How Anonymous is the Tor Network? A Long-Term Black-Box Investigation

R. Koch, M. Golling and G.D. Rodosek

How Anonymous is the Tor Network? A Long-Term Black-Box Investigation

Robert Koch, Mario Golling and Gabi Dreo Rodosek

***Abstract*—A popular choice for anonymous Internet communication, the Tor network uses entry, relay, and exit nodes to hide the traffic's origin. However, an investigation that involved running real applications and website requests through Tor revealed numerous agglomerations of exiting traffic that an attacker could exploit.**

***Index Terms*—Communication networks, Computer hacking, Computer security, IP networks, Internet, Privacy, Relays, Routing protocols, Surveillance, Tor, anonymous communication, deanonymization, hackers, network security, network surveillance, networking, networks, onion routing, routers, security**

I. Introduction

With many countries limiting both freedom of speech and the press [1], and with privacy concerns being paramount in less restrictive nations, assurances that anonymous Internet communication can indeed provide anonymity have become more important. At the same time, shadowy activities - such as drug trafficking through the Silk Road, publicizing classified information, and planning and coordinating terrorist activity like the November 2015 Paris attacks—have increased interest in breaking anonymized network communication.

In light of anonymity's two sides, both individuals and surveillance organizations are questioning the strength of popular communication services such as Tor. Both the US National Security Agency (NSA) and the UK surveillance agency Government Communications Headquarters (GCHQ) have initiated efforts to break the Tor network [2]. Multiple hidden Tor services like Silk Road 2.0 were shut down during Operation Onymous in November 2014, and documents have recently been disclosed that imply an NSA partnership with AT&T and Verizon to conduct Internet communication surveillance [3]. These developments have serious implications for the anonymity of Internet communication.

Anonymous communication depends on not identifying the originator's IP address and thus his or her location. In the Tor network, traffic is rerouted through several nodes: an entry node, which sends it to (typically one) relay node, which sends it to an exit node. The communication's origin is anonymous because the destination sees only the exit node's IP address.

Tor randomly selects exit nodes to hinder traffic-analysis attacks, but because it must minimize communication latency to avoid degrading performance, selection is not equally random, and thus does not produce uniformly distributed exit nodes. Instead, Tor weights the process of selecting exit nodes according to parameters such as the maximum number of exiting traffic streams pending. Consequently, the exit nodes actually used might be more heavily concentrated in a particular area or to a particular ISP.

To explore the consequences of this weighted random selection, we conducted an investigation based on 1.5 years of Tor data. Unlike previous Tor research efforts, we treated the Tor network as a black box and focused on identifying the information types that a large ISP or an intelligence service could gather. Our investigative results show a significant imbalance between the number of available exit nodes and those actually used. Moreover, many of the exit nodes used either belong to a small set of ISPs or are concentrated in a small area - characteristics that facilitate the collection of data on network traffic as part of a traffic-analysis or profiling attack. Consequently, the effects of exit-node distribution and selection could eventually erode network security and anonymity.

All authors are members of the Research Center CODE (Cyber Defence), Faculty of Computer Science, Universität der Bundeswehr München, D-85577 Neubiberg, Germany
E-mail: robert.koch@unibw.de

II. How Tor Works

Tor aims to eliminate the mapping between user and services or servers by hiding the user's IP address and thereby blocking user identification and communication tracking. To accomplish this, Tor generates an overlay network in which each node maintains a Transport Layer Security (TLS) connection to every other node [4]. Tor establishes a circuit—a random pathway through the network—by selecting entry, relay, and exit nodes.

Tor can extend the circuit by adding relay nodes, but a circuit typically has only one relay node so that communication latency remains at an acceptable level. To choose the exit node, Tor uses weighted random selection: it traverses the connection array and chooses a node to maximize the number of pending exit streams, optionally applying the exit node's required capacity and uptime as selection parameters. Section 5.3 of the Tor specification has more details (https://gitweb.torproject.org/torspec.git/tree/tor-spec.txt).

To avoid delays, Tor builds circuits preemptively and regularly as defined by the `NewCircuitPeriod` parameter, which defaults to every 30 seconds. Multiple TCP connections can share a Tor circuit. When an application sends a request to the Tor network, Tor attaches a new stream to an appropriate open circuit; if no existing circuit can handle the request, Tor launches a new one.

To avoid profiling attacks, Tor rotates used circuits regularly according to the `MaxCircuitDirtiness` parameter, which

TABLE I: Distribution of exit nodes in select countries from November 2013 to May 2015.

Country	No. of available exit nodes	Available exit nodes (%)	No. of times exit nodes were used	Share of all traffic exiting the Tor network (%)
Romania (RO)	42	1.67	893,728	7.15
Suisse (CH)	70	2.78	1,078,683	8.63
Russia (RU)	289	11.50	146,565	1.17
Netherlands (NL)	237	9.43	1,651,242	13.20
France (FR)	199	7.92	890,229	7.12
Germany (DE)	302	12.01	1,806,412	14.44
United States (US)	510	20.29	2,582,072	20.65

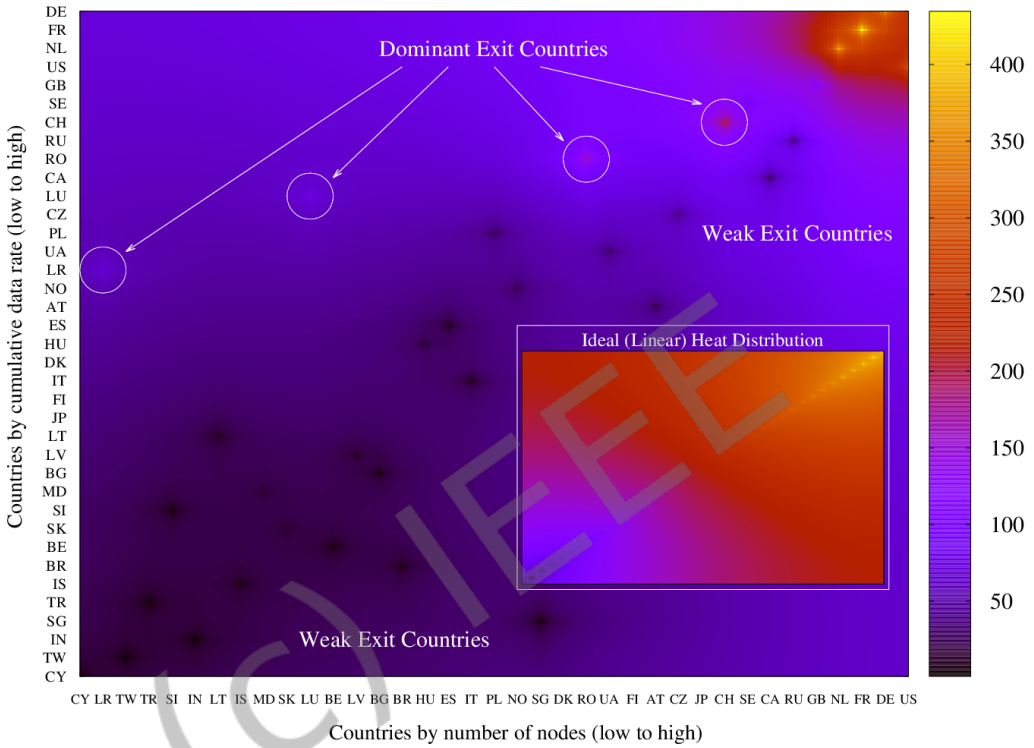

Fig. 1: Heat map of Tor exit-node use by country during three weeks of the investigation. Along the x-axis are countries ordered by number of Tor nodes in that country. Along the y-axis, country order is based on the nodes' cumulative data rate. Exit-node use is the number of times that Tor chose one of the country's nodes as an exit node. Dominant countries are overrepresented—their exit-node use is intense (yellow) and above average (orange). Weak countries (dark purple crosses) are underrepresented—their exit-node use is below average. The inset shows the ideal heat distribution, in which exit-node use is more intense overall.

defaults to every 10 minutes. Persistent single TCP streams, such as Internet Relay Chat (IRC) protocol connections, are not rotated and will remain on the same circuit indefinitely to counter profiling attacks (www.torproject.org/docs/faq.html.en#ChangePaths).

When sending data through Tor, the client encrypts it multiple times with the nodes' keys, including the predecessor's and successor's addresses for their respective nodes. Each node has the key only for one layer, uses the key to remove that layer, and then forwards the data. In this way, it sees only the IP address of where the packet came from and where it has to go. The exit node sends the packet to its final destination, which sees only the exit node's IP address. When the answer returns, each node adds its encryption layer and only the sender can finally remove them all and thus read the answer. This process is similar to peeling an onion (hence Tor's original name, The Onion Router).

III. Analyzing Exit-Node Selection and Use

To identify possible traffic-analysis and profiling attacks that exploit the selection and use of Tor exit nodes, we generated automated website requests and collected data about the exit nodes used and their characteristics from November 2013 to May 2015 (1.5 years). Our main goal was to determine the decrease in anonymity directly attributable to exit-node distribution and use.

We used the anonymizing relay monitor (arm; www.atagar.com/arm) to extract the IP addresses of used exit nodes and IPInfoDB - a combination of database and webservice that allows users to access IP geolocation information -to geolocate them by city. We used nslookup and dig to collect additional available information, such as DNS data, and whois to identify the ISPs that provided the IP addresses.

TABLE II: Exit-node distribution corresponding to Figure 1, actual exit-node use, and optimal exit-node use.

Country	Country's share of all available exit nodes (%)	Country's exit node share that was used (%)	Use of circuits exiting within the country (%)	Optimal use of circuits exiting within the country (%)	Aggregated % exits (%)
Liberia (LR)	0.27	100	2.79	0.02	3.34
Luxembourg (LU)	2.13	38.10	3.13	0.31	3.71
Singapore (SG)	0.27	1.64	0.04	0.90	0.04
Romania (RO)	3.47	19.70	6.69	0.97	8.15
Switzerland (CH)	3.73	10.45	8.67	1.97	10.54
Russia (RU)	3.73	4.46	0.74	4.62	0.83
Netherlands (NL)	12.53	9.69	15.63	7.13	18.29
France (FR)	15.2	8.93	17.64	9.38	6.77
Germany (DE)	8.53	2.22	14.16	21.21	16.34
United States (US)	24.0	5.53	12.07	23.92	13.60

To collect the data, we used Python and Bash scripts to set up parallel threads, each of which used a local configuration file and port to initialize its own Tor instance. A wrapper function in the script used wget to retrieve websites from established Tor network circuits. We then analyzed the exit nodes' IP addresses.

Because each circuit is used for 10 minutes and not rotated after each access, an exit node's IP address could be recorded multiple times, and the same exit node might be repeatedly selected. We deliberately allowed this to mimic real system and user behavior when the Tor browser is used to surf the Web. However, in our results, we count a used exit node's IP address only once; if numerous connections run over the same circuit and the exit IP address is counted multiple times, we label those "aggregated exits."

A. Global patterns

We began our investigation by identifying each country's role as a provider of Tor exit nodes during the 1.5 years, finding that 2,514 Tor exit nodes were used with more than 12.5 million aggregated exits. The recorded number of nodes is higher than the actual average number of Tor exit nodes because new nodes joined the Tor network during our observations, and others that we had counted were subsequently disabled.

As Table I shows, our investigation revealed some surprising imbalances in the degree to which countries used exit nodes. The US dominates the Tor network, hosting more than 20 percent of all Tor nodes; the percentage of used exit nodes is also 20 percent. Germany (DE) and France (FR) had similar percentages in the two categories. However, the percentages differ significantly for other countries. Russia (RU), for example, had nearly 12 percent of the available exit nodes (the seventh highest cumulative exit-node data rate), but only 1.17 percent of the connections ran over Russian Tor exit nodes. Consequently, Russia ended up being considerably underrepresented in the number of country-specific exit nodes used.

Because exit-node selection is weighted to favor nodes with higher capacity and data rate, it is natural to assume that a country's Internet data rate would heavily influence its exit-node use. However, our investigation showed many exceptions to that assumption. Sweden (SE), for example, provided very fast connections with an average speed of 16.1 Mbps, which was not reflected in that country's actual exit-node use. On the other hand, Liberia's (LR's) broadband connections are still below global average, yet we found an intense use of its single exit node. Countries like Luxembourg (LU) also had higher than expected exit-node use relative to their average Internet data rate, number of exit nodes, and exit-node cumulative data rate.

Although it might be tempting to attribute differences in exit-node use among countries to variations in the exit nodes' bandwidth, we found evidence of major deviations. To better visualize and understand deviations from logical assumptions and the findings of previous work, we generated a series of heat maps, such as that in Figure 1. Each map shows the distribution of exit nodes according to their number and use by applications exiting the Tor network during three weeks of our investigation period. We opted for a reduced dataset because displaying data for the entire 1.5 years would hide important details, such as nodes that were in the network for only a short time. Exit nodes regularly leave and join the network, and even a single node can severely affect the entire network. Liberia, for example, supported only one exit node during the three weeks, but that node was important enough to earn Liberia a dominant country ranking.

Moving from right to left along Figure 1's x-axis, during the three weeks portrayed, the US had the most Tor nodes (1,627), followed by Germany, France, the Netherlands (NL), Great Britain (GB), and Russia. However, the number of Tor nodes did not always determine a country's dominance in exit-node use. Liberia, the penultimate country in number of Tor nodes, had far fewer nodes than Norway (NO) at 56 or Denmark (DK) at 61, yet its capacity was higher than the cumulative node capacity in either of those countries. Thus, in terms of node capacity and the weighting in Tor's algorithm for selecting exit nodes, Liberia was one of the dominant countries—routing 3.34 percent of all application exits, according to these parameters.

The map also shows that countries like Luxembourg and Romania (RO) had intense use, while countries like Russia, which had the sixth highest number of exit nodes, are underrepresented.

B. Use by country versus optimal use

Table II shows the statistics we used to determine a country's expected average exit-node use. Circuit use clearly shows why particular countries were overrepresented or underrepresented. Exit nodes are always just that - exit nodes - but circuits

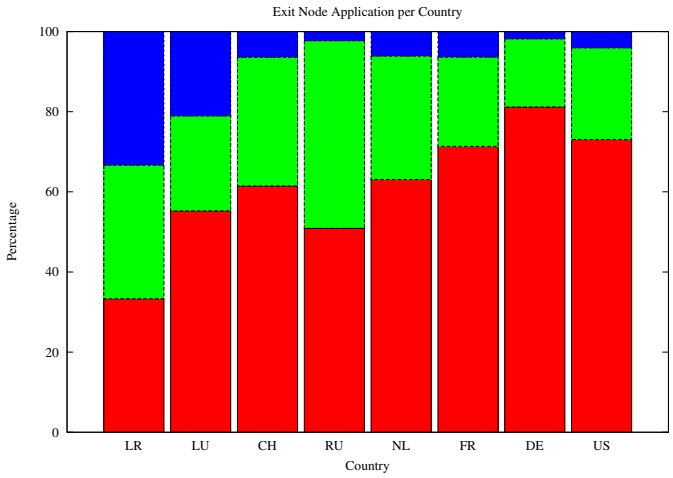

Fig. 2: Exit-node use per country relative to available exit nodes based on three weeks of data. Each bar shows the ratio of available Tor nodes (red) to nodes that Tor configured as exit nodes in that country (green) to used exit nodes, or ones that Tor selected (blue). Nearly a quarter of all nodes were located in the US, but Tor selected only 5.53 percent of these (blue section of US bar). Likewise, 8.53 percent of all exit nodes were located in Germany (green section of DE bar), but Tor selected only 2.22 percent of these (blue section of DE bar).

represent the path through the Tor network, which Tor builds by applying weighted random selection. Consequently, some exit nodes are practically never used, while others are used repeatedly.

In column 5 of Table II, the calculated optimal share is not necessarily optimal in Tor's weighted selection because the calculated share does not take the exit nodes' data rates into account. On the other hand, some organizations have sophisticated surveillance capabilities, so the broadest possible exit distribution would minimize observable traffic and hamper traffic-analysis attacks. Relative to a uniform distribution, actual circuit use (column 4 of Table II) in certain countries was considerably higher: 186 times in Liberia, 10.13 times in Luxembourg, 6.9 times in Romania, and 4.4 times in Switzerland (CH).

In contrast, the US - even with its considerable broadband capability - used only 0.5 times its optimal share - and Germany, with the second highest optimal share, used only 0.67 times. France was an interesting isolated case. Of all circuits constructed, 17.64 percent had exit nodes in France, which was 1.88 times its optimal share (9.38 percent), but its number of aggregated exits was quite low: only 6.77 percent of all application exits ran through French exit nodes. No other country demonstrated this behavior; typically, the number of aggregated exits was higher than the number of used circuits.

These unusual results led us to investigate how circuit stability relates to exit location. We subsequently determined that circuits ending in countries like Germany and the Netherlands were quite stable, typically running until they reached the `MaxCircuitDirtiness` value (10 min). However, circuits ending in France often showed an increased number of timed-out connections as well as early drops. On average, the circuit-use time in France was only about 7.5 min.

We found other imbalances in exit-node use per country. Countries like the US, France, and Germany provide numerous exit nodes, so theoretically exit-node use should be high. However, as Figure 2 shows, these countries actually used only a fraction of their available exit nodes.

A balanced distribution would reduce the risk of a traffic-analysis attack, in which every node is configured and used as an exit node, spreading the distribution and hampering network monitoring. The tradeoff is performance, because numerous nodes with low data rates must be used regularly, which slows Tor network traffic.

Moreover, configuring every available Tor node as an exit node is ideal for a country with many available Tor nodes, but some countries, such as Liberia, have only one or several, which greatly increases the risk of traffic monitoring. However, countries with a high number of Tor nodes could do more to even out their exit-node distribution. Germany, for example, has numerous Tor nodes (red in Figure 2) but only a small percentage is configured as exit nodes (green). Even more disconcerting is the fraction of exit nodes actually used (blue).

C. ISPs and exit-node use

The ratios in Figures 3 and 4 might explain the data pattern for some of the countries listed in Figure 2, such as Germany. According to a broad interpretation of German IT law, an exit-node provider can be held responsible for all illegal actions that the node executes. Not coincidentally, Figure 3 shows that only a few of Germany's ISPs support Tor exit nodes; indeed, we determined that a single ISP supported 52.3 percent of the exits based in Germany.

France is another country in which one ISP supported most of the exit nodes. France accounted for 6.77 percent of the aggregated traffic during our investigation, and 15.2 percent of all exit nodes were located there. However, only 8.93 percent

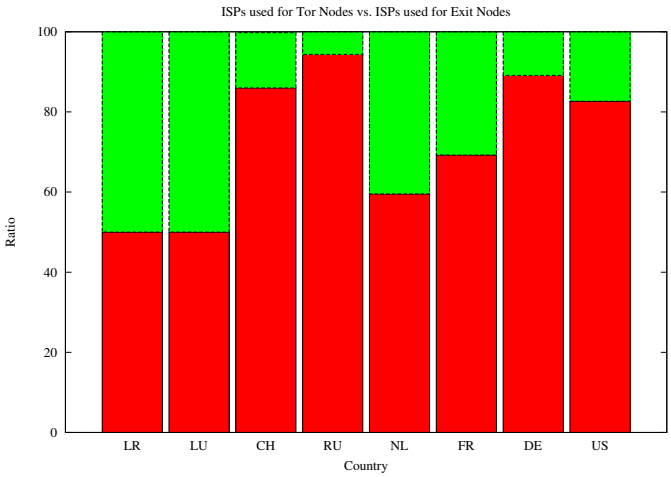

Fig. 3: Ratio of ISPs operating Tor nodes in a specific country (red) to the number of ISPs used for exit nodes (green) based on three weeks of data. For the most part, only a few ISPs per country are responsible for Tor exit nodes.

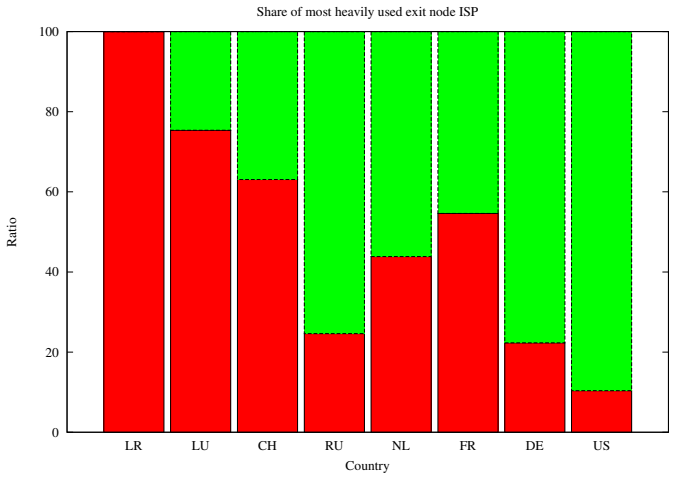

Fig. 4: Ratio of available exit nodes supported by the most heavily used ISP in that country (red) to all used Tor exit nodes in that country (green) based on three weeks of data. In countries like Luxembourg (LU), Switzerland (CH), and France (FR), single ISPs have a huge share of the overall exit traffic.

of France's exit nodes were used, and of those, 46.07 percent were supported by one ISP.

In the Netherlands, of the 10 percent of exit nodes used, 31.73 percent were supported by one multinational ISP, which also provided IP addresses for exit nodes running in other countries.

The data in Figure 3 suggests that in countries with many ISPs, only a fraction that route through the Tor network also support exit nodes - in essence, agglomerating Tor exit traffic. As Figure 4 shows, the US has the best distribution among ISPs; the dominant ISP routes only about 12 percent of all anonymized exit traffic.

D. Hot ISPs and exit points

In our investigation, we considered a "single ISP" as one that supported all Tor exit nodes used in that country. A single ISP could be the only ISP that supported Tor exit nodes in that country during our evaluation period; however, it could also be one of many ISPs in that country but the only one that supported the exit nodes on which all Tor exit traffic ran during that time. Thus, single ISPs were in smaller countries such as Liberia but also in larger countries such as Italy or Spain.

Some single ISPs provided multiple exit IP addresses, which in effect combined a sizable share of exit traffic. For example, traffic during our investigation period used 6 exit nodes in Portugal (PT) and 10 in Namibia (NA), each belonging to one ISP.

We labeled single ISPs that combined traffic in this way as "hot ISPs," and our investigation revealed many of these. In France, for example, one ISP supported 61 exit nodes - a number that translates to 2.96 percent of all exit traffic and more than 50 percent of all traffic exiting in France (calculated from three weeks of data not shown). In the US, one ISP hosted 37 exit IP addresses. In the Netherlands, 19 exit IP addresses belonged to a single ISP, which also provided exit IP addresses in the US as well as in Great Britain, Belgium (BE), and several other EU countries. One German ISP routed about 9.23 percent of all exit traffic during three weeks. Regrettably, these hot ISPs are not outliers, but are typical of what we saw in the Tor data we collected.

High-capacity exit nodes agglomerate exiting Tor traffic, which creates hot exit points—an intense use of individual exit nodes relative to the number of aggregated exits. Hot exit points transport excessive amounts of traffic, greatly increasing the risk of traffic-analysis attacks. The differences in actual and optimal circuit use reveal the presence of hot exit points in various countries (columns 4 and 5 in Table II).

A hot ISP could also be a hot exit point. For example, one ISP aggregated 19 exit nodes with 5.92 percent of the aggregated exits in the Netherlands. The combination of hot ISPs and organizations that create hot exit points makes a traffic-analysis attack even easier.

E. Hot regions

A geographical area can also be a source of hot exit points, particularly those areas with organizations and academic institutions that are conducting Tor-related research. Also, students in general are often willing to provide exit nodes to support their research activities or private communication, which can result in a high number of locally aggregated exits.

At a campus where numerous students are running Tor exit nodes, local concentration is obvious. However, we also noted clusters of Tor exit nodes in particular areas not identified as related to Tor research. Figure 5 shows Baar, a village in Switzerland, the location of 70 percent of the country's exit nodes. In the Netherlands, more than 80 percent of all exit nodes were located in Amsterdam. In Luxembourg, seven exit nodes were used, transporting about 4 percent of all exit traffic; six of these nodes belonged to the same autonomous system (AS) that was routing 98.74 percent of Luxembourg's exit traffic.

IV. Related Investigations on Tor Anonymity

Our work compares favorably with other efforts to explore the strength of anonymity in the Tor network. One group who also evaluated how Tor's path selection related to anonymity used a simulated Tor network running on PlanetLab to examine load and estimated maximum node capacity [5]. Another group analyzed Tor's selection of entry guards and concluded that clients use more than they should, which increases the likelihood of profiling attacks [6]. (Entry guards are a small subset of entry nodes that Tor randomly selects from all entry nodes that exist when a Tor client initializes. The entry guards remain for an extended period, even when relay and exit nodes change; www.torproject.org/docs/faq#EntryGuards.)

In other work, researchers noted that Tor might select a single AS for both entry and exit nodes, which would enable the monitoring of two ends of an anonymous communication path. To decrease this threat, they proposed an AS-aware path-selection algorithm [7].

On effort looked at how users might tune Tor to boost the performance or anonymity level by replacing self-reported values with an opportunistic bandwidth-measurement algorithm that users could tailor to the desired levels [8].

Few efforts have focused, as we did, on how applications use Tor's exit nodes. One 2013 paper describes an evaluation of providers and IP addresses [9] taken from the exit IP address lists available through the Tor Project. The authors merged locations of IP addresses and affiliate organizations with an attack IP address list from DShield - a system from the SANS Internet Storm Center that collects volunteer logs to analyze attack trends. They concluded that exit nodes are located in many countries and belong to many organizations. Our results show far more uneven distributions.

We believe our results differ from those reported in the 2013 paper for two main reasons. First, the other study was only for October 2012 - one month versus our 1.5 years. Second, the authors analyzed the published IP addresses only with regard to their distribution and location, not their application use. This limited view yields only a static display of exit-node distribution; it does not consider the actual use of exit nodes, which depends heavily on the path algorithm and its attributes. In contrast, our work analyzes the implications of how used exit nodes are distributed.

In recent work, another group concluded that Tor's path-selection algorithm is not ideal [10]. However, they did not examine exit-node selection by running real-world applications and requests over Tor as we did. They also neglected the role of ISPs in traffic agglomeration. Because of the increasing capabilities of intelligence services or monitoring possibilities of huge ISPs, the distribution and selection of exit nodes can seriously affect the risk of traffic-analysis and profiling attacks.

V. Conclusion

Our long-term analysis of Tor exit-node selection and use and the role of ISPs in agglomerating exit traffic identified the potential for expert surveillance entities such as intelligence services to exploit hot exit points. In contrast to earlier work that found little impact from exit-node distribution, we pinpointed numerous agglomerations in which multiple exit IP addresses belong to the same ISP or organization, and we identified large percentages of exit IP addresses within the same city.

Our heat maps make it easy to visualize countries in which exit-node use is significantly underrepresented or overrepresented. As combinations of hot ISPs, hot exit points, and hot regions emerge, the risk of surveillance and traffic analysis grows alarmingly, with a higher potential for cooperation between intelligence services and telecommunication providers.

These combinations strongly limit the Tor exit nodes that can be used with confidence in their anonymity. However, decreasing the already small number of exit nodes will end up pushing the lion's share of anonymized traffic into very few organizations and regions, which will only make traffic analysis easier. Recent discussion about the role of Carnegie Mellon University and the FBI in the 2014 Tor attacks addresses the danger of compromised exit-node selection (https://blog.torproject.org/blog/did-fbi-pay-university-attack-tor-users).

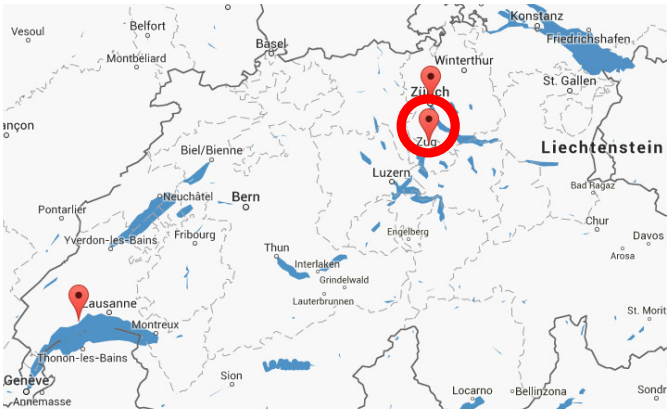

Fig. 5: Distribution of exit nodes in Switzerland based on three weeks of data. Baar (circled) holds more than 70 percent of all exit nodes in the country. The remaining two markers point to other node clusters. Restricting exit-node distribution to select areas increases the risk of a traffic-analysis attack.

To reduce the risk of surveilled exit nodes, we recommend applying end-to-end encryption with SSL/TLS to access websites. Although this strategy will impede content eavesdropping, it cannot guard against a deanonymization attack by someone who can monitor multiple exit nodes. The combined use of Tor with additional proxies or other anonymizing networks such as I2P might help in countering such an attack.

Acknowledgements

This work was partly funded by FLAMINGO, a Network of Excellence project (ICT-318488) supported by the European Commission under its Seventh Framework Programme.

References

[1] C. Bitso, I. Fourie, and T. J. Bothma, "Trends in transition from classical censorship to internet censorship: selected country overviews," *Innovation: journal of appropriate librarianship and information work in Southern Africa: Information Ethics*, no. 46, pp. 166–191, 2013.

[2] Greenwald, Glenn, "NSA and GCHQ target Tor network that protects anonymity of web users," http://www.theguardian.com/world/2013/oct/04/nsa-gchq-attack-tor-network-encryption, last visited on 29.03.2015.

[3] The New York Times, "Newly Disclosed N.S.A. Files Detail Partnerships With AT&T and Verizon," http://www.nytimes.com/interactive/2015/08/15/us/documents.html, last visited on 25.08.2015.

[4] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," DTIC Document, Tech. Rep., 2004.

[5] A. Panchenko, F. Lanze, and T. Engel, "Improving performance and anonymity in the tor network," in *Performance Computing and Communications Conference (IPCCC), 2012 IEEE 31st International*. IEEE, 2012, pp. 1–10.

[6] T. Elahi, K. Bauer, M. AlSabah, R. Dingledine, and I. Goldberg, "Changing of the guards: A framework for understanding and improving entry guard selection in tor," in *Proceedings of the 2012 ACM workshop on Privacy in the electronic society*. ACM, 2012, pp. 43–54.

[7] M. Edman and P. Syverson, "As-awareness in tor path selection," in *Proceedings of the 16th ACM conference on Computer and communications security*. ACM, 2009, pp. 380–389.

[8] R. Snader and N. Borisov, "A tune-up for tor: Improving security and performance in the tor network." in *NDSS*, vol. 8, 2008, p. 127.

[9] A. Schaap, "Characterization of tor exit-nodes," *Proc. 18th Twente Student Conf.*, 2013, http://referaat.cs.utwente.nl/conference/18/paper/7381/characterization-of-tor-exit-nodes.pdf.

[10] M. Backes, A. Kate, S. Meiser, and E. Mohammadi, "(nothing else) mator (s): Monitoring the anonymity of tor's path selection," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 513–524.