27].

2.2.2 Blockchain: Blockchain systems build directly on the principles of peer-to-peer (P2P) overlays, where communication and data exchange happen without the need for central authorities. In this model, the overlay acts as a virtual layer on top of the physical Internet, connecting nodes through logical links that determine how transactions and blocks are propagated [13]. The overlay structure is central to blockchain's promise of decentralization and robustness; however, it also means that performance and resilience depend heavily on how well the overlay is designed. If connectivity is poor or uneven, the entire system's ability to distribute data and maintain consensus can be compromised. In systems such as Bitcoin, the earliest and most widely recognized cryptocurrency, peers (miners) are connected in an overlay that propagates blocks and transactions. Unlike traditional P2P file-sharing systems, blockchain overlays ensure not only fast content propagation but also global consensus [22]. Blockchain systems are built directly on the principles of peer-to-peer (P2P) overlays, where communication and data exchange occur without the need for central authorities.

2.2.3 Anonymity Networks: Anonymity networks, such as Tor, I2P, and Freenet, are overlay networks that utilize layered encryption and randomized routing paths to provide anonymity for peers on the network[15, 24]. Tor primarily focuses on anonymous access to the Internet, while I2P is designed for anonymous communications and services on the dark web to protect peers' identities and circumvent censorship [16, 36, 37]. Freenet emphasized strong anonymity for whistleblowers and the distribution of sensitive content. The design of these overlays directly influences scalability, latency, and resilience to attacks [13, 22, 32], allowing for the support of diverse use cases, ranging from entertainment-driven file sharing to critical systems such as blockchains and anonymity networks. Each application domain leverages the structured and unstructured trade-off differently, highlighting the importance of overlay design in determining the outcomes of scalability, efficiency, and anonymity.

2.3 Overview of I2P

I2P is a message-centric, low-latency anonymous overlay network [39] that leverages the Kademlia-based DHT to facilitate anonymous communication between participants within the I2P ecosystem [25]. I2P employs garlic routing, a technique that combines numerous messages and their delivery instructions into a single, encrypted message using the recipient's key [41]. Figure 2 illustrates the garlic routing method employed in I2P to facilitate anonymous communication between peers. The routers register as nodes, making their bandwidth available for communication on the network [26]. I2P transmits messages through unidirectional tunnels. It incorporates bandwidth sharing, with a default setting of 80% of the user's bandwidth allocated to the network [38]. To enhance anonymity, the router information and services accessed by users are isolated from each other [25]. Below is a description of routers on I2P, including how tunneling occurs and the selection mechanism for choosing a router to participate in a tunnel.

2.3.1 Routers. Routers are a vital component of the 12P network. The I2P framework is centered around the I2P router, responsible for tasks such as managing peer statistics, handling encryption and decryption, and constructing tunnels [14]. Participants running an instance of the I2P routing software become peers or nodes in the network [40]. Routers on the I2P network have cryptographic identities with distinct hash values used for identification instead of IP addresses [16, 39]. A cryptographic identity is a hash value uniquely associated with a router, created at 12p software installation, and remains unchanged for its lifetime [16]. I2P classifies routers basically into three tiers based on their speed and capacity [4, 45]:

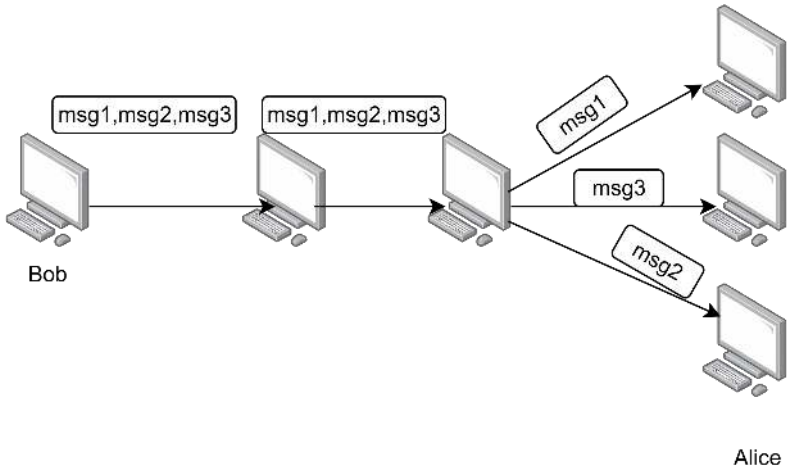

Fig. 2. Garlic Routing

- High Capacity: These routers can handle more tunnels than the average, offering reliable backbone traffic handling.
- Fast: Despite being high-capacity, these routers boast exceptional speed.
- Standard: These are the other routers available on the network.

2.3.2 Tunnels. Tunneling is a core concept in the I2P network, where each node or user deploys an I2P router to form hop paths among themselves [40]. Tunnels are categorized by transmission direction—outbound and inbound—and by usage into exploratory and client tunnels [4]. Exploratory tunnels are utilized for internal administrative tasks, network maintenance, service, and peer discovery, while client tunnels facilitate communication within the network [25]. Tunnels allow users to communicate anonymously by securely relaying each other's traffic, creating encrypted pathways through which data travels, bouncing between different users' devices, making it extremely difficult to track the data's origin or route [38]. To maintain anonymity, I2P tunnels have a short lifespan of ten minutes [8]. They must be recreated to mitigate the risk of de-anonymization from traffic analysis. I2P applications depend on the anonymizing tunnels created by the router to ensure privacy protection [14]. The number of hops in a tunnel can vary, depending on the user's configuration and the type of tunnel being created [38]. Increasing the number of hops enhances anonymity but can negatively impact performance [40]. However, the I2P development team recommends three or more hops to mitigate susceptibility to attack [19]. Figure 3 shows the inbound and outbound tunnels I2P users Bob and Alice used for anonymous communication.

2.3.3 Router Selection and Reputation Mechanism. Router selection in P2P overlay networks is not left to chance. Instead, it is shaped by mechanisms that assess the reliability, performance, and past behavior of participating peers. A common approach is the use of reputation-based systems, which encourage cooperation, reduce the risks of free-riding, malicious activity, or unreliable participation. By sharing information about a peer's trustworthiness, reputation mechanisms provide a foundation for more stable and resilient decentralized networks.

Reputation systems generally take two forms: local trust models and global trust models [17]. Local models rely on direct interactions and, in some cases, neighbor feedback to estimate a peer's trustworthiness. The Global trust

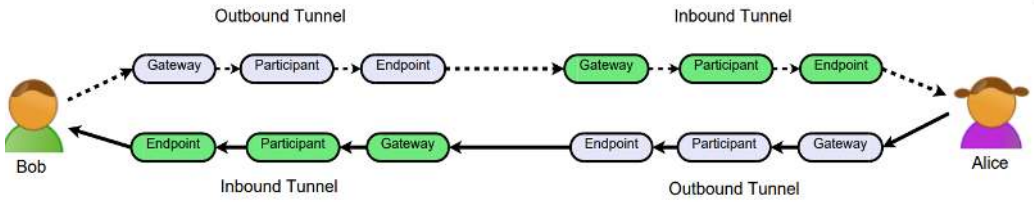

Fig. 3. I2P Tunnel [41]

model, by contrast, combines input from across the entire network to create a widely recognized reputation score [17]. Both designs aim to encourage honest participation and provide peers with guidance on whom to trust [11]. Unlike self-reported claims, reputation mechanisms depend on observable behavior. Interactions between peers are converted into numerical reputation scores that predict future reliability [20]. These scores are inherently subjective, dynamic, and context-dependent, meaning that assessments can vary across peers and change over time. In a decentralized environment, this adaptability is particularly valuable as peers must continuously update their trust judgments in response to evolving network behavior [20].

Within I2P, router selection relies on a reputation mechanism based on a performance-driven profiling system. Each peer tracks others' reliability through measures such as response times, tunnel success rates, and failure rates [1]. I2P takes a more dynamic approach by allowing each router to continuously observe and profile other peers based on their actual behavior, rather than just their claims. This profiling encompasses both direct interactions and indirect observations of other routers' performance in the network [25]. Router profiles are summarized into four main dimensions: speed, capacity, integration, and failure rate, all of which inform tunnel construction [1].

3 Network Description: Geographical Distribution

Mapping the I2P network is crucial for understanding the ecosystem, as previous studies have shown a limited amount of research on the network at the physical level. The study in [10] highlights that anonymity networks rely heavily on location diversity to defend against adversaries who can observe substantial portions of the network. Limited empirical and descriptive study on I2P exists. Peiping [25] identified Russia as the highest contributor at 33%, followed by America with 19.4%. The authors in [41] agree that the I2P peers are most active in Russia, the United States, France, and Germany. During their experiment, Russia and the United States consistently had over 1,500 routers active at any given time, with only occasional deviations. Russia contributed an average of 40% of the total participants in the network view obtained by the researcher out of 140 countries [41], signifying the country's dominance on the I2P network. The analysis of router distribution by [45] reveals that the Russian Federation, the United States, France, and Germany are the leading countries in active participation within the network. The Russian Federation constitutes nearly one-third (approximately 33%) of all routers observed from around 130 countries. [16] agrees with prior literature and found the US, Russia, and Britain to have the highest number of peers.By selecting a random sample of 250 I2P nodes and plotting their locations on Google Maps, [24] indicates that North America and Europe have the highest user populations. In general, literature shows that consistently over the years, Russia has emerged as the country with the highest percentage of peers in the network except [16], which reported that the United States is home to 28 thousand peers, thus contributing to the majority of peers, proving that the United States has overtaken Russia in peer population as of 2018. The metrics from I2P developers [18] indicate that there are approximately 52,012 thousand peers in the

United States alone as of April 22, 2024. [28] in a recent study supports the findings of [16, 45] that the United States, Russia, and Germany are still the primary contributors to the network, with the United States taking the lead. This geographic description offers valuable insights into the global reach and regional prominence of countries within the I2P network, providing a good starting point for our study to understand how geographical location may impact the network's performance and subsequently introduce vulnerabilities.

3.1 Monitoring Experiment: Single Location Perspective

In this monitoring experiment, a t2.medium virtual machine (VM) was provisioned on AWS with a specification of 2 vCPUs and 4 GB of memory. The experiment took place over 24 hours, from June 9th to June 10th, 2025. The collected data includes peer profiles and tunnel data. During this experiment, our router joined the network as a temporary, legitimate participant. We took strict measures to ensure our activities were non-intrusive and did not cause any harm. The data collected focused exclusively on publicly available peer profiles and tunnel metrics, allowing us to establish ground truth about geolocation and its relationship to network participation. We want to explicitly state that at no point did we engage in traffic de-anonymization. Furthermore, our presence was limited to the duration necessary for the experiment, after which the router was removed from the network. The primary goal was to observe and analyze network dynamics without compromising the anonymity or security of any user. For uniformity and clarity, we henceforth refer to our deployed routers as "router" and the nodes observed by each router as "peers". In this

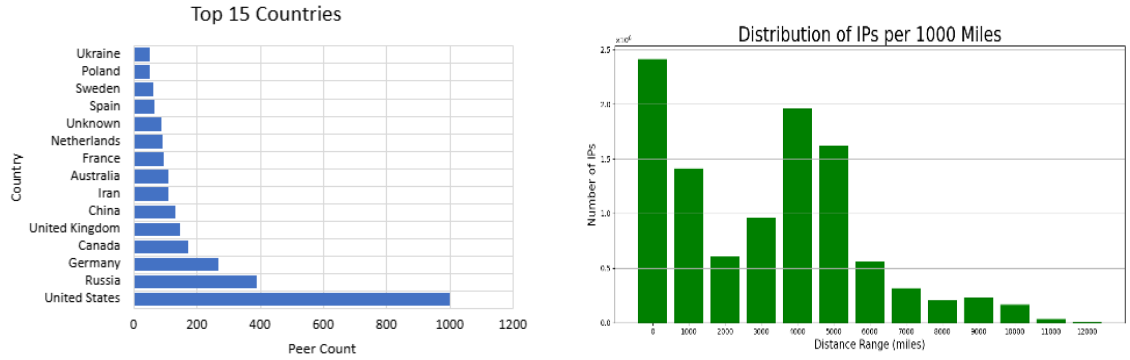

(a) Top 15 countries

(b) Distribution of IPs by geographic distance (in miles).

Fig. 4. (a) and (b) show global patterns of I2P network peer distribution.

part of our work, we focus on characterizing the Invisible Internet Project (I2P) from a single dominant location: the United States. Data collection began a day after the I2P router became active, ensuring a sufficient observation of the peer population and stable network behavior. Throughout the observation period, a total of 3,530 unique peers were identified, with 1,002 of these located within the United States. Notably, several peers did not disclose their country flags; consequently, we employed the Python library ipwhois to ascertain the geographic locations of these IP addresses. However, it is essential to note that the geographic locations of a small number of routers remain undetermined due to the absence of corresponding IP address information. Figure 4a shows the countries with the highest peer count observed in this experiment, with the United States, Russia, Germany, and Canada collectively representing 52% of the peers observed within the network. Additionally, countries such as Algeria and Slovakia each host two peers, while

countries including Venezuela, the Isle of Man, the Philippines, Kenya, Aruba, Monaco, Angola, Ecuador, Uruguay, Kyrgyzstan, and Morocco are represented by a single router each. Treating each observed peer as the "center of the world" allows for the categorization of peers into clusters based on their geographical distribution within intervals of 1,000 miles. Figure 4b provides a visual representation of these clusters, depicting the peers accessible to our router when delineating country classifications is disregarded.

3.2 Tunnel Participation and Selection Fairness

Tunnel selection fairness in I2P refers to how equitably the I2P network selects nodes for tunnel creation, indicating whether all eligible routers have a reasonably equal opportunity to be selected for use in tunnels. Tunnel selection fairness has a direct impact on the network's anonymity, resilience, and security. Centralization is likely to occur when there is unfair selection, a situation in which a small subset of high-performing nodes is disproportionately chosen, making the network more vulnerable to traffic correlation and targeted attacks. Ensuring fair participation among all eligible nodes enhances the robustness of anonymity guarantees, which strengthens the integrity of the network. Studying and improving fairness also provides valuable insights for optimizing peer selection algorithms and developing more secure, scalable anonymous communication systems. The tendency of peers to be selected for use in participatory tunnels was examined through the analysis of tunnel presence observed in our experiment. The analysis of tunnel presence was based on the number of peer profiles seen for each country, compared to the number of times a peer in that country is seen in the participating tunnel (t). For tunnel selection fairness evaluation, the first step is to divide the number of times a peer in the country is seen in the participating tunnel by 6, which is the total default hop length for incoming and outgoing communication in I2P. Next, we calculate the probability that a peer in the individual country will be selected for participation in the participating tunnel.

Countries with one or two routers in the collected data have less than a 0.04% chance of being selected for participation in the tunnel. In comparison, all countries with fewer than 30 routers have less than a 1% chance of being selected. For instance, Canada and the United States share borders, yet a peer in the US has a 30.47 percent chance of being selected. In comparison, their counterpart in Canada has a 4.85 percent chance of being selected, which is six times lower. We found that 75 out of 94 countries observed had less than a 1% chance of being chosen to be part of the participating tunnel. As shown in Figures 5a and 5b, the network is biased towards countries with a high number of peers, except in situations where the country's regime has implemented censorship or monitoring. Iran has 111, with 29 routers seen participating in the tunnels, resulting in a 1.07 percent chance of selection due to geolocation. [15] established in their study that China and Iran has censorship in place, validating the reason for low tunnel participation irrespective of the number of peers present in the location. It is noteworthy that only the peers in the United States and Russia have the highest chances of being selected to participate in a tunnel, which supports the observation that these two countries have consistently emerged as the locations with the most seen peer profiles. This imbalance has clear consequences for users in these areas, resulting in fewer mixzones and, consequently, smaller anonymity sets.The tunnel participation selection system reinforces the bias by favoring peers in well-populated areas, creating a situation where some users enjoy a stronger mixzone while others are left relatively exposed. On a mix-based anonymity network, every node acts as a mixzone, where messages are mixed together to dissociate the sender's identity from the traffic[34]. However, tunnel participation selection, such as the I2P selection mechanism, can disproportionately affect peers in regions with fewer routers and suboptimal infrastructure. As a result, these peers participates less in tunnels and consequently, faces diminished anonymity.

In the context of the client tunnel, Figure 6a depicts the peers identified by our router for the establishment of the tunnel, whereas Figure 6b delineates the selection fairness by country. Consistent with the findings observed in the participating tunnel, only the United States, Russia, and Germany achieved a utilization rate of 10% or higher for our router's client tunnel. China and Iran are indicated on the graph as not being selected. Despite China's peer count comparable to that of the United Kingdom or the Netherlands, the peers within China are selected for tunnel participation considerably less frequently. In contrast, countries such as Finland and Lithuania, which have a lower number of peers, were selected frequently for inclusion in our client tunnel. This monitoring experiment illustrates the substantial influence of router geographic location on its likelihood of being chosen as a participant in a peer's tunnel.

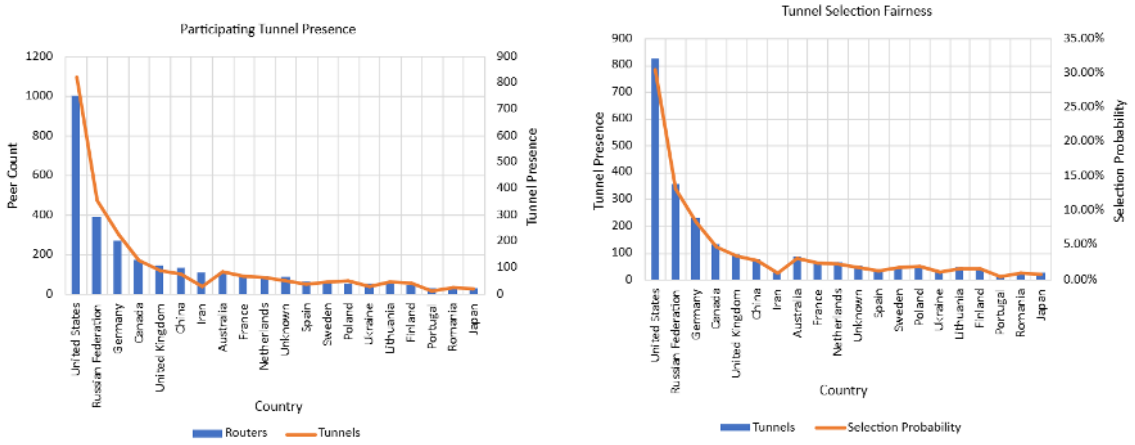

(a) Participating tunnel presence

(b) Tunnel selection fairness

Fig. 5. Comparison of tunnel presence and fairness across participating routers. (A) shows the presence distribution, while (B) illustrates fairness in tunnel selection.

3.3 Autonomous system Distribution

Autonomous System (AS) is a collection of networks managed by one or more network operators under a cohesive routing policy; these ASes are fundamental to enabling global Internet communication through diverse forms of interconnection [42]. An Autonomous System Number (ASN) primarily serves to identify the network of a specific Internet company [35]. Tozal [42] taxonomy categorizes ASes into four classes based on geometric terms: axial, pole, medial, and locus, to describe their role and position within the infrastructure. Locus ASes are peripheral networks, typically individual organizations or local ISPs, that lack customer ASes but have occasional internal peering. Medial ASes, such as national and regional ISPs, connect customer ASes to the internet through providers and often peer to manage costs. Axial or "core" ASes form the internet's backbone, lacking providers and peering directly with each other; they include major ISPs and infrastructure providers. Finally, Pole ASes are a subset of axial ASes forming the core's primary peering group [42]. Although IP addresses are frequently and dynamically rotated for residential users, peers often remain within the same AS or geographic region because newly assigned addresses typically come from the same subnet ranges [16]. This means that, while host-level identifiers may shift, the AS-level structure remains relatively stable, making it a crucial lens for analysis.

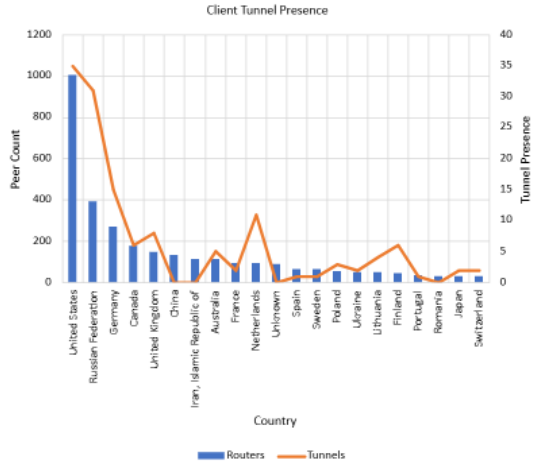

(a) Client Tunnel Presence

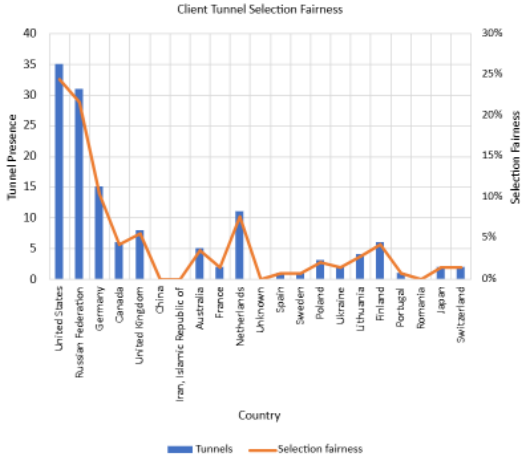

(b) Client Tunnel Selection Fairness

Fig. 6. Comparison of client tunnel presence (a) and tunnel selection fairness (b).

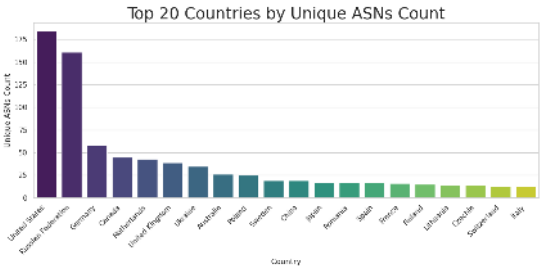

(a) Top 20 Countries by ASN count

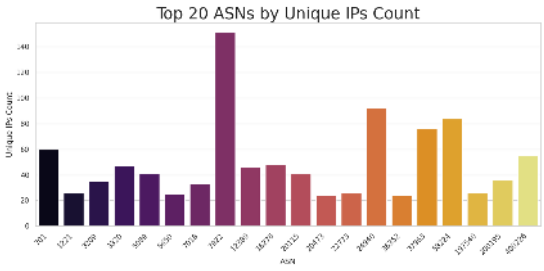

(b) Top 20 ASNs by Unique IPs count

Fig. 7. Geographic and ASN-level distribution of I2P peers. (a) Top 20 countries by ASN count. (b) Top 20 ASNs by unique IP count.

To achieve location independence in an anonymity network, it requires more than simply diversifying IP addresses; it also demands careful attention to the underlying Autonomous Systems (ASes) hosting those addresses, since geographic or AS-level concentration can undermine resilience, due to control by a small number of ASes which could potentially expose large portions of the network [10]. Our I2P measurements reveal this vulnerability. Figure 7a ranks the top 20 countries by ASN count, with the United States (184) and the Russian Federation (161) far ahead of others, followed by Germany (58), Canada (45), and the Netherlands (43). While the US and Russia exhibit high ASN diversity, this dominance concentrates network influence in a few ASes in some geographic regions.

Likewise, Figure 7b highlights even stronger centralization at the AS level. Although I2P peers are globally distributed, most activity clusters within a handful of ASNs. ASN 7922 (US) leads with 151 unique IPs, followed by ASN 24940 (Germany), ASN 58224 (Iran), ASN 37963 (China), and ASN 701 (US), each with 60–92 IPs. This pattern mirrors the concerns observed in other anonymity networks such as Mixmaster and Tor [10], where over-reliance on a small set of ASes, especially in specific countries, can weaken anonymity guarantees. Further analysis of the collected peer profiles reveals that only 13.3% of autonomous systems have more than five IP addresses, while the remaining 86.7% manage

fewer than five. This skewed distribution highlights the disproportionate influence of a minority of ASNs in the network, where a small group controls a considerably larger segment of IPs, thereby increasing the risk of surveillance and deanonymization. More ASN participation in underrepresented regions would help distribute trust, improve location diversity, and strengthen I2P's resistance to network-level surveillance or compromise.

Authors in [10] emphasizes that anonymity networks rely heavily on location diversity to defend against adversaries who can observe substantial portions of the network. Their analysis revealed that multiple nodes often reside within the same Autonomous System (AS) and that many network paths, including those to popular endpoints, traverse the same domain. This concentration at the AS level reduces location independence and increases vulnerability since control over a small set of ASes could enable large-scale observation or traffic correlation. These findings parallel patterns in our I2P measurements, where significant clustering of peers and IP addresses occurs within a few ASes, emphasizing the potential anonymity risks of AS-level centralization. Using each IP address observed in the first monitoring experiment as the center of the world, clusters of peers within 1000-mile distances are shown in Figure 3b. For instance, some parts of the USA, Canada, and Mexico are within 1,000 miles of each other. As such, the country-level dominance will be addressed, and a fairer selection will ensue, leveraging each country's ASN and ISP. Mapping peers at the AS level allows us to identify centralization risks, potential vulnerabilities, and the extent of geographic diversity. This is critically important in evaluating I2P, as it directly influences both resilience and anonymity guarantees.

4 Network Observation Deployment

This study hypothesizes that I2P adheres to spatial restrictions. To test the hypothesis, we conducted a monitoring experiment using Google Cloud. Six routers were deployed on the same day, in the same cloud, and with identical configurations. The Network Visibility Monitoring Architecture ensures all VMs are identically configured to eliminate any infrastructure bias. If location has no impact on the view of the network seen by the peer, all our deployed routers should get a similar view.

4.1 Network Visibility Comparison

Here, we conducted a monitoring experiment using Google Cloud, deploying one virtual machine (VM) in each of the six continents: the United States in North America, Germany in Europe, Australia in Australia, South Africa in Africa, and Chile in South America as shown in Figure 9a. This continental-level deployment allows us to validate whether peers, regardless of their geographic location, have a comparable view of the network and access to the fast set used for routing traffic within the I2P network. Figure 9b shows the United States router leading with an observation of 3,031 peers, surpassing all other countries, which suggests a dominant share of network participation. Australia follows with a view of 1,684 peers, while Latin America hosts 1,250 routers, indicating moderate involvement. Japan and South Africa have comparable totals, at 1,026 and 978, respectively. The router in the United States observed triple the number of peers compared to Japan, highlighting that the United States is a dominant hub within the I2P network. This trend is consistent across all deployed monitoring routers. In each country of observation, the United States is the most frequently observed peer, followed by the Russian Federation and Germany, denoting the most active participants in the network. To facilitate comparison, the countries with the highest peer counts, obtained from all our deployed routers, are plotted together in Figure 10a, providing a clear visualization of their relative participation levels on the network.

The total number of peers observed over 24 hours by the six routers deployed in this monitoring experiment was 8,440. The United States stands out with the highest peer concentration (2,373 peers), followed by the Russian Federation (902 peers) and Germany (905 peers). Mid-tier participation is observed in countries such as Canada, the United Kingdom, and

the Netherlands, each of which hosts several hundred peers. In contrast, Australia has 184 peers, while lower counts are recorded in several European countries, including Sweden (156), Spain (140), Poland (136), and Japan (116), all of which have fewer than 200 peers. The map in Figure 10b shows the total number of peers in observed in different countries by the monitoring architecture. The size of the bubbles indicates the number of peers, while darker colors represent higher values. Countries with fewer than 100 peers are not displayed on the map. This spatial visualization highlights the geographic centralization of I2P activity, revealing dominant hubs in North America and Europe, alongside peripheral regions with relatively modest peer presence. Cumulatively, we found that the United States accounts for 28.12%, the Russian Federation for 10.69%, and Germany for 10.72%. Prior studies, such as [16, 25, 28, 41, 45], support our finding that the US, Russia, and Germany are hubs of the network.

In a pairwise analysis, the Chi-square test of independence was applied separately to both the peer profile and fast set distributions of the observations of the different routers we deployed. The test revealed that the distributions within each set are statistically significantly different across countries. For the peer profile, we found that, as shown in Table 1, the views of various countries on the network differ in terms of the peers seen and the proportion of fast sets. Comparing the United States and South Africa, the test shows a p-value of 2.46×10^{-254} with 64 degrees of freedom. As this p-value is substantially less than the chosen significance level of 0.05, this result indicates a highly statistically significant difference, suggesting that the distribution of peer profiles seen by a user located in the US and South Africa differs significantly, providing support that the distribution is dependent on the country the peer is located in. This pattern does not only affect the USA and South Africa. It is seen in other countries, such as Chile and Germany, with a p-value of 2.41×10^{-107}, and Japan and Australia, with a p-value of 3.03×10^{-32}. Also, Table 1 displays the results of the chi-square test on the fast-set distribution available to our routers. In Asia and Australia, the p-value obtained was 4.60×10^{-9}; in Germany and Australia, it was 0.00017, indicating that the difference between the distribution of the fast-set available to the routers in our monitoring architecture is statistically significant. This demonstrates that peers located in different countries do not have equal access to fast sets for routing, highlighting spatial disparities within the I2P network.

The graphs in Figures 8a, 8b, 8c, 8d, 8e, and 8f show the distribution of peers and fast sets observed by each router in our monitoring system. A comparison of peer counts across the six locations reveals dominance in both the count of US-based peers and their "fast set." The United States, Russia, and Germany are consistently the most prominent peer profiles in their respective fast set categories. However, the router deployed in the United States exhibits a clear advantage, establishing connections with the highest number of US peers at 800, of which 80 are classified as fast set. In contrast, the Australian router observes a substantially lower total of 488 US peers, with only 26 of these belonging to the fast set. While Chile's router reports 379 US peers, it has the lowest count of fast set peers (18). We test the relationship between the peer profiles observed by each router and their corresponding fast sets. It is assumed that these two distributions are independent of each other. The chi-squared test result, with a p-value of 0.79 in Japan, indicates a 79% likelihood that the observed data could have occurred due to chance, which is not indicative of a statistically significant relationship. At a significance level of 0.05, the exception found was a dependent relationship between the available fast set and peer profiles in Chile and South Africa, with p-values of 0.0003 and 0.01, respectively. These two countries, with a dependent relation between the peer profile and fast set, have a lower number of peers, suggesting that router performance will be affected, as the available fast set is directly related to the number of available routers. For an anonymous network, the implication of the difference in peer distribution on anonymity is enormous, as it affects the density of their mixzones and thus provides a long-term passive observer with the ability to associate traffic with the source or destination for de-anonymization. The p-value obtained from the chi-square tests indicates that I2P

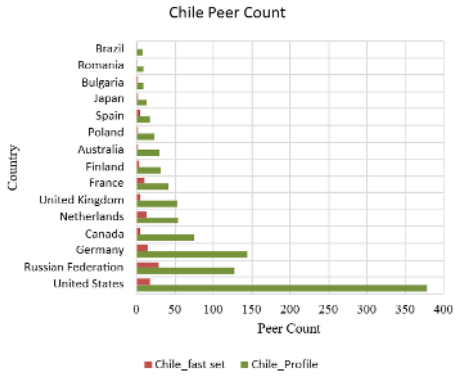

(a) Chile Peer Count

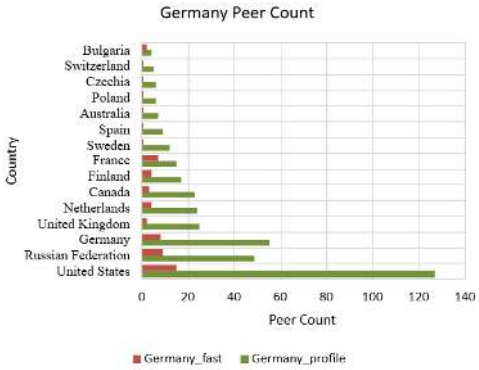

(b) Germany Router Peer Count

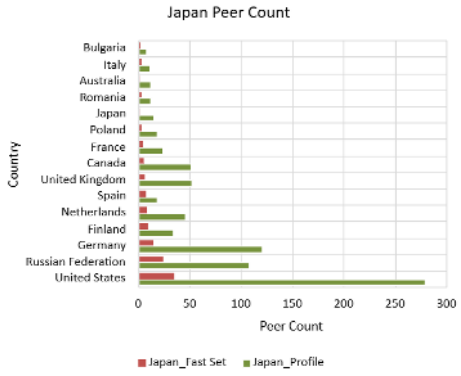

(c) Japan Peer Count

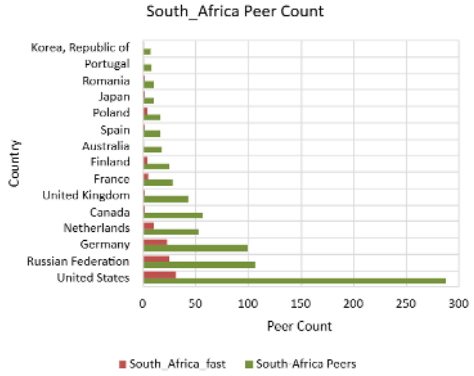

(d) South Africa Peer Count

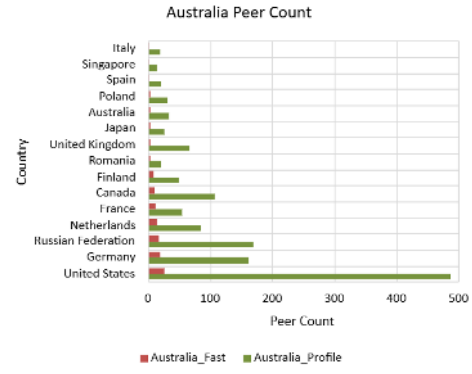

(e) Australia Peer Count

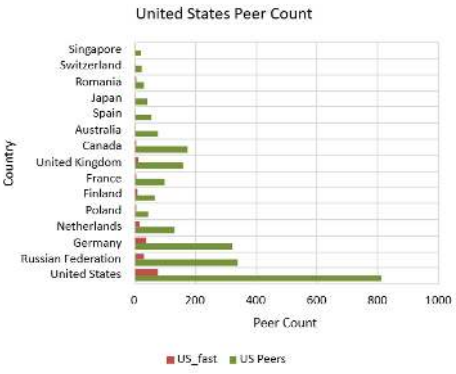

(f) United States Peer Count

Fig. 8. Peer counts observed by routers in different countries. (a) Chile, (b) Germany, (c) Japan, (d) South Africa, (e) Australia, (f) United States.

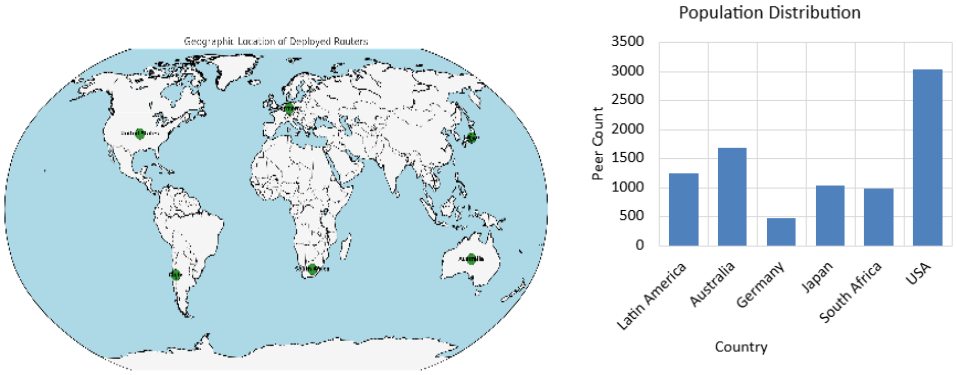

(a) Location of 6 routers on a map.

(b) Population distribution by location

Fig. 9. Deployed router locations and Population distribution.

adheres to spatial restriction, as the routers do not observe the same distribution of peer profiles or fast sets used in routing their traffic, nor do peers have an equal probability of selection to participate in tunnels. Thus, we reject the null hypothesis and accept the alternative hypothesis that I2P is subject to spatial restrictions, as geolocation influences the network's view and latency experienced by the peers.

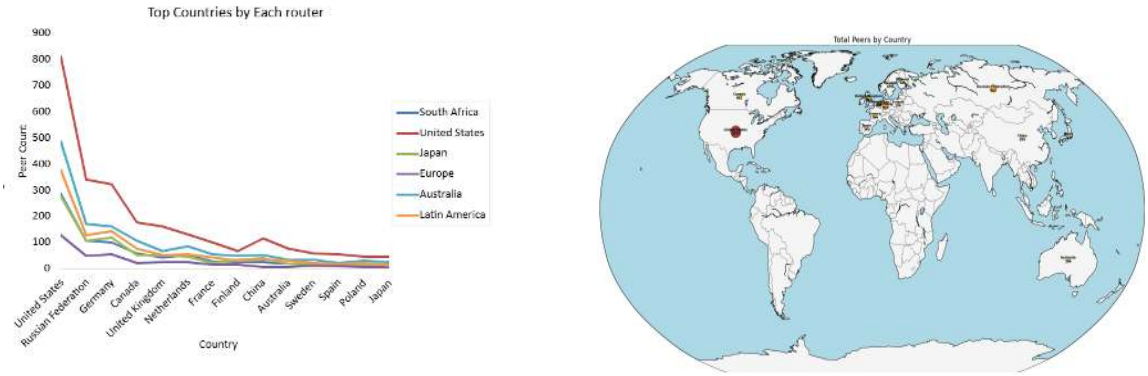

(a) Top countries by number of peers

(b) A global view of Peers.

Fig. 10. Peer distribution views: (a) Population distribution by location, (b) Comparative view of Top countries

5 Latency Measurement

Network latency refers to the total time delay for a message to traverse the network path, from its source to destination, including all intermediate devices.[23, 44]. Latency is a critical metric in modern network research as it has become a major basis for constructing de-anonymization schemes and a primary method for comparing network transmission rates [5]. Given this significance, our experiment deployed I2P routers in various countries to collect latency data for testing these hypotheses:

- Null Hypothesis: Geographical location has no impact on network latency in I2P.

Table 1. P-values for Peer Profile vs Fast Set Comparisons Across Regions

Peer Profile			Fast Set		
Country 1	Country 2	P-value	Country 1	Country 2	P-value
Chile	Germany	2.41×10^{-76}	Chile	Germany	0.00496
Japan	Australia	3.03×10^{-32}	Japan	Australia	4.60×10^{-9}
US	South Africa	2.46×10^{-254}	US	South Africa	3.96×10^{-7}
Germany	South Africa	2.28×10^{-111}	Germany	South Africa	5.06×10^{-16}
Germany	Japan	1.63×10^{-38}	Germany	Japan	0.00017
Chile	Australia	9.20×10^{-24}	Chile	Australia	0.00712

- Alternative Hypothesis: Geographical location impacts network latency in I2P.

5.1 Measurement Architecture and Deployment

To effectively evaluate the latency experienced by peers within the I2P network due to their geographical location, three Virtual Machines (VMs) labeled A, B, and C were provisioned by the Ohio Cyber Range Institute at the University of Cincinnati. Each VM operated on an identical Linux configuration, consisting of 8 GB RAM, 4 CPUs, and 80 GB of storage. VM A and VM B were configured to host the same eepsite, while VM A and VM C were tunneled through Germany. Each of the three VM hosts an I2P router, two of which host an eepsite. The eepsite content consisted of static text and images, simulating a lightweight web service for consistent latency measurements. Figure 11 illustrates the experimental architecture for latency measurement.

We developed a custom measurement script (see Appendix) tailored for this study. The script was deployed on VM C, which functioned as the client in our setup. Its primary task was to periodically get the hosted eepsites running on VM A and VM B, thereby simulating real-world client requests within the I2P environment. The script executed automated probes at five-minute intervals, issuing wget requests to the .b32.i2p addresses of the eepsites. It logged both round-trip times (RTTs) and HTTP response statuses to collect comprehensive latency data, ensuring consistent and repeatable measurements over the 24-hour observation window. By maintaining a steady probing rate, the system was able to capture temporal variations in latency, providing a reliable metric for performance comparison.

For a second iteration of the same experiment, the measurement framework was replicated with VM A and VM C tunneled through Nigeria, a country with an extremely low density of I2P routers. During this deployment process, it was observed that certain countries, particularly some in Africa, are associated with servers that I2P deems untrusted. Peers within these regions are therefore more likely to tunnel their outbound connections through alternative paths, introducing additional delay. This behavior can cause peers in these regions to appear slower and less reliable, thereby decreasing their likelihood of being selected for tunnel participation. As a result, reduced tunnel participation not only impacts performance but also diminishes the anonymity guarantees for these peers. The latency scripts were executed continuously over 24 hours to access the eepsites hosted on the servers.

To isolate the additional delay introduced by the extra tunneling, we measured the average latency when the U.S. router accessed its own hosted eepsite, which was then subtracted from the observed latency in Nigeria and Germany. Another iteration of the experiment was conducted to ensure that the extra tunneling was not responsible for any increased delay. The last setup involves three VMs: one hosted in the United States on the OCRI infrastructure, and two newly deployed routers located in South Africa, with no extra tunneling involved. In this configuration, one South

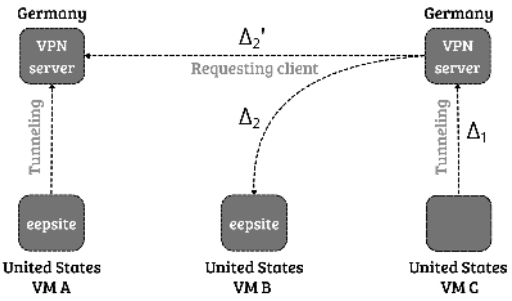

Fig. 11. Tunneling Configuration

African router functioned as the client, pinging the other two routers. This iterative approach enabled us to generate a dataset that reflects the actual influence of geographic location on latency in I2P with no infrastructure bias involved. By running the same client-side measurement script against different eepsite hosts, the overall setup allowed us to examine whether being located in a region with limited network infrastructure and few peers introduces additional delay.

5.2 Data Analysis

The data cleaning process began with the removal of values exceeding 10,000 milliseconds, which were identified as outliers. Subsequently, any data points that fell more than two standard deviations away from the mean were excluded. The resulting cleaned dataset was then utilized to calculate the mean latency for each router location. Throughout these experiments, the data analysis process remained consistent, incorporating the removal of outliers and the application of standardized filtering techniques. Figures 12a, b, and c illustrate the latency results across the three experimental setups.

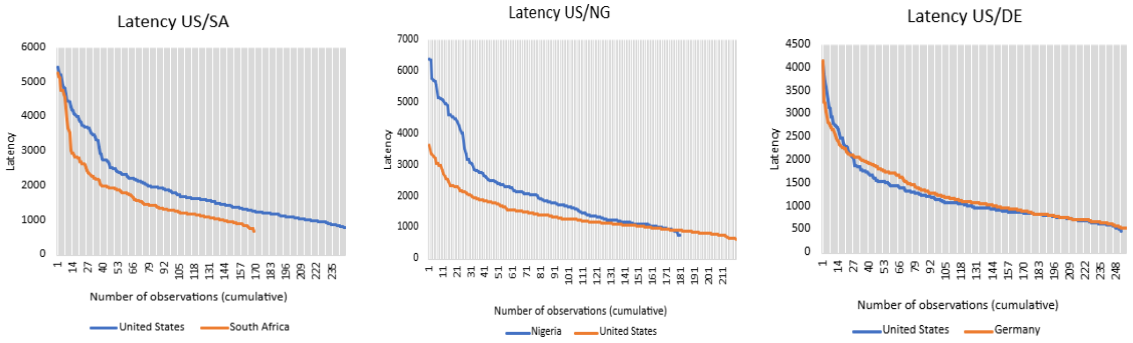

(a) Latency US/SA (b) Latency US/NG (c) Latency US/DE

Fig. 12. Comparison of latency across regions: (a) South America, (b) US/NG, (c) US/DE.

Statistically, we evaluate whether geographic location has a significant impact on network latency in the Invisible Internet Project (I2P), by applying the Kolmogorov–Smirnov (K–S) two-sample test. For the latency experienced between the United States (US) and Nigeria (NG), the Kolmogorov-Smirnov (KS) statistic achieved is 0.2778 with a p-value of 1.6088×10^{-6}. For Germany and the US, the KS statistic is 0.0968 with a p-value of 0.19607. The latency distribution

observed in South Africa (SA) and the US yields a KS statistic of 0.1246 with a p-value of 0.08. Hence, we found no statistically significant difference in the latency distribution between the US and Germany, as well as between the US and South Africa, in a pairwise analysis, shown in Figure 12. The discrepancy in the distribution for Nigeria and the US illustrates the possible influence of peer distribution in each geographical area. Nigeria has an extremely lower router population compared to the US, which explains the latency difference observed when communicating on the network from the respective locations.

The observed variations of router latency in our experiment should not occur if we assume that location does not influence latency. As a result, at a pre-established significance level α of 0.05, we reject the null hypothesis. This finding provides compelling statistical evidence that location has a significant influence on the latency experienced by a peer in the I2P. To avoid introducing bias in the latency measure experiments, the infrastructure, configuration, and analytical processes remained consistent. Therefore, we propose a peer selection process in which routers are chosen to participate in a tunnel based on their distance to the client rather than their country of origin. In summary, we accept the alternative hypothesis, as it has been demonstrated that the geographic location of the router impacts latency.

6 Experimentation Result and Discussion

I2P is P2P anonymous overlay network [4, 39, 41]. As a P2P overlay network, it is designed to utilize the internet's underlying infrastructure to support peer-to-peer connections within the network, disregarding geographical restrictions and imperfections. I2P has been described from a geographical perspective in prior literature; however, the impact of geolocation on network performance has not been studied to the best of our knowledge.

This study establishes that spatial independence, a core assumption of I2P overlay design, is not upheld. Our results demonstrate that I2P's peer-to-peer architecture is not evenly distributed globally. While the system is designed as a P2P overlay, in practice, the United States, Russian Federation, Germany, Canada, and the United Kingdom dominate the peer population. In our measurement experiment, we found that approximately 50 percent of observed networks using the 6 router monitoring architecture consist of peers located in the United States, the Russian Federation, and Germany, which supports the findings in the empirical study conducted in [28]. This clustering occurs within a limited number of Autonomous Systems (ASes) and can weaken location diversity. The presence of geographic and AS-level clusters spotlights that I2P does not fully achieve the spatial neutrality expected of an overlay network. Using Kolmogorov–Smirnov tests, we observed statistically significant differences in latency between the United States and Nigeria. This pattern suggests that regions with advanced infrastructure may experience low latency, while areas with fewer peers experience higher delays. These delays can serve as a distinguishing signal, allowing long-term observers to infer traffic origins or destinations, becoming a risk to anonymity.

The resultant topology, based on our findings that spatial restrictions apply to I2P, differs from the expected outcome if the restriction-free assumption holds. Although I2P's peer profiling system is designed to improve reliability by basing router selection on observed performance [1], the mechanism is not without bias. Routers located in regions with robust network infrastructure achieve higher reputation scores because they demonstrate lower latency and participate in more tunnels. This means that some routers in specific locations are chosen repeatedly, which further boosts their already favorable reputation scores. On the other hand, routers in regions with weaker infrastructure often struggle to build a strong profile, even when operating honestly. Factors such as limited bandwidth, higher latency, or inconsistent availability translate into lower performance ratings, making these routers less attractive to their peers for tunnel participation. Over time, this creates a self-reinforcing cycle where well-connected regions continue to dominate peer selection, while peers in low-population-density areas with weaker infrastructure remain sidelined

and are unable to improve their reliability within the network. Reputation-driven router selection thus achieves its goal of promoting trustworthy participation, but at the cost of reinforcing geographic clustering, which in turn affects anonymity guarantees.

Prior literature that has studied I2P has not considered the location effect on the network. Some others utilized simulation to describe the structure and resilience; however, simulation-based studies are not effective in describing I2P due to the limited data points available for use in the software and erroneous assumptions that every peer in the network experiences the same level of anonymity and security guarantee, irrespective of their location. Although [29] utilized live I2P data to simulate percolation for evaluating the network's resilience to attacks.

Taken together, our findings challenge the assumption that I2P is free from spatial restriction. In practice, its structure and anonymity guarantees are strongly shaped by geography and underlying network infrastructure. For anonymity networks, the study's findings have profound implications. Users outside significant hubs experience reduced mixzone density and latency-based leakage, which negatively impacts anonymity guarantees through passive traffic correlation attacks. For a network that continues to grow, these results highlight the need to redesign peer selection mechanisms to mitigate geographic bias. By addressing these structural imbalances, I2P may achieve fairer performance and stronger anonymity guarantees for users, regardless of their connection location.

7 Conclusion

This study challenges the long-standing assumption that I2P operates as a neutral, location-independent network. What we found instead is that the network is uneven, both in its structure and in its performance. The peer population is disproportionately concentrated in a handful of countries, with the United States, Russia, and Germany alone representing nearly half of all active peers. Statistical tests reveal that the peers a router can see depend heavily on its location, indicating that I2P is not as independent of geographical location as previously assumed. Similar pattern shows up in the latency experienced by peers. Our latency measurements reveal that where a router is located in the world may directly affects how quickly it can connect, creating unequal experiences for users. For a network that continues to expand and gain importance, this imbalance carries profound implications against anonymity guarantees. Future research will focus on measuring the extent to which anonymity is compromised due to the location of routers.

References

[1] [n. d.]. Peer Profiling and Selection in the I2P Anonymous Network. https://geti2p.net/en/docs/how/peer-selection. Accessed: 2025-08-19.

[2] [n. d.]. Servers; Tor Metrics. https://metrics.torproject.org/networksize.html. [Accessed 01-09-2025].

[3] [n. d.]. Tor Metrics. https://metrics.torproject.org/userstats-relay-table.html. [Accessed 01-09-2025].

[4] Jacques Bou Abdo and Liaquat Hossain. 2023. Modeling the Invisible Internet. In *International Conference on Complex Networks and Their Applications*. Springer, 359–370.

[5] Longy O. Anyanwu, Jared Keengwe, and Gladys Arome. 2010. Anonymity Leakage Reduction in Network Latency. In *Innovations and Advances in Computer Sciences and Engineering*, Tarek Sobh (Ed.). Springer Netherlands, Dordrecht, 561–565.

[6] Anirban Basu, Simon Fleming, James Stanier, Stephen Naicken, Ian Wakeman, and Vijay K. Gurbani. 2013. The state of peer-to-peer network simulators. *ACM Comput. Surv.* 45, 4, Article 46 (Aug. 2013), 25 pages. https://doi.org/10.1145/2501654.2501660

[7] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2004. Tor: The Second-Generation Onion Router. In *13th USENIX Security Symposium (USENIX Security 04)*. USENIX Association, San Diego, CA. https://www.usenix.org/conference/13th-usenix-security-symposium/tor-second-generation-onion-router

[8] Christoph Egger, Johannes Schlumberger, Christopher Kruegel, and Giovanni Vigna. 2013. Practical Attacks against the I2P Network. In *Research in Attacks, Intrusions, and Defenses*, Salvatore J. Stolfo, Angelos Stavrou, and Charles V. Wright (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 432–451.

[9] Yousef Fazea, Zainab S. Attarbash, Fathey Mohammed, and Ibrahim Abdullahi. 2021. Review on Unstructured Peer-to-Peer Overlay Network Applications. In *2021 International Conference of Technology, Science and Administration (ICTSA)*. 1–7. https://doi.org/10.1109/ICTSA52017.2021.

9406524

[10] Nick Feamster and Roger Dingledine. 2004. Location diversity in anonymity networks. In *Proceedings of the 2004 ACM Workshop on Privacy in the Electronic Society* (Washington DC, USA) *(WPES '04)*. Association for Computing Machinery, New York, NY, USA, 66–76. https://doi.org/10.1145/1029179.1029199

[11] Minaxi Gupta, Paul Judge, and Mostafa Ammar. 2003. A reputation system for peer-to-peer networks. In *Proceedings of the 13th International Workshop on Network and Operating Systems Support for Digital Audio and Video*. Association for Computing Machinery, New York, NY, USA, 144–152. https://doi.org/10.1145/776322.776346

[12] Hongmu Han, Jie He, and Cuihua Zuo. 2010. A hybrid P2P overlay network for high efficient search. In *2010 2nd IEEE International Conference on Information and Financial Engineering*. 241–245. https://doi.org/10.1109/ICIFE.2010.5609293

[13] Parikshit Hegde and Gustavo de Veciana. 2022. Performance and efficiency tradeoffs in blockchain overlay networks. In *Proceedings of the Twenty-Third International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing* (Seoul, Republic of Korea) *(MobiHoc '22)*. Association for Computing Machinery, New York, NY, USA, 221–230. https://doi.org/10.1145/3492866.3549730

[14] Michael Herrmann and Christian Grothoff. 2011. Privacy-Implications of Performance-Based Peer Selection by Onion-Routers: A Real-World Case Study Using I2P. In *Privacy Enhancing Technologies*, Simone Fischer-Hübner and Nicholas Hopper (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 155–174.

[15] Nguyen Phong Hoang, Sadie Doreen, and Michalis Polychronakis. 2019. Measuring {I2P} censorship at a global scale. In *9th USENIX Workshop on Free and Open Communications on the Internet (FOCI 19)*.

[16] Nguyen P. Hoang, Panagiotis Kintis, Manos Antonakakis, and Michalis Polychronakis. 2018-10-31. An Empirical Study of the I2P Anonymity Network and its Censorship Resistance. ACM.

[17] Jianli Hu, Quanyuan Wu, and Bin Zhou. 2008. Distributed and Effective Reputation Mechanism in P2P Systems. In *2008 International Conference on Computer Science and Software Engineering*, Vol. 3. 110–115. https://doi.org/10.1109/CSSE.2008.777

[18] I2P. [n. d.]. I2P Router Distribution. https://i2p-metrics.np-tokumei.net/router-distribution. Accessed: 2024-09-30.

[19] I2P. 2011. Tunnel Routing. https://geti2p.net/en/docs/how/tunnel-routing Accessed: 2024-09-14.

[20] Cornelius Ihle, Dennis Trautwein, Moritz Schubotz, Norman Meuschke, and Bela Gipp. 2023. Incentive Mechanisms in Peer-to-Peer Networks — A Systematic Literature Review. *ACM Comput. Surv.* 55, 14s, Article 308 (July 2023), 69 pages. https://doi.org/10.1145/3578581

[21] Nanami Imada and Kazunori Ueda. 2016. Peer-to-Peer Network System and Application Design on Multiple Virtual Networks. In *2016 19th International Conference on Network-Based Information Systems (NBiS)*. 298–302. https://doi.org/10.1109/NBiS.2016.66

[22] Suhan Jiang and Jie Wu. 2023. Approaching an Optimal Bitcoin Mining Overlay. *IEEE/ACM Trans. Netw.* 31, 5 (Jan. 2023), 2013–2026. https://doi.org/10.1109/TNET.2023.3235307

[23] Simran Preet Kaur, Manojit Ghose, Ananya Pathak, and Rutuja Patole. 2024. A survey on mapping and scheduling techniques for 3D Network-on-chip. *Journal of Systems Architecture* 147 (2024), 103064. https://doi.org/10.1016/j.sysarc.2024.103064

[24] Likun Liu, Hongli Zhang, Jiantao Shi, Xiangzhan Yu, and Haixiao Xu. 2019. I2P anonymous communication network measurement and analysis. In *International Conference on Smart Computing and Communication*. Springer, 105–115.

[25] Peipeng Liu, Lihong Wang, Qingfeng Tan, Quangang Li, Xuebin Wang, and Jinqiao Shi. 2014. Empirical Measurement and Analysis of I2P Routers. *Journal of Networks* 9, 9 (-09-07 2014).

[26] Roberto Magán-Carrión, Alberto Abellán-Galera, Gabriel Maciá-Fernández, and Pedro García-Teodoro. 2021. Unveiling the I2P web structure: A connectivity analysis. *Computer Networks* 194 (2021), 108158. https://doi.org/10.1016/j.comnet.2021.108158

[27] R. Matei, A. Iamnitchi, and P. Foster. 2002. Mapping the Gnutella network. *IEEE Internet Computing* 6, 1 (2002), 50–57. https://doi.org/10.1109/4236.978369

[28] Siddique Abubakr Muntaka, Jacques Bou Abdo, Kemi Akanbi, Sunkanmi Oluwadare, Faiza Hussein, Oliver Konyo, and Michael Asante. 2025. Mapping The Invisible Internet: Framework and Dataset. *arXiv preprint arXiv:2506.18159* (2025).

[29] Siddique Abubakr Muntaka and Jacques Bou Abdo. 2025. Resilience of the invisible internet project: A computational analysis. *Internet Technology Letters* 8, 5 (2025), e70119.

[30] Saiba Nazah, Shamsul Huda, Jemal Abawajy, and Mohammad Mehedi Hassan. 2020. Evolution of Dark Web Threat Analysis and Detection: A Systematic Approach. *IEEE Access* 8 (2020), 171796–171819. https://doi.org/10.1109/ACCESS.2020.3024198

[31] Fawn T. Ngo, Catherine Marcum, and Scott Belshaw. 2023. The Dark Web: What Is It, How to Access It, and Why We Need to Study It. *Journal of Contemporary Criminal Justice* 39, 2 (2023), 160.

[32] Aristodemos Paphitis, Nicolas Kourtellis, and Michael Sirivianos. 2023. Resilience of Blockchain Overlay Networks. In *Network and System Security*, Shujun Li, Mark Manulis, and Atsuko Miyaji (Eds.). Springer Nature Switzerland, Cham, 93–113.

[33] Dongyu Qiu and R. Srikant. 2004. Modeling and performance analysis of BitTorrent-like peer-to-peer networks *(SIGCOMM '04)*. Association for Computing Machinery, New York, NY, USA, 367–378. https://doi.org/10.1145/1015467.1015508

[34] Mahdi Rahimi. 2024. MALARIA: Management of Low-Latency Routing Impact on Mix Network Anonymity. In *2024 22nd International Symposium on Network Computing and Applications (NCA)*. 193–202. https://doi.org/10.1109/NCA61908.2024.00038

[35] Nematullo Rahmatov and Hoki Baek. 2024. Exploring scale-free networks: Survey on autonomous system dynamics and connectivity analysis. *Computer Networks* 248 (2024), 110454. https://doi.org/10.1016/j.comnet.2024.110454

[36] Abhiram S and Subhasri Duttagupta. 2023. Anonymous Decentralized Chat Engine Over Overlay Networks. In *2023 4th International Conference for Emerging Technology (INCET)*. 1–6. https://doi.org/10.1109/INCET57972.2023.10170194

[37] Roman Schlegel and Duncan S. Wong. 2012. Anonymous overlay network supporting authenticated routing. *Information Sciences* 210 (2012), 99–117. https://doi.org/10.1016/j.ins.2012.04.042

[38] Khalid Shahbar and A. Nur Zincir-Heywood. 2017. Effects of Shared Bandwidth on Anonymity of the I2P Network Users. In *2017 IEEE Security and Privacy Workshops (SPW)*. 235–240. https://doi.org/10.1109/SPW.2017.19

[39] Marco Simioni, Pavel Gladyshev, Babak Habibnia, and Paulo Roberto Nunes de Souza. 2021. Monitoring an anonymity network: Toward the deanonymization of hidden services. *Forensic Science International: Digital Investigation* 38 (2021), 301135. https://doi.org/10.1016/j.fsidi.2021.301135

[40] Juan Pablo Timpanaro, Thibault Cholez, Isabelle Chrisment, and Olivier Festor. 2015. Evaluation of the anonymous I2P network's design choices against performance and security. In *2015 International Conference on Information Systems Security and Privacy (ICISSP)*. 1–10.

[41] Juan Pablo Timpanaro, Isabelle Chrisment, and Olivier Festor. 2012. A Bird's Eye View on the I2P Anonymous File-Sharing Environment. In *Network and System Security*, Li Xu, Elisa Bertino, and Yi Mu (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 135–148.

[42] Mehmet Engin Tozal. 2016. The Internet: A system of interconnected autonomous systems. In *2016 Annual IEEE Systems Conference (SysCon)*. 1–8. https://doi.org/10.1109/SYSCON.2016.7490628

[43] Mehr un Nisa, Fasiha Ashraf, Ateeqa Naseer, and Shaukat Iqbal. 2019. Comparative Analysis of Unstructured P2P File Sharing Networks *(ICISDM '19)*. Association for Computing Machinery, New York, NY, USA, 148–153. https://doi.org/10.1145/3325917.3325952

[44] Wenye Wang, Yi Xu, and Mohit Khanna. 2011. A survey on the communication architectures in smart grid. *Computer Networks* 55, 15 (2011), 3604–3629. https://doi.org/10.1016/j.comnet.2011.07.010

[45] X. B. Wang, P. P. Liu, C. L. Li, and Q. F. Tan. 2013. Towards measurement on the I2P network, Vol. 427-429. 2223–2228.

A Latency Measure Script

```
#!/bin/bash
ROUTER_NAME_CURRENT="Router2"
EEPSITE_B32_1="oapkzxgcmvhclybqk4ifaimg6zqnwoq4ia4xntnyzjnxxxxxxxxx.b32.i2p" #USA
EEPSITE_B32_2="26tbnldvh4bvmbtp7ozjvzyljl7e6c6soilueokmpzyxxxxxxxxx.b32.i2p" #DE
LOG_FILE_1="$HOME/i2p_eepsite_US_rtt.csv"
LOG_FILE_2="$HOME/i2p_eepsite2_DE_rtt.csv"
TIMESTAMP=$(date +"%Y-%m-%d␣%H:%M:%S")

# Function to test eepsite and log results
test_eepsite() {
    local eepsite="$1"
    local log_file="$2"
    local site_name="$3"

    # Create log file header if it doesn't exist
    [ ! -f "$log_file" ] && echo "Timestamp,Pinging_Router,Target_Service,Target_B32,RTT_ms,Status" > "$log_file"

    # Measure RTT using date command (millisecond precision)
    local start_time=$(date +%s%3N)

    if wget -q -e use_proxy=on -e http_proxy=127.0.0.1:4444 --timeout=30 -O /dev/null "http://$eepsite/"; then
        local end_time=$(date +%s%3N)
        local rtt_ms=$((end_time - start_time))
        echo "$TIMESTAMP,$ROUTER_NAME_CURRENT,$site_name,$eepsite,$rtt_ms,SUCCESS" >> "$log_file"
        echo "$site_name␣SUCCESS:␣${rtt_ms}ms"
    else
        echo "$TIMESTAMP,$ROUTER_NAME_CURRENT,$site_name,$eepsite,N/A,FAILED" >> "$log_file"
        echo "$site_name␣FAILED"
    fi
}

# Run both tests in parallel using background processes
echo "Starting␣parallel␣eepsite␣tests..."
```

```
test_eepsite "$EEPSITE_B32_1" "$LOG_FILE_1" "Eepsite1" &
PID1=$!

test_eepsite "$EEPSITE_B32_2" "$LOG_FILE_2" "Eepsite2" &
PID2=$!

# Wait for both background processes to complete
wait $PID1
wait $PID2

echo "Both eepsite tests completed."
echo "Results saved to:"
echo "  Site 1: $LOG_FILE_1"
echo "  Site 2: $LOG_FILE_2"
```