On the Challenges of Geographical Avoidance for Tor

Katharina Kohls*, Kai Jansen*, David Rupprecht*, Thorsten Holz* and Christina Pöpper†

*Ruhr University Bochum, Germany
{katharina.kohls, kai.jansen-u16, david.rupprecht, thorsten.holz}@rub.de
†New York University Abu Dhabi, UAE
christina.poepper@nyu.edu

***Abstract*—Traffic-analysis attacks are a persisting threat for Tor users. When censors or law enforcement agencies try to identify users, they conduct traffic-confirmation attacks and monitor encrypted transmissions to extract metadata—in combination with routing attacks, these attacks become sufficiently powerful to de-anonymize users. While traffic-analysis attacks are hard to detect and expensive to counter in practice, *geographical avoidance* provides an option to reject circuits that might be routed through an untrusted area. Unfortunately, recently proposed solutions introduce severe security issues by imprudent design decisions.**

In this paper, we approach geographical avoidance starting from a thorough assessment of its *challenges*. These challenges serve as the foundation for the design of an empirical avoidance concept that considers actual transmission characteristics for justified decisions. Furthermore, we address the problems of untrusted or intransparent ground truth information that hinder a reliable assessment of circuits. Taking these features into account, we conduct an empirical simulation study and compare the performance of our novel avoidance concept with existing approaches. Our results show that we outperform existing systems by $22\,\%$ fewer rejected circuits, which reduces the collateral damage of overly restrictive avoidance decisions. In a second evaluation step, we extend our initial system concept and implement the prototype *TrilateraTor*. This prototype is the first to satisfy the requirements of a practical deployment, as it maintains Tor's original level of security, provides reasonable performance, and overcomes the fundamental security flaws of existing systems.

I. Introduction

Tor enables anonymous communication on the Internet as it allows to separate one's identity from what is being read, watched, bought, or shared. Such protection is put to good use in cases where this additional layer of anonymity helps journalists, whistleblowers, or everyday supporters of the digital freedom to stay safe under oppressive regimes and Internet censorship. At the same time, Tor's anonymity holds opportunities for illegal activities. Both cases serve as motivation for censorship authorities [54] as well as law enforcement agencies [1] to hinder the use of Tor and to monitor what is going on in the "dark parts" of the Internet.

We can circumvent blocked Tor access in different ways [17], but users never know if someone analyzes their traffic [18], [30], [36]. Low-cost countermeasures do not sufficiently protect metadata [13] and obfuscating traffic against correlation leads to per-packet delays [28]. However, we gain trust in a connection by *avoiding* paths through critical countries. Such circumvention becomes even more important since we know that, e. g., monitoring a circuit's middle relay is already sufficient to identify onion services [21]. Sophisticated path selection [4], [8] is a starting point for this approach, but systems tend to focus on performance features [48] rather than geographical characteristics.

DeTor [31], proposed by Li et al. in 2017, makes an attempt to provide *provable geographical avoidance* of untrusted countries. Provable avoidance means that it is impossible for an established Tor circuit to have traversed a forbidden area. This does not only apply to the avoidance of *relays* located in a specific country, but also considers the *Internet routing* between the client and server. DeTor uses an approach comparable to the principle of distance bounding: instead of depending on hardware solutions [5] or the modification of routing protocols [38], it uses the Round-Trip Time (RTT) of a connection and compares it to a theoretical lower bound for reaching the forbidden area. The lower bound is estimated using the geographical locations of relays in the circuit and utilizes the fact that transmissions through the Internet can never be faster than approximately ⅔ of the speed of light [31].

Unfortunately, several design flaws hinder DeTor from providing a convincing solution for geographical avoidance. (i) DeTor does not consider the diverse network infrastructure of Tor and the underlying network, e. g., it applies one static decision threshold for *all* circuits. Tor's skewed distribution of relays leads to various circuit lengths that cannot offer the same performance features for all users. Applying the same threshold even for varying connection characteristics leads to overly restrictive avoidance decisions. Furthermore, (ii) DeTor makes false assumptions on the available ground truth information. In particular, it assumes symmetric routes, miscalculates the distance within the lower bound detection mechanism, and ultimately accepts connections traversing forbidden areas. This contradicts the "provable" security guarantee for geographical avoidance. Furthermore, DeTor accepts external GeoIP information without any further verification and overlooks the chances of using false locations as the foundation for a decision. Finally, (iii) DeTor was designed

Network and Distributed Systems Security (NDSS) Symposium 2019
24-27 February 2019, San Diego, CA, USA
ISBN 1-891562-55-X
https://dx.doi.org/10.14722/ndss.2019.23402
www.ndss-symposium.org

without considering the constraints of real-world deployment. By sending timing probes through the *entire* circuit, the system reveals the connection endpoints even before we can be sure about the security of this circuit. This opens up new attack vectors instead of protecting against potential threats. We argue that all of these flaws are unnecessary and introduce strict security and performance issues that render the system hardly usable for an actual deployment.

In our work, we approach the general problem of geographical avoidance systematically and begin with a definition of its *challenges*. We introduce three classes of challenges that we identify as the general pitfalls of geographical avoidance, namely, the demanding characteristics of Tor's (i) *network diversity*, the lack of trusted (ii) *ground truth* information, and the requirements of a real-world (iii) *deployment*. Tackling these challenges, we propose a new timing-based avoidance system that overcomes design flaws of existing systems. We back up the theory of these challenges with a preliminary assessment of Tor's network infrastructure and the transmission characteristics of the underlying network. Our results show that the skewed distribution of Tor relays that we measured empirically not only leads to different levels of anonymity for users, but also affects the essential end-to-end timing of messages sent through the network. Ignoring this diversity means to over-simplify the decision process with consequences for either Tor's security or performance. We find that accepting external GeoIP information as ground truth for relay positions is error-prone and can impact geographical avoidance decisions. False locations would imply propagation speeds that exceed the speed of light and, with that, are provably wrong from a physical perspective. We verify the GeoIP information and identify false locations by applying this physical proof to improve the information through trilateration [19].

The assessment of challenges is our foundation to propose technical solutions and design a new, empirical avoidance concept. *Empirical avoidance* has two main benefits. First, it allows considering hop-individual transmission characteristics rather than one static threshold for different connections. Consequently, we can apply avoidance decisions concerning the various performance characteristics of Tor and step back from the collateral damage of overly restrictive decisions. Second, we derive the hop-individual timing estimates from distributed measurements of several reference points. This distributed approach adds another level of security, as it allows to represent single connections through empirical data that cannot be manipulated by an adversary prolonging messages [47]. In a first evaluation step, we analyze the performance of our novel avoidance concept and compare it to existing approaches.

In a second step, we introduce the prototype implementation *TrilateraTor* that is the first also to satisfy the requirements of a real-world *deployment*. *TrilateraTor* introduces a novel measurement technique that derives a circuit's end-to-end timing directly from the `NTor` handshake in Tor's circuit establishment procedure. As the establishment of several ready-to-use circuits is part of Tor's startup procedure, the use of *TrilateraTor* neither induces any delays through preliminary probing nor information leaks. Along with additional verification steps for untrusted ground truth information and the less restrictive empirical avoidance concept, *TrilateraTor* provides realistic answers to the challenges of geographical avoidance. We analyze the performance of our prototype implementation in another empirical simulation study, discuss the organizational steps of a practical deployment, and provide a detailed security analysis. Our contributions are as follows.

- **Challenges of Geographical Avoidance.** We assess the problem of geographical avoidance in Tor and identify three classes of challenges. These classes address the *diversity* of Tor's infrastructure and the underlying network, the lack of *ground truth* information, and the constraints arising from the real-world *deployment* of an avoidance system.
- **Preliminary Measurements.** We conduct an empirical evaluation of Tor's infrastructure to confirm the relevance of the above challenges. Our results show that a skewed relay distribution cannot provide the *same anonymity for all* and can limit the success of an avoidance system. Furthermore, we identify a false assumption that hinders DeTor [31] from providing *provable* avoidance.
- **Experimental Evaluation**. Starting from the assessment of the given network infrastructure, we introduce solutions for the set of challenges and compare their performances in an empirical simulation study. In a second step, we propose, implement, and evaluate *TrilateraTor*, our approach to take the constraints of a real-world deployment into account.

II. Background

Before we define fundamental challenges for avoidance and provide possible solutions for a system deployment, we introduce some background on the context of geographical avoidance. This background explains why traffic-analysis attacks harm the anonymity of Tor users, how routing attacks can render this situation a real-world threat, and introduces the technical background of trilateration. Furthermore, we discuss which attacker model we consider and briefly summarize the functionality of DeTor.

A. Motivation: Traffic-Analysis Attacks

Tor does not protect the metadata that is present in the TCP and IP packet headers or that can be derived from time relations. An adversary that is capable of monitoring transmissions can thus analyze patterns, such as the inter-arrival times of packets that result in individual features for different streams, and use the information to match streams, thus learning the relations between transmissions and users. Countermeasures against traffic analysis impose a high overhead as they require the obfuscation of metadata, which can only be achieved at the expense of performance impairments.

Passive attacks monitor transmissions to perform end-to-end matches between clients and servers [30], [43], [58]. They use *correlation metrics* that estimate the similarity of ingress and egress traffic or analyze the statistical characteristics of streams [11], [12], [26], [34] to make an educated guess on potential relations. Active attacks extend this by targeted manipulations of traffic, e. g., watermarking [49]–[51] or coding [32], [33], [42], [56] approaches may inject characteristic patterns at one end of the connection to increase the correlation of transmissions. Such attacks can reduce the required monitoring overhead, but tend to be less stealthy due to their active interference.

All correlation attacks have one aspect in common: A successful attack requires to monitor traffic through the involved nodes and will be more successful for adversaries that can put themselves in advantageous positions. Geographical avoidance helps to avoid such areas, but its circumvention is challenged by active routing attacks.

B. Amplifier: Routing Attacks

When a user at location A accesses a website hosted at location B, all messages must be routed through a set of nodes and organizational units to reach their target and to be sent back. The path of a circuit does not only depend on the choice of relays, but also on the routing conditions between the relays. Adversaries use routing attacks [9], [39] to manipulate such paths, forcing traffic through areas that are under adversarial control.

The Internet is divided into autonomous systems (AS), large organizational units that provide the service of forwarding messages to the desired destination. Routing between ASes is managed by the Border Gateway Protocol (BGP), which defines rule sets for ASes that agreed to connect to each other. Hence, BGP manages available and unavailable connections. As a result, the BGP routing tables define the paths a message will take when it is transmitted from A to B. As soon as an adversary manages to alter the routing tables, she is capable of forcing traffic through regions she is in control of.

AS-level adversaries can manipulate Tor routing [9] or perform BGP hijacks [39], [44] to force traffic to be routed to the adversarial prefix, e. g., to a Tor exit node that forwards large amounts of traffic. BGP hijacks blackhole traffic and allows for the observation of transmitted data, but this results in dropped connections that may reveal the adversarial activity. To overcome this, more sophisticated BGP interception attacks [7] force the adversarial AS to be at an intermediate point of the path. In this case, the connection is kept alive and allows the attacker to monitor the transmissions. BGP hijacks and interceptions empower traffic-analysis attacks, increase the probability of successfully de-anonymizing users in the network, or help learn the positions of critical nodes on the network. Nevertheless, routing changes influence the RTT of a transmission, and we can exploit this fact for the design of an avoidance system.

C. Technique: Trilateration

As a means to geographically localize Tor relays, we make use of trilateration. This technique is based upon measured distances to multiple known reference points. This widespread approach is used in, e. g., satellite navigation systems (such as GPS) or to determine the location of mobile phones in radio cells, and utilizes time or signal strength differences between reference points [19]. In the context of geographical avoidance, we utilize the round-trip times from one node to multiple reference servers to derive hop-individual time references.

The underlying theoretical model can be summarized as follows. The unknown location of a relay $\overrightarrow{R}$ is denoted by (x, y)[1]. As references, we use RTT measurements from n other nodes $\overrightarrow{S_1}, \overrightarrow{S_2}, \ldots, \overrightarrow{S_n}$ to $\overrightarrow{R}$. As a result, we obtain n RTTs $t_1, t_2, \ldots, t_n$ between known references and the respective relay. These timings are related to geographic distances considering a typical transmission speed v of up to $2/3$ of the speed of light. Having three or more geographic distances allows us to pinpoint the target by intersecting circles with radii corresponding to the estimated distances.

Due to noise in RTT measurements, the resulting circles do not intersect in a distinct point but rather mark a target area. To find the most likely position $\overrightarrow{R}$, we use a weighted root-mean-square error approach. This correction technique optimizes the result towards the minimal error with respect to all reference measurements:

$$\arg\min_{\overrightarrow{R}} \sqrt{\frac{\sum_{i=1}^{n}\left[\left(dist(\overrightarrow{R}, \overrightarrow{S_i}) - t_i \cdot v\right) \cdot \omega_i\right]^2}{n}}, \qquad (1)$$

where ω_i is a normalized weighting factor based on the distance to the reference. In particular, smaller RTTs are expected to be less affected by noise and consequently have higher weight in the error minimization process. The output of Equation 1 is the most likely relay position with minimal error.

D. Scope: Attacker Model

Throughout this work, we follow attacker models proposed in the literature and assume an AS-level adversary who can cover between 40 % (single malicious AS) and 85 % (state level adversary, colluding ASes) [40] of nodes in an attack. The attacker is capable of performing routing attacks, e. g., BGP hijacks [44] for redirecting user traffic, and traffic-analysis attacks [18], [21], [36] with the goal of learning sensitive information about Tor users. This may be achieved by having access to relays in the Tor network (by contributing as a volunteer relay operator), to layer three or four switches (network nodes that forward IP or TCP/UDP traffic), or by monitoring Internet exchange points (IXP). We assume that the adversary can manipulate time measurements, i. e., can hold back replies to increase the measured RTT of a connection. Note that a global adversary serves as a theoretical upper bound and can capture traces at arbitrary nodes; in this case, geographical avoidance is without effect.

E. System: DeTor

In 2017, Li et al. [31] proposed DeTor as a system to provide *provable geographical avoidance* in Tor. The core principle is comparing the measured RTT of a Tor circuit with a lower bound threshold that includes the trip to the forbidden area. If the measured RTT does *not* exceed the threshold, the respective forbidden area could not have been reached. In other words, the additional distance, and hence time, required to traverse the forbidden area is higher than the measured RTT would allow. This concept was originally introduced in the context of Alibi Routing [29], where single hops were checked and later extended to three-hop connections to fit the needs of Tor.

When estimating the lower bound, DeTor first calculates the minimal geographical distance D_{min} required for routing

[1] A position is defined by its latitude and longitude coordinates and neglects altitude information for the sake of simplification.

TABLE I. CHALLENGES OF GEOGRAPHICAL AVOIDANCE

Class	Challenge	Solution/Design Goal	Section
Network Diversity	*Relay Distribution*	Prevent Collateral Damage	§III-A1
	Connection Lengths		§III-A2
	Connection Failures		§III-A3
Ground Truth	*Relay Locations*	ICMP Reference, Update	§III-B1
	Asymmetry	Single Extension	§III-B2
	Transm. Characteristics	Individual Estimates	§III-B3
Deployment	*Performance*	Evaluation	§VI-D
	Information Sources	Circuit Establishment Timing Distributed Measurements	§V-C §VI-A
	Security	Security Analysis	§VI-C

through the forbidden area and, second, relates it to a transmission speed of $\frac{2}{3}c$, which is an estimation of the maximal speed of Internet connections. Considering an established Tor circuit, DeTor calculates the following threshold:

$$R_{min} = \frac{3}{2c} \cdot min \begin{cases} 2 \cdot D_{min}(c, F, e, m, x, s) \\ 2 \cdot D_{min}(c, e, F, m, x, s) \\ 2 \cdot D_{min}(c, e, m, F, x, s) \\ 2 \cdot D_{min}(c, e, m, x, F, s) \end{cases}, \qquad (2)$$

where c, s are client and server, e, m, x are entry, middle, and exit relays of Tor, and F is the forbidden area. To obtain geographical positions, DeTor performs a Geo IP lookup with the respective IP address of relays.

When deciding whether a Tor circuit avoided a forbidden area, a binary decision on the measured RTT R_{e2e} is performed against the calculated threshold R_{min}:

$$avoided = \begin{cases} 1, & (1+\delta) \cdot R_{e2e} < R_{min} \\ 0, & (1+\delta) \cdot R_{e2e} \geq R_{min} \end{cases}, \qquad (3)$$

with δ being a static overhead parameter in the range between $[0, 1]$ designated to compensate network inconsistencies and measurement noise. Whenever a measured RTT R_{e2e} is shorter than the DeTor-estimated threshold R_{min}, the circuit is proven to avoid a forbidden area.

III. CHALLENGES OF GEOGRAPHICAL AVOIDANCE

We begin our work with a systematic evaluation of the challenges of geographical avoidance, i. e., we identify fundamental influencing factors that define the performance *and* security of an avoidance system. We introduce three classes of challenges (see Table I), namely (i) *network diversity* that leads to heterogeneous transmission characteristics, (ii) a lack of *ground truth* information that complicates avoidance decisions, and (iii) the restrictions of a realistic *deployment*. In the following, we provide a detailed introduction of these three classes of challenges and complement our theoretical claims with the results of a preliminary measurement study that addresses the characteristics of Tor and the underlying network. Throughout this work, the set of challenges will guide our design of a new avoidance concept and later also dictate the requirements that a practical prototype implementation must satisfy.

TABLE II. CONSENSUS STATISTICS

		EU							NA	
		DE	**FR**	**NL**	**RU**	**GB**	**SE**	**UA**	**US**	**CA**
Relays	[%]	19.4	13.0	8.1	4.6	3.4	2.2	1.5	18.0	2.4
Bandwidth	[%]	22.1	22.5	13.1	1.9	3.1	2.3	-	10.6	1.5

A. Network Diversity

Tor's circuit establishment procedure and the transmission characteristics of the underlying network directly influence the end-to-end timing of transmissions. Diverse network infrastructures can be one crucial influencing factor for such varying conditions that have major consequences for timing-based avoidance systems. We identify the following avoidance challenges related to network diversity.

1) **Skewed Relay Distribution.** The worldwide distribution of Tor relays is skewed towards countries with a higher number of Tor supporters. *The biased relay distribution can induce performance impairments when an avoidance decision excludes a high number of nodes.*
2) **Connection Lengths.** The network infrastructure might enforce certain routes for a connection, e. g., in cases where the path between countries is forced to travel through a trans-atlantic cable. *Along with the skewed distribution of relays this influences the length of a connection, which also affects its timing characteristics.*
3) **Connection Failures.** Permanent and temporary partitions in the network infrastructure affect the availability of routes and different circuits. *Such partitions force traffic through specific routes and can hinder the avoidance of a forbidden area.*

All these characteristics address the complex infrastructure of Tor. An avoidance system must incorporate such varying transmission characteristics and provide a flexible decision mechanism that reduces the negative effects of incorrect decisions. An incorrect decision can either lead to collateral damage in cases where conservative security is preferred over performance, or accept critical connections that might harm a user.

1) Skewed Relay Distribution: We analyze the characteristics of one consensus file and derive the distribution of relays and their performance (please refer to Table III for a detailed overview of all measurement setups). The majority of Tor relays runs in Europe, where multiple countries are located within a comparably small area (see Table II). This influences the choice of relays and renders avoiding specific countries within Europe more challenging. We see that $72\,\%$ of all relays are operated in the `EU`[2] and $21\,\%$ run in `NA`; the remaining $7\,\%$ are distributed over all other continents. The same applies to the bandwidth offered, i. e., `EU` provides $81.5\,\%$ of the overall bandwidth, `NA` has $17\,\%$, and all other continents provide no more than the remaining $1.5\,\%$. As the Tor relay selection is weighted towards higher bandwidth nodes, we find the most prominent choices for relays in Europe.

[2] **Continents:** `NA` - North America, `EU` - Europe, `AS` - Asia, `SA` - South America, `OC` - Oceania, **Countries:** `DE` - Germany, `FR` - France, `UA` - Ukraine, `NL` - Netherlands, `GB` - United Kingdom, `SE` - Sweden, `US` - USA, `CA` - Canada, `RU` - Russia, `IN` - India, `SG` - Singapore, `BR` - Brazil

TABLE III. EXPERIMENTAL SETUPS

Protocol	Target	Servers	Nodes	Duration	Num. Results	Section
ICMP	Relays	16	6042	20 h	1,837,761	§III-B1 §III-B1 §III-B3
		8	6042	20 h	27,274 62,643	§VI-B §VI-B
TCP	Weighted Circuits	8	Random	7 d	135,924	§III-A2
	Art. Circuits	8	150,150,150	14 d	360,395 360,395 134,370 223,070	§IV-A §IV-B §VI-B §VI-B
Ntor	Art. Circuits	8	1945,3724,893	4 d	104,889	§V-D §III-A3

Protocol Network protocol used in measurements. *ICMP* messages sent as standard ping, *TCP* messages sent through Tor circuits using a reply server, *Handshakes* are offsets between initial and final `NTor` handshakes.
Target Where probes were sent to. *Relays* are single relay nodes from the consensus; *weighted circuits* are Tor standard circuits where we do not interfere with the relay selection; *art. circuits* are artificial circuits we build from selected relays using the control port.
Servers Number of servers we use for conducting measurements. *ICMP* measurements originate from all servers, as they do not depend on a reply server; *TCP* measurements depend on a reply server and we split the set of servers into senders and receivers
Nodes Number of nodes addressed in a measurement; we use a filtered consensus where all relays provide the `Stable` and `Running` flags. We use this filtered consensus to analyze the distribution of relays (§III-A1). *ICMP* measurements send pings directly to these nodes; *TCP* measurements use Tor circuits, hence, the nodes summarize the relays used to build circuits. Weighted circuits are built from whatever Tor selects, artificial circuits are built from permutations of $m \times n \times l$ entries, middles, exits.
Duration Time elapsed between first and last measurement in a batch of measurements. Might include several repetitions of the same measurement.
Num. Results Total number of individual results.

TABLE IV. CIRCUIT LENGTH

		EU-EU	EU-NA	NA-EU	EU-AS	NA-NA
Median	[km]	4,384	11,117	12,394	12,897	19,210
Minimum	[km]	318	8,425	6,411	10,329	16,907
Maximum	[km]	40,630	45,436	44,807	44,094	51,092

The distribution of relays influences the overall length of circuits, which determines the transmission times between clients and servers. Furthermore, a higher density of relays *and* countries makes it harder to distinguish between different countries. This leads us to an evaluation of the expected *connection lengths*.

Design Goal: Overly restrictive avoidance decisions cause collateral damage in regions with a high density of relays and countries. We must provide a decision mechanism that does not exclude too many relay choices.

2) Connection Lengths: The length of a circuit depends on the client/server location and the distribution of relays involved; the overall distance traveled in a circuit influences the RTT of a transmission. We define the length of a circuit as the distances $(client, entry)+(entry, middle)+(middle, exit)+(exit, server)$ and use approximate direct connections between all nodes from client to server of Tor standard circuits (see Table III for experimental details).

The shortest circuits are built within Europe and the longest circuits in North America (see Table IV for reference, combinations of continents describe the client and server locations). A closer look at the relay locations for all `(NA,NA)` circuits reveals that none of the entry relays was located in `NA`, only 14 % of circuits had a middle relay in `NA`, 27 % had an exit in `NA`, and only 4 % of all circuits went through a middle *and* exit in `NA`. We must assume that even though we established a connection that was limited to `NA`, the circuit traversed the Atlantic twice, which results in a high average circuit length.

Design Goal: Varying circuit lengths lead to individual timing characteristics. The decision threshold of an avoidance system should consider individual characteristics for precise decisions on different connections.

3) Connection Failures: Partitions in the network infrastructure influence the circuit establishment procedure on the application layer and the selection of routes on the network layer. This influences the transmission path of a circuit and eventually its end-to-end timing features. As temporary and permanent connection issues can be caused by a wide range of reasons that are intransparent for an avoidance system, we limit our analyses to a summary of monitored circuit establishment failures (see Table III).

In our measurements, overall 10.65 % of circuit establishments failed (12,500 out of 105,889 circuits). We use the consensus archives to check whether a relay was unavailable during the circuit establishment and distinguish two cases: A relay might *not* be documented in the consensus and we consider it as completely unavailable, or the relay occurs in the list of relays and we can check its flags. Our results reveal 20.32 % of relays that caused the connection failure to have the `Stable` flag set (router is suitable for long-lived circuits), whereas all failure relays are flagged as `Running` (router is currently usable). We find 9 % of failing relays to be entry guards and 29 % to be flagged as exit. On average, failing relays provide 4.85 MB/s of advertised bandwidth (on average a relay provides 7.83 MB/s) and are 3,122 km away from the preceding node in the circuit (see Tab. IV as reference for total circuit lengths). The overall rate of circuit failures is non-negligible and further influences the circuit buildup procedure.

Design Goal: Connectivity issues and partitions amplify the effects of Tor's skewed relay distribution. Decision thresholds must be flexible to respect diverse performance features.

B. Ground Truth Information

Transmissions through Tor and the underlying network infrastructure are not transparent. Therefore, we lack trusted ground truth information about precise relay locations, all hops of the transmission path, or performance features given at the time of transmission. Nevertheless, we depend on a certain set of information to provide a profound avoidance decision. The lack of ground truth information introduces the following challenges.

1) **Relay Locations.** We have no reliable information about the actual positions of relays. GeoIP databases claim to provide accurate information on a country level, nevertheless, such databases are an untrusted source of information [45]. Manipulated or false entries that suggest an incorrect country code are a security threat. *Reference measurements help to verify GeoIP information and provide an additional source of information to identify false country codes.*

2) **Asymmetric Paths.** Routing between the client and server is not performed on symmetric paths, but routes can differ on the way forth and back. *Assuming symmetric paths induces an error in the application of an avoidance decision.*
3) **Transmission Characteristics.** Dynamic routing might change the paths of a circuit between individual transmission sessions. Furthermore, varying network conditions can influence transmissions through the effects of congestion. *Assuming static characteristics introduces inaccuracies in the avoidance system.*

Even though we cannot gain full transparency in the transmissions, preliminary measurements and verification steps help to achieve more trust in the available data. In the following, we introduce a verification mechanism for GeoIP location information, identify the security threat of assuming symmetric paths, and estimate the dynamics of varying transmission characteristics.

1) Relay Locations: Prior avoidance systems use lower bounds to decide whether it is possible that a circuit traverses a forbidden area and for this, the locations of relays must be known. The consensus does not provide coordinates for relays, so the best way to estimate their position is an IP address lookup in a GeoIP database. Unfortunately, such databases provide untrusted information that might lead to false locations. We conduct reference measurements, similar to the approach of Weinberg et al. [52], to *verify* GeoIP locations and identify potential errors.

We measure the ICMP round-trip time between different remote servers and all relays of one consensus (please refer to Table III for details on the experimental setup). In a first evaluation step, we compare the transmission time with the great circle distances between servers and relays to approximate the transmission speed in each measurement (cf. Fig. 1). We use this speed estimate to identify provably false GeoIP locations, i. e., locations that lead to propagation speeds faster than the speed of light. Such a violation occurs in cases where the GeoIP location documents a position that is further away from our reference[3] point than indicated. Consequently, the measured time is too short to travel the entire distance to the recorded position. As we use multiple worldwide server instances, we receive reference measurements from opposing points and identify false information as soon as at least one server indicates a speed of light violation.

From all tested relays, we find approximately $5.5\,\%$ relays (330 out of $6{,}042$) to exceed the maximum allowed propagation speed and, consequently, to be represented through false GeoIP information. Using trilateration, as introduced in Section II, we utilize the reference measurements from our server instances to update the position of obviously false relay locations. Besides the improved location, we update the country code of $3.2\,\%$ (194) of the relays.

Solution: We use ICMP reference measurements to verify the general correctness of untrusted GeoIP locations and identify obvious false locations that violate the speed of light. Trilateration allows us to improve the location data by removing provably false information.

2) Asymmetric Paths: We acknowledge the general approach of the recently proposed system DeTor [31], but find—besides further security and performance issues—a critical overestimation in its lower bound decisions. DeTor bases its mechanism on *symmetric* routes, which is not a valid assumption as has been discussed and demonstrated by Sun et al. [44]. We use this as an example of the consequences of a false asymmetry assumption. In particular, DeTor calculates its decision threshold based on a detour to the forbidden area on *both* directions of a round-trip. This is a critical misconception introduced when the authors transitioned their technique from one-way connections [29] to Tor circuits. A negligent doubling of the necessary distance overestimates the required time. To fix the symmetric routes apparent in DeTor's time estimation (Equation 2), we consider asymmetric routes and obtain:

$$R_{min} = \frac{3}{2c} \cdot min \begin{cases} D_{min}(c, F, e, m, x, s) \\ D_{min}(c, e, F, m, x, s) \\ D_{min}(c, e, m, F, x, s) \\ D_{min}(c, e, m, x, F, s) \end{cases} + D(c, e, m, x, s). \tag{4}$$

The amount of overestimation done by DeTor can be quantified as:

$$R_{error} = \frac{3}{2c} \cdot [D(A, F, B) - D(A, B)], \tag{5}$$

where A and B are the hops with an extension to reach F. DeTor overestimates the decision threshold by R_{error}, which represents the range of false decisions. The greater the distance to F, the higher the overestimation done by DeTor. This constitutes a critical security flaw as connections are falsely labeled secure, creating the illusion of protection from being monitored, and putting users to risk. DeTor uses an uncertainty parameter δ that can be used as a factor to adjust the measured RTT (see Eq.3), nevertheless, this does not fix the system-intrinsic overestimation.

Solution: We consider only one forbidden area extension for the entire connection, i. e., assume asymmetric paths.

3) Transmission Characteristics: Varying transmission characteristics influence the end-to-end timing of a connection, e. g., congestion in relays or routers prolongs the transmission times and can lead to false avoidance decisions. Consequently, the timing characteristics of a circuit depend on the distances between hops and the amount of routing that takes place in between. We utilize the ICMP reference measurements to review real-world timing characteristics and derive the experienced propagation speeds.

Three "clouds" of points (cf. Fig. 1) summarize the transmission distances from remote servers to all relays in the consensus and indicate sparse areas like oceans and continents with only few relays. We apply a nonlinear least squares fit (NLS) to these measurement results and receive the propagation speed as a function of the transmission lengths. The fitting function indicates a dynamic propagation speed rather than a fixed threshold, e. g., we find varying transmission characteristics for different circuit and hop lengths. Typical transmission speeds are in the range of 0.22c to 0.67c [25], whereas we find a maximum speed of $0.381c$ and a mean of $0.342c$ in the NLS fit of all ICMP measurements.

[3] By reference points we refer to remote server instances that we use to conduct measurements. Details of the experimental setups are documented in Table III.

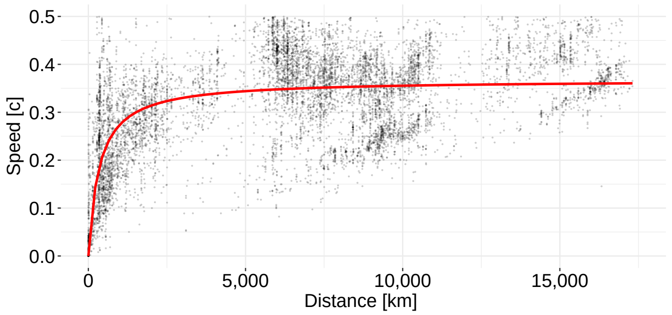

Fig. 1. Measured propagation speeds versus traveled distances (points) of ICMP measurements from 16 reference servers to 6,042 Tor relays; nonlinear least squares fit (NLS) of the relation between transmission distance and propagation speed (red line).

Solution: We use hop-individual timing estimates for all possible pairs of relays and step back from a static speed assumption for all connections.

C. Deployment

Deploying an avoidance system in the real-world means that we need reasonable sources of information for an accept or reject decision. Furthermore, we must maintain the original security and performance features of Tor, as otherwise new attack vectors open up. We define the following requirements for a realistic deployment.

1) **Gather Required Information.** All information required to perform a profound avoidance decision must be made available for Tor users. *For a realistic deployment we need a reasonable source of this information and generate trust in its content.*
2) **Security.** Gathering information for the avoidance decision must maintain the original security features of Tor. *The system must avoid actions that reveal any sensitive information about users or the network.*
3) **Performance.** Additional security through geographical avoidance might justify minor performance impairments, nevertheless, *it remains a design goal to maintain the original performance.*

A real-world deployment leads to additional requirements for the features of an avoidance system, e. g., they define the amount of information we can (or cannot) incorporate in the decision process, and they also dictate the security and performance features that must be satisfied. Even though they are the conditions for the practical deployment of a system, the *deployment* challenges are still independent of the general concept of an avoidance system. In other words, it is possible to propose a general avoidance concept that answers the challenges of missing *ground truth* information and follows the design goals associated with *network diversity* in a first step. As soon as the general avoidance concepts can satisfy these challenges, we can approach the subsequent step of deriving a prototype implementation that also serves the all real-world conditions.

We organize our evaluation procedure according to this two-step workflow. In a first evaluation step, we introduce an empirical avoidance concept and rate its detection capabilities and potential collateral damage in comparison with recent proposals in this context. This initial assessment provides an overview of how different concepts can manage the challenges. In a second step, we extend the experimental setup by real-world constraints and introduce a prototype implementation that utilizes the empirical avoidance concept of Section IV.

IV. Comparison of Avoidance Concepts

The assessment of challenges (§III) is our starting point to evaluate building blocks for a realistic avoidance system. In the following, we introduce an empirical avoidance concept and its system model and compare it with recent work in this context.

A. Empirical Avoidance Decisions

From the preliminary measurements we learned that Tor not only provides a skewed distribution of relays (§III-A1), but also that varying transmission characteristics (§III-B3) and circuit lengths (§III-A2) have a fundamental influence on the end-to-end timings of circuits. Consequently, we lose information when applying a *static* threshold in the avoidance decision. In the following, we propose an alternative approach to estimate the timing characteristics of each hop individually.

1) Relay Hop Time Estimation: Our goal is to obtain a realistic estimation of transmission times between individual hops. We do so by extracting dependencies of circuits that share the same hops. In particular, we analyze RTT measurements of Tor circuits that we gather from remote probing servers. We build these circuits from permutations of entry, middle, and exit relays such that they share pairwise identical hops. This redundancy of circuit segments allows us to estimate the timing distribution that each hop contributes to the overall circuit's RTT. We aim to create a map of RTT relations between all possible combinations of Tor relays:

$$\mathit{optimize} \begin{pmatrix} RTT(c \rightarrow \mathbf{e_1} \rightarrow m_1 \rightarrow x_1 \rightarrow s) \\ RTT(c \rightarrow \mathbf{e_2} \rightarrow m_1 \rightarrow x_1 \rightarrow s) \\ \cdots \\ RTT(c \rightarrow e_1 \rightarrow \mathbf{m_1} \rightarrow x_1 \rightarrow s) \\ RTT(c \rightarrow e_1 \rightarrow \mathbf{m_2} \rightarrow x_1 \rightarrow s) \\ \cdots \\ RTT(c \rightarrow e_1 \rightarrow m_1 \rightarrow \mathbf{x_1} \rightarrow s) \\ RTT(c \rightarrow e_1 \rightarrow m_1 \rightarrow \mathbf{x_2} \rightarrow s) \\ \cdots \end{pmatrix}, \tag{6}$$

where e_m, m_n, x_l are combinations of relays and c, s are the remote servers we measure circuits from. In the above notation, the hops partially overlap, which allows us to define equal segments throughout all measurements. The dependency between measurements allows us to assign portions of the total RTT to individual hops. Notably, the measurements take the Tor and other network overhead into account, resulting in hop time estimations already including realistic overhead metrics.

We define an objective function, which minimizes the error for all combinations of measurements, as shown in Equation 7:

$$\min_x f(x) = \|\mathbf{A}x - b\|, x \geq 0, \tag{7}$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a design matrix we arrange from our measurements, and $b \in \mathbb{R}^m$ is the vector of observations [27],

i. e., the measured RTTs. The design matrix is arranged as follows:

$$\begin{array}{c|cccc} & n_1 \rightarrow n_2 & n_1 \rightarrow n_3 & \dots & n_y \rightarrow n_z \\ \hline m_1 & 1 & 1 & \dots & 1 \\ m_2 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ m_x & 0 & 1 & \dots & 0 \end{array}$$

Here, the rows contain individual measurements m and the columns represent *all* pairs of nodes n, i. e., hops between relays that occurred in the measurements. A 1 denotes that the measured circuit contained this specific hop, whereas a 0 is assigned to all other hops. In total, a maximum of four 1 values can exist in each row, as this represents the number of hops from the client to the server.

Equation 7 represents a non-negative least-squares (NNLS) problem, which is a constrained version of a least squares problem. The sheer size of the problem, i. e., several thousand measurements and tens of thousands of relay combinations, exceeds the processable complexity by magnitudes. Nevertheless, applying a highly optimized solver [27], [37] and the fact that we are dealing with a very sparse design matrix allows us to handle large-scale problems. We implement such a solver to calculate the timing distribution of all hops minimizing the squared error. As a result, we receive a lookup table that provides pairwise estimates for all relays in the consensus.

2) Forbidden Area Decision: The hop estimations are our basis to calculate the time it would take to send data through a forbidden area. In particular, we measure the RTT R_{e2e} for a newly built circuit and identify the involved relays. From these relays, we compute the decision threshold R_{est} that summarizes the expected transmission time for the current circuit. Our approach follows the concept of DeTor (Eq. 2), but uses the empirical estimates instead of translating great circle distances into a lower bound transmission time. We compute the decision threshold R_{est}, the shortest possible extension ext_F to the forbidden area, the hop estimates to send data from the client to the server (excluding the hops involved in the extension), and the estimates for the way back from the server to the client. First, we compute the shortest possible extension:

$$ext_F = min \left\{ \frac{D(A, F, B)}{avg[S(A, F), S(F, B)]} \right\}, \tag{8}$$

where $D(A, F, B)$ denotes the great circle distance D from a node A over the forbidden area to the next hop B. As we cannot know the exact propagation speed for the extension to the forbidden area F and nodes A, B, we approximate the extension using the average empirical speed $avg[S(A, F), S(F, B)]$ of all RTT measurements between the respective countries that summarizes the propagation speeds to $S(A, F)$ and from $S(F, B)$ the forbidden area. If for example the shortest extension takes place between an entry relay in `NL` and a middle relay in `FR` with `UK` as forbidden area, then we use the average propagation speeds of `NL`→`UK` and `UK`→`FR` and apply it to the extension distance. Accordingly, we receive an empirical result for the extension time to F on the shortest possible trip for a circuit. We use the approximate extension time to now define the decision threshold:

$$R_{est} = ext_F + est(c, s \setminus ext_F) + est(e, s), \tag{9}$$

where ext_F is the shortest possible extension (Eq. 8), $est(c, s \setminus ext_F)$ are the estimates of all hops except those involved in the extension, and $est(s, e)$ summarizes all estimates for hops on the way back from the server to the client. In other words, we take a detour to the forbidden area on the trip from the client to the server and have two nodes of the circuit involved in this extension. These two nodes represent the fastest possible option to reach the forbidden area, whereas all other nodes use direct connections. Consequently, we make lookups on the hop estimates for all pairs of nodes *not* involved in the extension to receive the transmission time from the client to the server. On the way back, we follow the asymmetry assumption and now use the estimates for lookups of *all* hops in the circuit. We receive the transmission time from server to client and can add up all components to the decision threshold R_{est}.

In the decision process, we relate the measured time R_{e2e} to our estimated time R_{est}, which we define as our *time ratio* Δ:

$$\Delta = \frac{R_{est}}{R_{e2e}}. \tag{10}$$

The reject/accept decision can now be performed directly against this time ratio. A time ratio of 1 marks the equality of our estimated threshold and the round-trip measurement. A lower Δ is calculated when the measured RTT exceeds the estimation and indicates insecure circuits. On the other hand, a higher Δ results from measurements faster than the estimate. To account for a trade-off between security and performance, we can shift the decision threshold to either end. This allows to establish higher security guarantees or to keep more circuits for the sake of performance. Furthermore, we follow the lower bound threshold of $2/3\ c$ for use cases where provable avoidance is preferred over an empirical decision.

3) System Components: We build our empirical avoidance concept from the above decision mechanism and apply the design goals and solutions introduced with the assessment of challenges in Section III. Our avoidance concept consists of two organizational units, i. e., on the network side we conduct distributed measurements as an information input for the avoidance decision. This includes the ICMP reference measurements for verifying relay locations through trilateration (§III-B1), and the TCP measurements for the computation of pairwise hop estimates (Eq. 6). On the client side, we conduct the circuit measurements where repeated TCP probes through an established circuit reveal the RTT R_{e2e} of the connection. We compare this measured RTT with the empirical threshold R_{est}, which leads to the time ratio between the measured and predicted transmission time (Eq. 10). The final avoidance decision uses this time ratio to rate the current circuit and, eventually, reject or accept its usage. All network side measurements are conducted offline; we discuss realistic ways to realize this following the example of Tor's consensus along with the proposal of the prototype (§V). Client-side measurements must be conducted before an established circuit transmits user data.

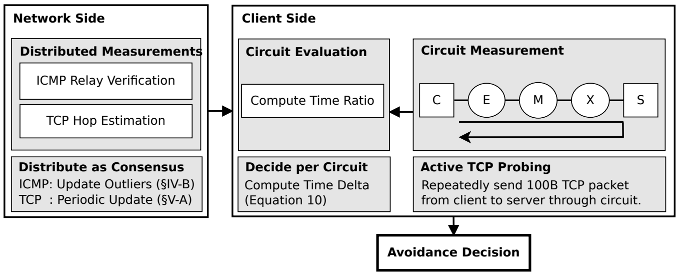

Fig. 2. High level overview of the empirical avoidance concept. Network-side components provide offline information, client-side computes the time ratio from the measured time R_{e2e} and the threshold R_{est}.

B. Experimental Evaluation

We compare our empirical avoidance concept with two other approaches that we distinguish by their consideration of the design challenges in the classes of *network diversity* and *ground truth* (a detailed evaluation of the requirements of a real-world *deployment* follows in § V). We address the original version of DeTor that assumes symmetric paths and static transmission characteristics (Eq. 3); we refer to this as *symmetric* avoidance concept. Furthermore, we compare this to an updated version of DeTor, referred to as *asymmetric* concept, which functions in the exact same way but assumes asymmetric routes to correct the logical flaw of DeTor (Eq. 4). Both static concepts use unverified GeoIP information. Finally, we introduce a novel *empirical* concept that uses hop-individual estimates and verified GeoIP locations. For the sake of comparability, we apply the decision mechanisms to *full* circuits from client to server that were measured by active TCP probing. We are interested in the detection capabilities of two static (symmetric, asymmetric) and one novel empirical avoidance concept. Our evaluation first focuses on the general performance concerning the number of rejected circuits and the avoided advertised bandwidth.

1) Measurement Setup: Our experiments are based upon empirical RTT measurements from the live Tor network, i. e., we use actual transmission characteristics for the computation of hop estimates and use the RTTs of full circuits to simulate avoidance decisions for all three concepts.

Test Set. We perform RTT measurements from eight server instances (CA, NL, US, IN, SG, GB, DE, BR) that send 20 TCP ping messages of $100\,\mathrm{B}$ length through an established Tor circuit. After each message, we wait $1\,\mathrm{s}$ until the next $100\,\mathrm{B}$ are sent to avoid any interaction. In case a reply was not received within the timeout limit of $2\,\mathrm{min}$, we assume a failed connection. From 1,670 entries, 2,712 middles, and 735 exits, we build a total of 70,081 individual circuits and perform 275,509 measurements; the selection of relays is randomized and biased towards higher bandwidth nodes that provide $12.564\,\mathrm{MB/s}$ advertised bandwidth on average ($16.513\,\mathrm{MB/s}$ in entries, $2.735\,\mathrm{MB/s}$ in middles, $18.445\,\mathrm{MB/s}$ in exits).

To ensure that our artificial circuits resemble similar transmission characteristics as weighted standard circuits, we build 135,924 additional weighted circuits using the NEWNYM command from the same remote server instances and compare their characteristics to those of the artificial circuits. The results of a Kolmogorov-Smirnov (KS) test of the probability distributions of both circuit lengths show that artificial (NA, EU), (NA, NA) circuits tend to be shorter than the measured Tor standard circuits, while we find a higher similarity for the other combinations (EU, EU), (EU, NA), (EU, AS), (NA, AS).

Simulation Methodology. We use the RTT measurements of all artificial circuits and compare the detection mechanisms of the three avoidance concepts. For each circuit, we iterate the top nine relay-providing countries (DE,US,FR,UA,RU,NL,GB,SE,CA) as hypothetical forbidden areas, using the following simulation methodology:

1) For all circuits, we identify the shortest possible extension to the current forbidden country F, compute the extension time, and identify its position in the circuit.
2) For the empirical approach, we perform a lookup on the estimated RTT for each hop in the current circuit and approximate the transmission time for the extension hop to the forbidden area F. Using this information, we compute the RTT threshold R_{est} (Eq. 9) and the time ratio Δ (Eq. 10) of the circuit.
3) For the symmetric and asymmetric decision, we follow the detection mechanism proposed in DeTor and compute the time consumption of each hop using the great circle distance between relays and a static speed of ⅔c. We estimate the RTT threshold for a circuit following the definitions of Eq. 2 for the symmetric approach and Eq. 4 for the asymmetric approach. Again, we derive the time ratio Δ for the circuit.
4) We apply a decision threshold of $\Delta \geq 1$ to accept a circuit and handle all other time ratios as a reject decision.

C. Results

To evaluate our results, we analyze the relative number of circuits an avoidance concept rejects for a forbidden area F. Furthermore, we estimate the loss in advertised bandwidth that results from the avoidance decision.

1) Detection Capabilities: The reject and accept rates of a system indicate the restrictions in the choice of circuits when avoiding a specific geographical area. Table V (top) summarizes the reject rates, i. e., the relative number of circuits that were rejected because the measured RTT exceeded the respective threshold. When comparing the symmetric and asymmetric approaches, we see only minor differences for forbidden countries within Europe (DE, FR, UA, NL, GB, SE), but a significantly increased reject rate for the asymmetric approach for US and CA. This is caused by the higher extension distance to North America, i. e., remote forbidden areas emphasize the overestimation of DeTor's symmetric approach (Eq. 5). In comparison, with the hop-individual decision we reject overall approximately $22.64\,\%$ fewer circuits, as a result of the individual consideration of hop RTTs to be less conservative with the comparison threshold.

2) Performance Impairments: Being too conservative with the reject decision can cause severe performance impairments, especially in cases where large user groups decide to circumvent a certain area. The empirical approach manages to reject fewer circuits and can maintain on average $27\,\mathrm{MB/s}$ more (advertised) bandwidth per circuit. Table V (bottom) summarizes the relative bandwidth loss in a worst-case scenario, in which $100\,\%$ of users avoid a certain country. Example: As

TABLE V. Loss [%] of available circuits and bandwidth

Relative Reject	System	DE	US	FR	UA	RU	NL	GB	SE	CA	Average
Circuits	**Symmetric**	89.80	61.08	93.64	88.71	88.02	92.66	94.25	92.27	65.99	**85.17**
	Asymmetric	90.22	73.39	93.75	90.83	86.46	92.82	94.35	92.53	79.02	**88.04**
	Empirical	70.86	59.78	76.22	48.52	55.83	79.10	72.51	60.35	-	**65.40**
Bandwidth	**Symmetric**	85.73	60.16	91.66	85.41	83.58	88.84	91.79	89.06	66.06	**82.48**
	Asymmetric	86.10	73.05	91.74	87.09	84.86	88.99	91.93	89.38	77.54	**85.63**
	Empirical	74.21	63.01	76.03	42.80	51.47	79.57	65.24	56.39	-	**63.59**

we know from the usage statistics [2], approximately $45\,\%$ of Tor's advertised bandwidth is consumed on a daily basis. If we take the $13\,\%$ average daily users of the United States as an example [3] and assume `UA` as the forbidden region, this translates into an overall load factor of $50.56\,\%$ for the individual decision ($56.1\,\%$ for the symmetric, $56.32\,\%$ for the asymmetric decision). Even though our results predict a worst-case scenario, it is likely that a majority of users is motivated to avoid the *same* country due to censorship activities. Losing bandwidth in the range of $85\,\%$ brings us close to an overloaded situation and is unacceptable.

3) Collateral Damage: Conservative reject decisions not only result in performance impairments for a user, but can also cause collateral damage to the entire network. While highly sensitive use cases should maintain a restrictive lower bound threshold, less demanding cases allow for a tradeoff between detection capability and performance. We can adjust the security of the individual implementation by reducing the original decision threshold of 1 for lower time ratios. This increases the chances of routing through the forbidden area, but helps to reduce the reject rates drastically. We have a close look on the potential of using alternative decision threshold in the prototype evaluation (§V).

Summary. Utilizing an empirical decision allows to incorporate hop-individual timings. This reduces the error of a static lower bound threshold that can only represent best case propagation speeds and neglects the varying transmission characteristics of real-world connections. Using an empirical threshold allows to tradeoff performance and security while reducing the collateral damage through overly restrictive decisions. Our results indicate that such collateral damage has an enormous impact on Tor's performance that affects all users and, therefore, cannot be an acceptable trade for security.

V. Prototype Implementation: *TrilateraTor*

TrilateraTor is the prototype implementation of an empirical avoidance system that takes all three classes of challenges (§III) into account. In particular, we extend the empirical avoidance concept (§IV) by features that satisfy the conditions of a practical *deployment* scenario. In the following, we detail the system model, improve the security of the circuit RTT measurement technique, and provide an experimental analysis of the prototype's performance. Finally, we discuss the possible ways to realize the deployment of *TrilateraTor*.

A. System Model

TrilateraTor's geographical avoidance consists of the same organizationl units as the previously introduced empirical concept (cf. Fig. 2). The network-side measurements are conducted in case of changes in the consensus (update ICMP verification) or on a periodical basis (TCP hop estimates). We discuss options to share the information of the `Distributed Measurements` among Tor users in §VI-A and suggest an infrastructure with Tor's bandwidth authority. The first adaption to a real-world deployment takes place on the client side, where we apply a novel measurement technique to recover the end-to-end RTT R_{e2e} of a circuit. In contrast to the generic empirical concept, where we sent TCP ping probes to the entire connection, we now limit our measurements to the Tor nodes in the circuit, i. e., the connection from the client to the exit relay. Again, we compare the measured RTT R_{e2e} to the predicted time R_{est} and derive the time ratio Δ. The time ratio suggests an avoidance decision following the desired tradeoff between performance and security in which we can shift the decision point towards higher (more security) or lower (more performance) thresholds.

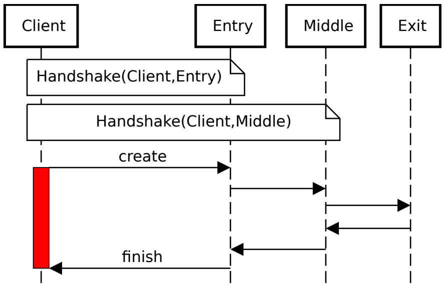

Fig. 3. When building a new circuit, the client performs three handshakes with the entry, middle, and exit relay of the circuit. We are interested in the offset between the `create` and `finish` messages (red bar) that summarizes the RTT for a message to be transmitted from client to exit and back.

B. Avoidance Decision

In the evaluation of avoidance concepts we were able to use hop estimates for full circuits, i. e., our RTT measurements provided us with estimates that also cover hops between the client and entry/exit and server. In a realistic setup, such hops are highly individual and cannot be covered. Furthermore, we switch from actively sending TCP ping probes through an established circuit (introduced by DeTor) to measure the time offset in the key establishment for organizational and security benefits, as we will introduce throughout this section. These changes lead us to an updated decision threshold R_{est} that consists of the shortest possible extension to the forbidden area ext_F, an approximation of the transmission time between client and entry $app_{c,e}$, and the pairwise hop estimates for remaining hops in the circuit. We first define the $c \rightarrow e$ approximation:

$$app_{c,e} = \frac{D(c,e)}{avg(S(c,e))}, \tag{11}$$

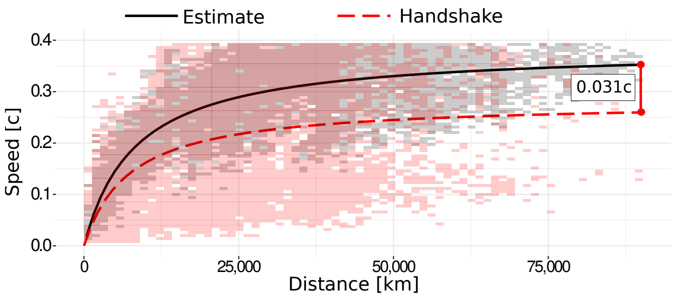

Fig. 4. Comparison of propagation speeds in estimates (black) and `NTor` handshakes (red). Results summarize the spectrum of measured times for full circuits (areas) and the NLS propagation speed fit (lines).

where $D(c, e)$ is the great circle distance from the client to the entry, and $avg(S(c, e))$ is the average measured propagation speed from the client's country to the entry's country. Using this approximation of the first hop, we now define two cases for the definition of the predicted transmission time R_{est}:

$$R_{est} = \begin{cases} 2 \cdot app_{c,e} + ext_F + est(e, x, e \setminus ext_F) & , c \notin ext_F \\ ext_F + est(c, x, c) + app_{c,e} & , c \in ext_F \end{cases},$$

where we distinguish an extension to F that happens without including the client ($c \notin ext_F$), or an extension that takes place between the client and the entry node ($c \in ext_F$). In the first case, we approximate the hop from the client to the entry twice for both directions of the connection and estimate the remaining hops and the extension as usual. In the second case, the shortest possible extension is between the client and the entry, and we only need to approximate this hop for the way back from the exit to the client. Just as in the empirical concept, we estimate only hops that are not involved in the extension ($est(e, x, e \setminus ext_F)$); the estimation $est(c, x, c)$ includes all hops from client c to exit x and back to c. For the final avoidance decision, we compare the predicted time R_{est} with the measured RTT between the client and the exit R_{e2e} and reject or accept a circuit according to the time ratio and desired decision threshold. We derive the circuit timing from a novel measurement technique that we introduce as follows.

C. Circuit Establishment Timing

On each new circuit establishment, Tor performs three cryptographic handshakes with the entry, middle, and exit relay of a connection (cf. Fig. 3). Each of these handshakes traverses parts of the circuit and delivers the end-to-end timing information R_{e2e}. We measure the handshake timings in the `NTor` [46] implementation, which provides Tor's cryptographic primitives since version `0.2.4.x`. More precisely, we measure the offsets (red bar) between the `create` and the final `finish` messages. The client repeats the handshake procedure for each relay in the circuit and finally delivers the total transmission time between the client and the exit. In contrast to active TCP probing, the Tor client performs the required cryptographic handshakes at each circuit establishment, i. e., we can derive all relevant information without any active interference. We will see later how this benefits the usability of the avoidance system and overcomes one existing security issue of DeTor.

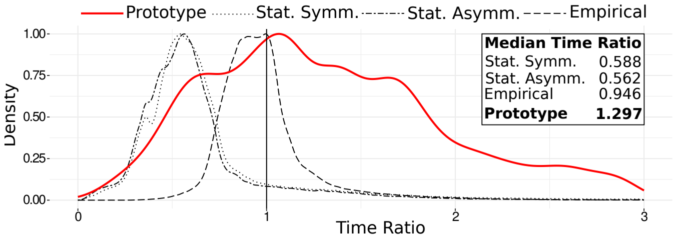

Median Time Ratio	
Stat. Symm.	0.588
Stat. Asymm.	0.562
Empirical	0.946
Prototype	**1.297**

Fig. 5. Distribution of time ratios for theoretical avoidance concepts (Stat. Symm., Stat. Asymm., Empirical) in comparison with the prototype implementation *TrilateraTor*. Higher time ratios indicate higher accept rates.

D. Experiments

In our prototype evaluation, we first address the timing characteristics of RTTs derived from the circuit establishment procedure and compare them to the characteristics we observed for TCP pings (§III). Furthermore, we analyze the performance of *TrilateraTor* in comparison to the theoretical avoidance concepts of §IV.

1) Experimental Setup: We use eight worldwide server instances (`CA,NL,US,US,IN,SG,GB,DE`) and measure a total of 16,500 individual relay combinations (1,945 entries, 3,724 middles, 893 exits) for the exit handshake offsets. For each measurement, we draw 100 top bandwidth relays from the first consensus of the day and form random circuits from this; measurements are repeated every 10 min within a period of three days. We document the handshake timings for all successful circuits along with the relay at which the buildup procedure failed in case of an unfinished circuit establishment. Besides the handshake measurements, we repeat the ICMP reference measurements (§III-B1) and TCP ping measurements (§IV-A) to provide recent information for the verification of relay positions and the estimation of hop relations.

2) Timing Characteristics: We analyze the robustness of timings from the key agreement procedure, as in comparison to repeated TCP Ping measurements, a much smaller data basis for the decision is given. To do so, we measure the median deviation of handshake times between identical hops. Our results indicate that exit handshakes differ by 6.54 % between measurements, which results in an average variance of 32 ms.

Furthermore, we compare the propagation speed of handshake timings with the TCP ping hop estimates (cf. Fig. 4). Obviously, the cryptographic computations of the handshake procedure induce an additional overhead that leads to an overall reduced propagation speed in comparison to the hop estimates. As this computational overhead is not related to the transmission characteristics of a connection, we speed up the handshake measurements by the difference $0.031c$ between the estimates and the handshakes.

3) Performance: We now compare the performance of *TrilateraTor* with the theoretical concepts of §IV and analyze the spectrum of time ratios that results from avoiding the top nine relay providing countries (cf. Fig. 5). We see that the static avoidance concepts lead to smaller time ratios, which supports the finding that a fixed speed assumption leads to overly restrictive decisions. In comparison, the empirical approach achieves a median time ratio of 0.946 and with that is close to a balanced distribution of decisions.

Our prototype implementation improves this result further, e. g., we achieve a median time ratio of 1.297 that allows accepting a high number of circuits. At this point it is important to once more emphasize the structural differences between the theoretical concepts and *TrilateraTor*: While we can analyze *full* circuits in the theoretical concepts, the prototype is limited to connections from the client to the exit. Nevertheless, the comparison of time ratios still delivers an essential perspective on the differences between the theoretical concepts and the practical implementation.

VI. Deployment of *TrilateraTor*

For successfully deploying *TrilateraTor*, we depend on reliable sources for all information that we consider in the avoidance decision. Furthermore, we must maintain Tor's original level of security and limit the performance impairments that the avoidance feature induces. In the following, we discuss the organizational aspects of deployment, analyze the security features of *TrilateraTor*, and estimate potential performance impairments.

A. Information Sources

We depend on three different sources of information for an avoidance decision in *TrilateraTor*: Distributed ICMP and TCP measurements improve the trust in GeoIP information and deliver the empirical timing information for individual hops between relays. Handshake measurements allow us to derive a circuit's RTT without any active probing.

1) ICMP and TCP Measurements: The strength of the ICMP and TCP measurements lies in the fact that we use multiple worldwide server instances that either reliably identify false relay locations in the case of ICMP, or generate representative empirical estimates of the timing characteristics between hops (TCP). While we will see later (§VI-C) how this adds another layer of security, we are now interested in ways to organize these distributed measurements in case of deployment.

We assume fixed relay locations, i. e., ICMP reference measurements only require updates for changes in the consensus. Throughout 2017, a fluctuation of approximately 17 % occurred when existing relays disappeared or new relays appeared in the consensus. For an average number of 7,283 relays in the consensus, this translates to approximately 1,238 nodes that require updates (in a worst case, only new relays need to be verified). The situation is different for TCP measurements, e. g., we do not only need to cover fluctuations in the consensus but must also consider varying transmission characteristics (§III-B3). Therefore, periodic updates help to improve the data basis for the pairwise hop estimations. For both the ICMP and TCP information, we advocate a consensus-centric infrastructure that allows users to access all relevant information.

The overhead through distributed measurements is negligible in comparison to Tor's daily usage and the provided capacities. We can assume approximately 2.8 Mio. daily Tor users, and an average consumed bandwidth of 121.5 Gbit/s. Under the assumption that an average user builds at least three circuits (this is a minimum estimate, numbers should be much higher), all experimental circuits represent approximately $4 \times 10^{-4}\,\%$ of Tor's daily circuits. To send 500 messages (as an upper bound for the number of probes sent) with a length of 100 B, we require 758 kbit/s per day, which is only $6.24 \times 10^{-7}\,\%$ of the daily bandwidth consumption in Tor.

TABLE VI. Measurement statistics [ms]

Type	Iteration	Mean	Median	SD	Duration	#Results
TCP	1	287.35	288.46	157.51	5 days	223,070
	2	358.89	335.28	179.58	7 days	134,370
	3	327.39	294.66	185.26	8 days	275,509
ICMP	1	98.95	67.35	97.68	1 day	27,274
	2	55.79	17.5	76.79	1 day	62,643
	3	135.85	128	102.42	2 days	1,837,761

B. Reproducibility

Our experimental setups can only represent snapshots of Tor's network infrastructure and describe the period in which empirical data was gathered. Differences might arise from varying network conditions (congestion, outages, attacks), the selection of measurement points (server locations), the hypothetical forbidden areas, etc. Nevertheless, our selection of experimental components represents worldwide server positions, top bandwidth relays provide the majority of Tor's performance capacities, and the number of conducted measurements delivers a profound data basis. Table VI summarizes the characteristics of our repeated measurements. In both sets, the first two iterations were conducted within one month, whereas the third iteration serves as a reference from measurements gathered six months later.

Results show a high standard deviation (**SD**) within all measurement sets, but are in a comparable range between iterations (27.75 ms delta between TCP results; 25.64 ms for ICMP). The variance of results once more confirms the findings of our challenges assessment, i. e., transmission characteristics depend on the infrastructure and health of the network and change, accordingly. Nevertheless, results are sufficiently comparable even through longer measurement periods.

C. Security

Any behavior that leaks information can open new attack vectors and, consequently, harm users that depend on additional protection mechanisms. Furthermore, overly restrictive decisions reduce one of Tor's core security features, the anonymity set size, and even facilitate traffic-analysis attacks. In the following, we discuss security implications that could arise from deploying an avoidance system.

1) Fingerprinting: An adversary that actively monitors the circuit establishment procedure might recognize deviations from expected patterns and derive fingerprinting information from this. The ability to fingerprint actions of the avoidance system can reveal the endpoints of a connection, help to derive sensitive information because of unexpected user behavior, or help to reduce the anonymity set.

Revealing Connection Endpoints. Measuring the timing characteristics of a circuit through TCP pings requires sending messages along the full transmission path between the client and the server. Under the assumption of a strong AS- or state-level adversary, such messages can reveal the endpoints of a connection, as the RTT measurements also include the destination of the connection. *TrilateraTor does not leak such information, as it utilizes the crypto handshake of the circuit buildup procedure.*

Unexpected Behavior. Active TCP ping measurements add unexpected traffic to the standard transmission patterns of a user. An adversary can monitor batches of TCP probes sent out by the avoidance system and derive additional information from this. Such information includes the presence of an avoidance system and might help to predict the choice of relays. *TrilateraTor attaches to already existing functions of Tor and does not depend on active probing, i. e., it maintains the original circuit buildup behavior.*

Reducing the Anonymity Set. Rejecting a majority of circuits helps an adversary to predict the remaining set of relays that are suitable candidates to circumvent a forbidden area. As a consequence, traffic-analysis attacks become more likely, and the measurement overhead is reduced—both factors would otherwise only enable very powerful adversaries to succeed. *TrilateraTor manages to reduce the number of rejected circuits and, furthermore, allows to apply a context-sensitive tradeoff between security and performance.*

2) Measurement Manipulation: A powerful, nation-state adversary can manipulate [57] the distributed measurements (ICMP, TCP ping) of an avoidance system by holding back probes. This results in an overall increased transmission time that would manipulate the relay verification and computation of hop estimates. *TrilateraTor* inherently limits the impact of such attacks. All network side measurements are conducted from multiple reference points, i. e., the scenario is comparable to verifiable trilateration as proposed by Čapkun et al. [47]. Prolonging one distance to a reference would require the shortening of at least one other distance to a different reference. However, this would need accelerating packets beyond typical Internet transmission speeds mitigating the manipulation success while leaving conspicuous attack fingerprints. *TrilateraTor protects against measurement manipulations that otherwise would affect a timing-based avoidance decision.*

D. Performance

Two influencing factors have the potential to impair Tor's original level of performance. First, timing-based avoidance systems depend on RTT measurements for a tested circuit. Prior work introduced active TCP probing where a client sends messages through the established circuit and measures the offset until the response was received. This approach forces users to wait until the measurement procedure is finished and hinders from directly using a (safe) circuit. We overcome this by using the circuit establishment handshake as an information source, i. e., we induce no additional waiting time. Second, restrictive avoidance decisions limit Tor's available resources. Our worst case evaluation (§IV and Table V) proves that *TrilateraTor* manages to reduce this source of collateral damage by preserving approximately 22 % more of Tor's advertised bandwidth. This is an important result, as a congested infrastructure also affects users that do not make use of *TrilateraTor*. From an individual perspective, users must always accept slight performance impairments through geographical avoidance, as rejecting the most prominent relay choices often leads to weaker circuits.

VII. Related Work

There are alternative ways besides geographical avoidance to circumvent the threats of traffic-analysis attacks and censorship. Sophisticated path selection in Tor focuses on a dynamic selection of nodes that takes network characteristics like congestion into account. Other ways of avoidance use traffic obfuscation to overcome targeted blocking of Tor.

A. Path Selection

Tor selects the relays for a circuit according to the bandwidth they can offer, and nodes with better performance are preferred over smaller and often less stable ones. This does not only influence the performance we experience when using Tor, but it also has an impact on the anonymity set [6], [15], [23] in which we hide. Recent work suggested different strategies to improve Tor's circuit establishment. Better congestion management [20], [48] can help to improve the load balancing in Tor and increase the number of relay candidates for a circuit through a better distribution of traffic.

Another important factor are autonomous systems (AS) [4], [8], [14], [22], as an AS-level adversary is in a powerful position to perform traffic-analysis attacks. Similar to geographical avoidance, the core principle of AS awareness is the circumvention of untrusted areas, while the analysis of paths must take place in a different layer of the network stack.

B. Censorship Circumvention

Geographical avoidance is just an indirect solution to the problem of Internet censorship and traffic-analysis attacks in particular, and we will introduce two classes of alternative circumvention approaches. Decoy routing is a technique that combines obfuscation and the support of a proxy to access content that is otherwise censored. Pluggable transports are an obfuscation extension to Tor that help make standard Tor traffic look like something else, e. g., any other traffic that is not the target of monitoring and blocking. Please note that both types of circumvention do not consider *geographical* avoidance.

Decoy Routing: Decoy Routing [16], [24], [55] circumvents censorship and blocking by routing critical traffic through servers outside the censored area. For this context, we assume to be located in a country where specific sites are prohibited and all requests made to such contents are blocked or even reported. To overcome this situation, so-called decoy routers are used that provide accepted content outside the censored area, hence, sites that are not hosted in your country but tolerated by the censor. They act as said man-in-the-middle and forward requests to blocked sites as well as send the contents back to the client, all obfuscated through techniques that hide the actual payload of a transmission. Countermeasures like RAD (routing around decoys) [41] try to avoid the functionality of decoy routers by forcing routes on alternate paths that cannot traverse the decoy router. In general, decoy routing is another possibility to circumvent censorship, but it does not consider routing attacks and potential consequences through traffic analysis.

Pluggable Transports: Pluggable transports [10], [35], [53] are another way of censorship circumvention as they provide access to Tor even in case standard circuits and bridges are not an option because of blocking. Despite the wide range of pluggable transport types, the general principle uses obfuscation to make Tor traffic look like some other, benign, protocol that is not the target of blocking or monitoring through the censor. In the context of geographical avoidance, the use of pluggable transports is complicated. They cannot guarantee secure routes, even though the obfuscation techniques make it much harder for an adversary to learn sensitive information from traffic metadata. That said, *it depends*. Random patterns in the obfuscation help to disrupt relations between traffic streams and detecting a relation to Tor becomes more difficult, but there is no protection against routing attacks.

VIII. Conclusion

In this work, we assessed challenges of geographical avoidance for data transmissions and used it as a foundation to introduce a novel empirical avoidance concept. To this end, our concept considers hop-individual transmission characteristics instead of static thresholds for individual connections, limiting the collateral damage through overly restrictive avoidance decisions. In a two-fold experimental study, we first compared the performance of our empirical avoidance concepts to existing work and managed to outperform other approaches by rejecting $22\,\%$ fewer circuits and maintaining on average $27\,\mathrm{MB/s}$ more advertised bandwidth. In a second step, we introduced the prototype implementation *TrilateraTor* that considers the requirements of a real-world deployment in addition to the challenges of Tor's diverse network infrastructure and untrusted ground truth information. *TrilateraTor* is the first to provide *practical* geographical avoidance and overcomes fundamental security issues of prior systems.

Acknowledgment

This work was supported by Intel as part of ICRI-CARS. In addition, this work was supported by the German Federal Ministry of Education and Research (BMBF Grant 16KIS0664 SysKit) and the Franco-German BMBF project BERCOM (13N13741).

References

[1] "SecureList: Law Enforcement Agencies in Tor," https://securelist.com/law-enforcement-agencies-in-tor-impact-over-the-dark-web/67574/, accessed: 2018-05-04.

[2] "Tor Metrics: Bandwidth Usage for October 2017," https://metrics.torproject.org/bandwidth.html?start=2017-10-01&end=2017-10-31, accessed: 2018-02-01.

[3] "Tor Metrics: Users by Country," https://metrics.torproject.org/userstats-relay-table.html, accessed: 2018-05-04.

[4] M. Akhoondi, C. Yu, and H. V. Madhyastha, "LASTor: A Low-Latency AS-Aware Tor Client," in *IEEE Symposium on Security and Privacy*, ser. SP '12. San Francisco, CA, USA: IEEE, May 2012, pp. 476–490.

[5] D. G. Andersen, H. Balakrishnan, N. Feamster, T. Koponen, D. Moon, and S. Shenker, "Accountable Internet Protocol (AIP)," in *Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '08. Seattle, WA, USA: ACM, Aug. 2008, pp. 339–350.

[6] M. Backes, A. Kate, S. Meiser, and E. Mohammadi, "(Nothing else) MATor(s): Monitoring the Anonymity of Tor's Path Selection," in *ACM Conference on Computer and Communications Security*, ser. CCS '14. Scottsdale, AZ, USA: ACM, Nov. 2014, pp. 513–524.

[7] H. Ballani, P. Francis, and X. Zhang, "A Study of Prefix Hijacking and Interception in the Internet," in *Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, ser. SIGCOMM '07. Kyoto, Japan: ACM, Aug. 2007.

[8] A. Barton and M. Wright, "DeNASA: Destination-Naive AS-Awareness in Anonymous Communications," *Proceedings on Privacy Enhancing Technologies*, vol. 2016, no. 4, pp. 356–372, Oct. 2016.

[9] K. Bauer, D. McCoy, D. Grunwald, T. Kohno, and D. Sicker, "Low-Resource Routing Attacks Against Tor," in *ACM Workshop on Privacy in Electronic Society*, ser. WPES '07. Alexandria, VA, USA: ACM, Oct. 2007, pp. 11–20.

[10] D. J. Bernstein, M. Hamburg, A. Krasnova, and T. Lange, "Elligator: Elliptic-Curve Points Indistinguishable from Uniform Random Strings," in *ACM Conference on Computer and Communications Security*, ser. CCS '13. Berlin, Germany: ACM, Nov. 2013, pp. 967–980.

[11] G. Danezis, "Statistical Disclosure Attacks: Traffic Confirmation in Open Environments," in *Security and Privacy in the Age of Uncertainty: IFIP TC11 International Conference on Information Security*, ser. SEC '03. Athens, Greece: Kluwer, May 2003, pp. 421–426.

[12] G. Danezis, C. Diaz, and C. Troncoso, "Two-Sided Statistical Disclosure Attack," in *International Workshop on Privacy Enhancing Technologies*, ser. PET '07. Ottawa, ON, Canada: Springer, Jun. 2007, pp. 30–44.

[13] K. P. Dyer, S. E. Coull, T. Ristenpart, and T. Shrimpton, "Peek-a-Boo, I Still See You: Why Efficient Traffic Analysis Countermeasures Fail," in *IEEE Symposium on Security and Privacy*, ser. SP '12. San Francisco, CA, USA: IEEE, May 2012, pp. 332–346.

[14] M. Edman and P. Syverson, "AS-Awareness in Tor Path Selection," in *ACM Conference on Computer and Communications Security*, ser. CCS '09. Chicago, IL, USA: ACM, Nov. 2009, pp. 380–389.

[15] A. Edmundson, R. Ensafi, N. Feamster, and J. Rexford, "Characterizing and Avoiding Routing Detours Through Surveillance States," *CoRR*, vol. abs/1605.07685, May 2016. [Online]. Available: http://arxiv.org/abs/1605.07685

[16] D. Ellard, C. Jones, V. Manfredi, W. T. Strayer, B. Thapa, M. Van Welie, and A. Jackson, "Rebound: Decoy Routing on Asymmetric Routes Via Error Messages," in *IEEE Conference on Local Computer Networks*, ser. LCN '15. Clearwater Beach, FL, USA: IEEE, Oct. 2015, pp. 91–99.

[17] D. Fifield, C. Lan, R. Hynes, P. Wegmann, and V. Paxson, "Blocking-Resistant Communication through Domain Fronting," *Proceedings on Privacy Enhancing Technologies*, vol. 2015, no. 2, pp. 46–64, 2015.

[18] Y. Gilad and A. Herzberg, "Spying in the Dark: TCP and Tor Traffic Analysis," in *Privacy Enhancing Technologies Symposium*, ser. PETS '12. Vigo, Spain: Springer, Jul. 2012, pp. 100–119.

[19] K. Jansen, M. Schäfer, D. Moser, V. Lenders, C. Pöpper, and J. Schmitt, "Crowd-GPS-Sec: Leveraging Crowdsourcing to Detect and Localize GPS Spoofing Attacks," in *IEEE Symposium on Security and Privacy*, ser. SP '18. San Francisco, CA, USA: IEEE, May 2018, pp. 1018–1031.

[20] R. Jansen, J. Geddes, C. Wacek, M. Sherr, and P. Syverson, "Never Been KIST: Tors Congestion Management Blossoms with Kernel-Informed Socket Transport," in *USENIX Security Symposium*, ser. USENIX Security '14. San Diego, CA, USA: USENIX Association, Aug. 2014, pp. 127–142.

[21] R. Jansen, M. Juarez, R. Galvez, T. Elahi, and C. Diaz, "Inside Job: Applying Traffic Analysis to Measure Tor from Within," in *Network and Distributed System Security Symposium*, ser. NDSS '17. San Diego, CA, USA: Internet Society, Feb. 2017.

[22] A. Johnson, C. Wacek, R. Jansen, M. Sherr, and P. Syverson, "Users Get Routed: Traffic Correlation on Tor by Realistic Adversaries," in *ACM Conference on Computer and Communications Security*, ser. CCS '13. Berlin, Germany: ACM, Nov. 2013, pp. 337–348.

[23] J. Juen, A. Johnson, A. Das, N. Borisov, and M. Caesar, "Defending Tor from Network Adversaries: A Case Study of Network Path Prediction," *Proceedings on Privacy Enhancing Technologies*, vol. 2015, no. 2, pp. 171–187, Jun. 2015.

[24] J. Karlin, D. Ellard, A. W. Jackson, C. E. Jones, G. Lauer, D. P. Mankins, and W. T. Strayer, "Decoy Routing: Toward Unblockable Internet Communication," in *USENIX Workshop on Free and Open Communications on the Internet*, ser. FOCI '11. San Francisco, CA, USA: USENIX Association, Aug. 2011.

[25] E. Katz-Bassett, J. P. John, A. Krishnamurthy, D. Wetherall, T. Anderson, and Y. Chawathe, "Towards IP Geolocation Using Delay and Topology Measurements," in *ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '06. Rio de Janeiro, Brazil: ACM, Oct. 2006, pp. 71–84.

[26] D. Kesdogan, D. Agrawal, and S. Penz, "Limits of Anonymity in Open Environments," in *International Workshop on Information Hiding*, ser. IH '02. Noordwijkerhout, The Netherlands: Springer, Oct. 2002, pp. 53–69.

[27] D. Kim, S. Sra, and I. S. Dhillon, "A Non-Monotonic Method for Large-Scale Non-Negative Least Squares," *Optimization Methods and Software*, vol. 28, no. 5, pp. 1012–1039, Oct. 2013.

[28] K. Kohls and C. Pöpper, "DigesTor: Comparing Passive Traffic Analysis Attacks on Tor," in *European Symposium on Research in Computer Security*, ser. ESORICS '18. Springer, Sep. 2018.

[29] D. Levin, Y. Lee, L. Valenta, Z. Li, V. Lai, C. Lumezanu, N. Spring, and B. Bhattacharjee, "Alibi Routing," in *Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '15. London, United Kingdom: ACM, Aug. 2015, pp. 611–624.

[30] B. N. Levine, M. K. Reiter, C. Wang, and M. Wright, "Timing Attacks in Low-Latency Mix Systems," in *International Conference on Financial Cryptography*, ser. FC '04. Key West, Florida, USA: Springer, Feb. 2004, pp. 251–265.

[31] Z. Li, S. Herwig, and D. Levin, "DeTor: Provably Avoiding Geographic Regions in Tor," in *USENIX Security Symposium*, ser. USENIX Security '17. Vancouver, BC, Canada: USENIX Association, Aug. 2017, pp. 343–359.

[32] Z. Ling, X. Fu, W. Jia, W. Yu, D. Xuan, and J. Luo, "Novel Packet Size-Based Covert Channel Attacks Against Anonymizer," *IEEE Transactions on Computers*, vol. 62, no. 12, pp. 2411–2426, Dec. 2013.

[33] Z. Ling, J. Luo, W. Yu, X. Fu, D. Xuan, and W. Jia, "A New Cell Counter Based Attack Against Tor," in *ACM Conference on Computer and Communications Security*, ser. CCS '09. Chicago, IL, USA: ACM, Nov. 2009, pp. 578–589.

[34] N. Mathewson and R. Dingledine, "Practical Traffic Analysis: Extending and Resisting Statistical Disclosure," in *International Workshop on Privacy Enhancing Technologies*, ser. PET '04. Toronto, ON, Canada: Springer, May 2004, pp. 17–34.

[35] H. Mohajeri Moghaddam, B. Li, M. Derakhshani, and I. Goldberg, "SkypeMorph: Protocol Obfuscation for Tor Bridges," in *ACM Conference on Computer and Communications Security*, ser. CCS '12. Raleigh, NC, USA: ACM, Oct. 2012, pp. 97–108.

[36] S. J. Murdoch and G. Danezis, "Low-Cost Traffic Analysis of Tor," in *IEEE Symposium on Security and Privacy*, ser. SP '05. Oakland, CA, USA: IEEE, May 2005, pp. 183–195.

[37] J. Myre, E. Frahm, D. Lilja, and M. Saar, "TNT-NN: A Fast Active Set Method for Solving Large Non-Negative Least Squares Problems," in *International Conference on Computational Science*, ser. ICCS '17. Zurich, Switzerland: Elsevier, Jun. 2017, pp. 755–764.

[38] J. Naous, M. Walfish, A. Nicolosi, D. Mazières, M. Miller, and A. Seehra, "Verifying and Enforcing Network Paths with ICING," in *International Conference on emerging Networking EXperiments and Technologies*, ser. CoNEXT '11. Tokyo, Japan: ACM, Dec. 2011, pp. 30:1–30:12.

[39] R. Nithyanand, O. Starov, A. Zair, P. Gill, and M. Schapira, "Measuring and Mitigating AS-level Adversaries Against Tor," in *Symposium on Network and Distributed System Security*, ser. NDSS '16. San Diego, CA, USA: Internet Society, Feb. 2016.

[40] ——, "Measuring and Mitigating AS-level Adversaries Against Tor," in *Symposium on Network and Distributed System Security*, ser. NDSS '16. San Diego, CA, USA: Internet Society, Feb. 2016.

[41] M. Schuchard, J. Geddes, C. Thompson, and N. Hopper, "Routing Around Decoys," in *ACM Conference on Computer and Communications Security*, ser. CCS '12. Raleigh, NC, USA: ACM, Oct. 2012, pp. 85–96.

[42] H. Sengar, Z. Ren, H. Wang, D. Wijesekera, and S. Jajodia, "Tracking Skype VoIP Calls Over The Internet," in *IEEE Conference on Computer Communications*, ser. INFOCOM '10. San Diego, CA, USA: IEEE, Mar. 2010.

[43] V. Shmatikov and M.-H. Wang, "Timing Analysis in Low-Latency Mix Networks: Attacks and Defenses," in *European Symposium on Research in Computer Security*, ser. ESORICS '06. Hamburg, Germany: Springer, Sep. 2006, pp. 18–33.

[44] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, and P. Mittal, "RAPTOR: Routing Attacks on Privacy in Tor," in *USENIX Security Symposium*, ser. USENIX Security '15. Washington, D.C., USA: USENIX Association, Aug. 2015, pp. 271–286.

[45] T. Tickets. Check Maxmind GeoIPLocation Database before distributing. [Online]. Available: https://trac.torproject.org/projects/tor/ticket/26240

[46] Torspec. NTor Handshake Proposal. [Online]. Available: https://gitweb.torproject.org/torspec.git/tree/proposals/216-ntor-handshake.txt

[47] S. Čapkun and J. P. Hubaux, "Secure Positioning of Wireless Devices with Application to Sensor Networks," in *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, vol. 3, March 2005, pp. 1917–1928 vol. 3.

[48] T. Wang, K. Bauer, C. Forero, and I. Goldberg, "Congestion-Aware Path Selection for Tor," in *International Conference on Financial Cryptography and Data Security*, ser. FC '12. Kralendijk, Bonaire: Springer, Feb. 2012, pp. 98–113.

[49] X. Wang, S. Chen, and S. Jajodia, "Tracking Anonymous Peer-to-Peer VoIP Calls on the Internet," in *ACM Conference on Computer and Communications Security*, ser. CCS '05. Alexandria, VA, USA: ACM, Nov. 2005, pp. 81–91.

[50] ——, "Network Flow Watermarking Attack on Low-Latency Anonymous Communication Systems," in *IEEE Symposium on Security and Privacy*, ser. SP '07. Oakland, CA, USA: IEEE, May 2007, pp. 116–130.

[51] X. Wang and D. S. Reeves, "Robust Correlation of Encrypted Attack Traffic Through Stepping Stones by Manipulation of Interpacket Delays," in *ACM Conference on Computer and Communications Security*, ser. CCS '03. Washington, D.C., USA: ACM, Oct. 2003, pp. 20–29.

[52] Z. Weinberg, S. Cho, V. Sekar, and P. Gill, "How to Catch when Proxies Lie: Verifying the Physical Locations of Network Proxies with Active Geolocation," in *Proceedings of the 18th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '18. New York, NY, USA: ACM, 2018.

[53] Z. Weinberg, J. Wang, V. Yegneswaran, L. Briesemeister, S. Cheung, F. Wang, and D. Boneh, "StegoTorus: A Camouflage Proxy for the Tor Anonymity System," in *ACM Conference on Computer and Communications Security*, ser. CCS '12. Raleigh, NC, USA: ACM, Oct. 2012, pp. 109–120.

[54] P. Winter and S. Lindskog, "How the Great Firewall of China is Blocking Tor," in *USENIX Workshop on Free and Open Communications on the Internet*, ser. FOCI '12. Bellevue, WA, USA: USENIX Association, Aug. 2012.

[55] E. Wustrow, C. M. Swanson, and J. A. Halderman, "TapDance: End-to-Middle Anticensorship without Flow Blocking," in *USENIX Security Symposium*, ser. USENIX Security '14. San Diego, CA, USA: USENIX Association, Aug. 2014, pp. 159–174.

[56] W. Yu, X. Fu, S. Graham, D. Xuan, and W. Zhao, "DSSS-Based Flow Marking Technique for Invisible Traceback," in *IEEE Symposium on Security and Privacy*, ser. SP '07. Oakland, CA, USA: IEEE, May 2007, pp. 18–32.

[57] D. J. Zage and C. Nita-Rotaru, "On the Accuracy of Decentralized Virtual Coordinate Systems in Adversarial Networks," in *Proceedings of the 14th ACM Conference on Computer and Communications Security*, ser. CCS '07. New York, NY, USA: ACM, 2007, pp. 214–224. [Online]. Available: http://doi.acm.org/10.1145/1315245.1315273

[58] Y. Zhu, X. Fu, B. Graham, R. Bettati, and W. Zhao, "On Flow Correlation Attacks and Countermeasures in Mix Networks," in *International Workshop on Privacy Enhancing Technologies*, ser. PET '04. Toronto, ON, Canada: Springer, May 2004, pp. 207–225.