

Nation-State Hegemony in Internet Routing

Anne Edmundson
Princeton University

Nick Feamster
Princeton University

Roya Ensafi
University of Michigan

Jennifer Rexford
Princeton University

ABSTRACT

While the growth of the Internet has fostered more efficient communications around the world, there is a large digital divide between Western countries and the rest of the world. Countries such as Brazil, China, and Saudi Arabia have questioned and criticized America's Internet hegemony. This paper studies the extent to which various countries rely on the United States and other Western countries to connect to popular Internet destinations in those countries. Unfortunately, our measurements reveal that underserved regions are dependent on North American and Western European regions for two reasons: local content is often hosted in foreign countries (such as the United States and the Netherlands), and networks within a country often fail to peer with one another. Fortunately, we also find that routing traffic through strategically placed relay nodes can in some cases reduce the number of transnational routing detours by more than a factor of two, which subsequently reduces the dependence of underserved regions on other regions. Based on these findings, we design and implement Region-Aware Networking, RAN, a lightweight system that routes a client's web traffic around specified countries with no modifications to client software (and in many cases with little performance overhead).

ACM Reference Format:

Anne Edmundson, Roya Ensafi, Nick Feamster, and Jennifer Rexford. 2018. Nation-State Hegemony in Internet Routing. In *COMPASS '18: ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS)*, June 20–22, 2018, Menlo Park and San Jose, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3209811.3211887>

1 INTRODUCTION

As the Internet continues to grow, the increasing social, economic, and cultural hegemony of Western regions over the rest of the world have led to a digital divide. This divide inhibits connectivity, transparency, and the equal exchange of ideas [14]. In recent years, we have seen various countries and regions, such as Brazil, China, and Saudi Arabia, push back against the United States hegemony in the global Internet [18, 30, 37].

This paper specifically studies the routing differences between American/Western European regions and underserved regions by

measuring and analyzing the international routing detours exhibited when accessing popular content. The consequences of these international routing detours include performance degradation, increased costs, surveillance, and censorship. Previous work has shown that *tromboning* paths—paths that start and end in the same country, but also traverse a foreign country—are common [24, 47], especially in underserved regions.

In this paper, we study two questions: (1) Which countries do *default* Internet routing paths traverse?; (2) What methods can help governments (or citizens, ISPs, etc.) better control transnational Internet paths and reduce dependence on the United States and Europe to transit Internet traffic? We *actively measure* the paths originating in five countries to the most popular websites in each of these respective countries. Our analysis focuses on five countries—Brazil, Kenya, India, the Netherlands, and the United States. Brazil, Kenya, and India are representative of underserved regions in the world, and the Netherlands and the United States represent more dominant global powers.

In contrast to previous work, we measure router-level forwarding paths that traffic actually traverses, as opposed to analyzing Border Gateway Protocol (BGP) routes [31, 47], which can provide at best an indirect estimate of country-level paths to sites. Although BGP routing can offer some information about paths, it does not necessarily reflect the path that traffic actually takes, and it only provides AS-level granularity, which is often too coarse to make strong statements about which countries that traffic is traversing. In contrast, we measure routes from RIPE Atlas probes [45] in each country to the Alexa Top 100 domains for each country; we directly measure the paths not only to the websites corresponding to themselves, but also to the sites hosting any third-party content on each of these sites.

Although using direct measurements provides these benefits, there are a number of challenges associated with determining which countries a client's traffic is traversing. First, performing direct measurements is more costly than passive analysis of BGP routing tables; RIPE Atlas, in particular, limits the rate at which one can perform measurements. As a result, we had to be strategic about the origins and destinations that we selected for our study. We study five geographically diverse countries, focusing on countries in each region that are either underserved or more dominant in the global Internet. Second, IP geolocation—the process of determining the geographic location of an IP address—is notoriously challenging, particularly for IP addresses that represent Internet infrastructure, rather than end-hosts [21]. We address this inaccuracy by making conservative estimates of the extent of routing detours, and by recognizing that our goal is not to pinpoint a precise location for an IP address as much as to achieve accurate reports of *significant* off-path detours to certain countries or regions. (Section 3 explains

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

COMPASS '18, June 20–22, 2018, Menlo Park and San Jose, CA, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5816-3/18/06...\$15.00

<https://doi.org/10.1145/3209811.3211887>

our method in more detail; we also explicitly highlight ambiguities in our results.) Finally, the asymmetry of Internet paths can also make it difficult to analyze the countries that traffic traverses on the reverse path from server to client; our study finds that country-level paths are often asymmetric, and, as such, our findings represent a lower bound on transnational routing detours.

We first *characterize the current state of transnational Internet routing detours* (Section 3). We explore hosting diversity by first measuring the Alexa Top 1000 domains and comparing the location of path endpoints to that of the Alexa Top 100 domains; we find that there is no significant difference between the results in the two domain sets, and therefore focus on the Alexa Top 100 domains *and all third party domains*. We find that only 45% of the Alexa Top 100 domains in Brazil are hosted in more than one country (other countries studied showed similar results); in many cases, that country is one that clients may want to avoid. Second, even if hosting diversity can be improved, routing can still force traffic through a small set of countries. Despite strong efforts made by some countries to ensure their traffic does not transit certain countries [7–10, 27], their traffic still does so. For example, over 50% of the top domains in Brazil and India are hosted in the United States, and over 50% of the paths from the Netherlands to the top domains transit the United States. About half of Kenyan paths to the top domains traverse the United States and Great Britain (but the same half does not traverse both countries). Much of this phenomenon is due to “tromboning”, whereby an Internet path starts and ends in the same country, yet transits an intermediate country; for example, about 13% of the paths that we explored from Brazil tromboned through the United States. Infrastructure alone is not enough. ISPs in respective regions need better incentives to interconnect with one another to ensure that local traffic stays local.

Second, we explore the extent to which clients can avoid relying on certain countries to popular destinations by using overlay network relays to route Internet traffic around a given country (Section 4). Our results demonstrate that this technique can be effective for clients in certain countries; of course, the effectiveness of this approach naturally depends on where content is hosted for that country and the diversity of Internet paths between ISPs in that country and the respective hosting sites. For example, our results show that clients in Brazil can completely avoid Spain, Italy, France, Great Britain, Argentina, and Ireland (among others), even though the default paths to many popular Brazilian sites traverse these countries. We also find that some of the most independent regions are also some of the least avoidable regions. For example, many countries depend on ISPs in the United States for connectivity to popular sites and content. Additionally, overlay network relays can increase performance by keeping local traffic local: by using relays in the country, fewer paths trombone out of the client’s country.

Finally, we *design, implement, and deploy Region-Aware Networking, RAN, a system of overlay network relays that allows a client to access web content while minimizing the her dependence on a specified country* (Section 5). We implemented RAN for end-users, but ISPs can also deploy RAN proxies to gain more routing independence as a service to its customers. Our evaluation shows that RAN can effectively route around many different countries and introduces minimal performance overhead.

2 RELATED WORK

Nation-state routing analysis. Shah and Papadopoulos recently measured international routing detours—paths that originate in one country, cross international borders, and then return to the original country—using public Border Gateway Protocol (BGP) routing tables [47]. The study discovered 2 million detours each month out of 7 billion paths. Our work differs by *actively* measuring Internet paths using traceroute, yielding a more precise (and accurate) measurement of the paths, as opposed to analyzing BGP routes. Obar and Clement analyzed traceroutes that started and ended in Canada, but tromboned through the United States [38]. Karlin *et al.* developed a framework for country-level routing analysis to study how much influence each country has over interdomain routing [31]. This work measures country centrality using BGP routes and AS-path inference; in contrast, our work uses active measurements and measures avoidability of a given country. Several studies have also characterized network paths *within* a country, including Germany [49, 50] and China [52], or a country’s interconnectivity [5, 19, 24, 46]; these studies focus on paths within a country, as opposed to paths that traverse multiple countries.

Routing overlays and Internet architectures. Alibi Routing uses round-trip times to prove that that a client’s packets did not traverse a forbidden country or region [34, 51]; RAN differs by measuring which countries a client’s packets traverse. Our work uses active measurements to determine the best path for a client wishing to connect to a server, whereas Alibi Routing uses the speed of light and circle distance to calculate minimum RTTs for a packet to traverse a region. Because Alibi Routing uses this technique, it is almost impossible for a client to provably avoid a neighboring region or country. Lastly, we see RAN as a compliment to Alibi Routing; Alibi Routing may result in no paths *provably* avoiding a given region, but that does not mean there are no paths that avoid a given region. RAN is a good alternative to determine if there are any paths, based on active measurements, that do not traverse a given region. Recently, researchers proposed DeTor, which applies Alibi Routing techniques to Tor [15], to prove that a Tor circuit does not traverse a forbidden region [51]. As DeTor uses Alibi Routing techniques, it suffers from the same limitations as these techniques; because DeTor is applied to the Tor network, it raises the chances of overloading certain relays that are in a position to be on many circuits for clients wishing to avoid certain countries. Additionally, DeTor fails to consider reverse paths in the calculation of region avoidance, causing the results to state provable avoidance, when the forbidden region may actually be on the reverse path. RON, Resilient Overlay Network, is an overlay network that routes around failures [1], whereas our overlay network routes around countries. ARROW introduces a model that allows users to route around ISPs [40], but requires ISP participation, making it considerably more difficult to deploy than RAN; additionally, ISPs may span multiple countries. Zhang *et al.* presented SCION, a “clean-slate” Internet architecture that provides route control, but requires fundamental changes to the Internet architecture.




Figure 1: Measurement pipeline to study Internet paths from countries to popular domains.

3 CHARACTERIZING DETOURS

We describe our measurement methods, the challenges in conducting them, and our findings concerning the transnational detours of default Internet paths.

3.1 Measurement Approach

Figure 1 shows the process that we use to discover end-to-end Internet paths from our respective vantage points to various domains. We first use VPNs to establish various vantage points in the countries from which we measure; then, we use `curl` to download corresponding webpages for each of those popular domains, including all subdomains that are embedded in the site’s top-level webpage (1,2). We extract all of these domain names (3) and resolve them to their corresponding IP addresses (4); we then perform traceroutes to each IP address (5). Figure 2 describes how we translate an IP-level traceroute to a country-level path. We geolocate each IP address, removing unknown hops; we then de-duplicate the country-level path. Although it is seemingly straightforward, this approach entails a number of limitations and caveats, which we describe in the rest of this section.

3.1.1 Resource Limitations. We focus on five countries due to resource limitations. The iPlane [35] and Center for Applied Internet Data Analysis (CAIDA) [12] projects maintain large repositories of traceroute data, neither of which are suitable for our study. iPlane has historical data as far back as 2006. Unfortunately, because iPlane uses PlanetLab [41] nodes, which are primarily hosted on the Global Research and Education Network (GREN), iPlane measurements may not be representative of typical Internet users’ traffic paths [4]. CAIDA runs traceroutes from different vantage points around the world to randomized destination IP addresses that cover all /24s; in contrast, we focus on paths to popular websites from a particular country.

We run active measurements that better represent paths of a typical Internet user. To do so, we run DNS and traceroute measurements from RIPE Atlas probes, which are hosted all around the world in many different types of networks, including home networks [45]. RIPE Atlas probes can use the local DNS resolver, which provides the router-level path to a destination that a user is likely to see in that country.

Conducting measurements from a RIPE Atlas probe costs a certain amount of “credits”, which restricts the number of measurements that we can run. RIPE Atlas also imposes rate limits on the number of concurrent measurements and the number of credits that an individual user can spend per day. We address these challenges in two ways: (1) we reduce the number of measurements we must

run on RIPE Atlas probes by conducting traceroute measurements to a single IP address in each /24 (as opposed to all IP addresses returned by DNS) because all IP addresses in a /24 belong to the same AS, and should therefore be located in the same geographic area; (2) we use a different method—VPN connections—to obtain a vantage point within a foreign country, which is still representative of an Internet user in that country. We are forced to use an alternative vantage point to RIPE Atlas probes because these probes do not support all operations that our methods require (such as requesting the webpage). Although VPN connections provide the necessary functionality in the correct country, RIPE Atlas probes are more representative of typical Internet users, as they are often hosted in home networks, therefore we decide to use RIPE Atlas probes when possible and VPN connections when necessary.

3.1.2 Path Asymmetry. The reverse path (*i.e.*, the path from the server to the client) is just as important as (and often different from) the forward path. Previous work has shown that paths between Internet endpoints are often asymmetric [26]. Most work on path asymmetry has been done at the AS level [20, 25, 26, 39], but not at the country level; our measurements can consider only the forward path (from client to domain or relay), not the reverse path from the domain or relay to the client.

We also (separately) measured path asymmetry at the country granularity. If country-level paths were symmetric, then the results of our measurements would be representative of the forward *and* reverse paths. If the country-level paths are asymmetric, then our measurement results only provide a lower bound on the number of countries that traffic between two endpoints may traverse. Using 100 RIPE Atlas probes and eight Amazon EC2 instances, we ran traceroute measurements from every probe to every EC2 instance and from every EC2 instance to every probe (the EC2 instances were located in the United States, Brazil, Canada, Ireland, Germany, Japan, Australia, and Singapore). After mapping the IP addresses to countries, we analyzed the paths for symmetry. First, we compared the set of countries on the forward path to the set of countries on the reverse path; we found that about 30% of the paths were symmetric at the country level. We compared the number of countries on the forward and reverse paths to determine how many reverse paths were a subset of the respective forward path; this situation occurred for 55% of the paths. This level of asymmetry suggests that our results are a lower bound on how many countries a client’s path traverses en route to a web site. It also suggests that while providing lower bounds on transnational detours is feasible, designing systems to *completely* prevent these detours on both forward and reverse paths is challenging. If tools that shed light on




Figure 2: Mapping country-level paths from traceroutes.

the reverse path between endpoints (e.g., Reverse Traceroute [33]) see more widespread deployment, the characterizations and avoidance techniques that we develop in this paper could be extended to include reverse paths.

3.1.3 Traceroute Origin and Destination Selection. Most of the countries studied host 75 to several hundred RIPE Atlas probes. Because of resource restrictions, we could not use all of the probes in each country. We selected the set of probes that had unique ASes in the country to get the widest representation of origination points.

To determine how many websites we must measure to sufficiently capture client paths to popular websites in a country, we first compare the country-level paths from a small set of vantage points to the Alexa Top 100 domains and to the Alexa Top 1000 domains. The proportion of paths that transited (and ended in) each country are similar in both cases; the paths to the top 1000 domains exhibit a longer tail of countries that transit or host content, likely because these domains are less popular and therefore hosted in more obscure locations. Otherwise, the results are similar. Figure 3 shows this comparison (for simplicity, we have removed the long tail of countries that are the endpoint for less than 1% of the measured paths). Due to the similarity of characteristics between the Alexa Top 100 and Alexa Top 1000 and because of resource constraints, we used the Alexa Top 100 domains in each of the respective countries as our destinations, as well as the third-party domains that are requested as part of an original web request.

To obtain the third-party domains that are hosted on each popular website, we use curl to retrieve the homepage for each respective domain from within the country that is hosting the vantage point in question. RIPE Atlas probes do not support these types of Web requests; instead, we establish a VPN connection within each of these countries to curl each domain and extract the third-party domains; we curl from the client’s location in case web sites are customizing content based on the region of the client.

3.1.4 Inferring Country-Level Paths. Accurate IP geolocation is challenging [16, 17, 22, 23, 28, 32, 42]. We use MaxMind’s geolocation service to map IP addresses to their respective countries [36]. This database is known to contain inaccuracies, particularly for IP addresses that correspond to Internet infrastructure, as opposed to end hosts; fortunately, our aim is to discover the countries that Internet paths traverse, and previous work has found that geolocation at a country-level granularity is more accurate than it is at finer granularities [29]. In light of these concerns, we post-processed our IP to country mapping, as shown in Figure 2. The method first

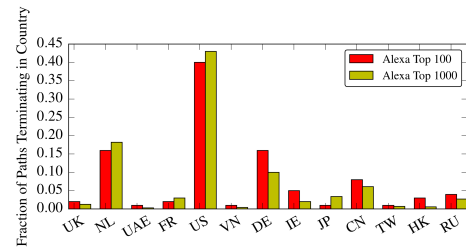


Figure 3: Comparison of path endpoints between the Alexa Top 100 and 1000.

Terminating in Country	Brazil	Netherlands	India	Kenya	United States
Brazil	.169	-	-	-	-
Canada	.001	.007	.015	.006	-
United States	.774	.454	.629	.443	.969
France	.001	.022	.009	.023	.001
Germany	.002	.013	.014	.028	.001
Great Britain	-	.019	.021	.032	.002
Ireland	.016	.064	.027	.108	.001
Netherlands	.013	.392	.101	.200	.024
Spain	.001	-	-	-	-
Kenya	-	-	-	.022	-
Mauritius	-	-	-	.004	-
South Africa	-	-	-	.021	-
United Arab Emirates	-	-	-	.011	-
India	-	-	.053	.002	-
Singapore	-	.002	.103	.027	-

Table 1: Fraction of paths terminating in a country.

removes all IP addresses that resulted in a ‘None’ response when querying MaxMind, which causes our results to provide a conservative estimate of the number of countries that paths traverse. Note that removing ‘None’ responses will always produce a conservative estimate.

3.1.5 Traceroute Accuracy and Completeness. Our study is limited by the accuracy and completeness of traceroute. Anomalies can occur in traceroute measurements [2], but most traceroute anomalies do not cause an overestimation in the number of countries on a path. The incompleteness of traceroutes, where a router does not respond, causes our results to underestimate the number of states that interfere with network traffic.

3.2 Results

Table 1 shows five of the countries that we studied along the top of the table and the countries that host their content along in each row. A “-” represents the case where no paths ended in that country. For example, the United States is the endpoint of 77.4% of the paths that originate in Brazil, and no Brazilian paths terminated in South Africa. Table 2 shows the fraction of paths that transit (or end in) certain countries, with a row for each country that is transited. We report on measurements conducted on January 31, 2016, and we are continuing to run these measurements and publish the data. We have published our data to an anonymized repository [13].

Transiting Country	Brazil	Netherlands	India	Kenya	United States
Brazil	1.00	-	-	-	-
Canada	.013	.007	.016	.008	.081
United States	.844	.583	.715	.616	1.00
France	.059	.102	.104	.221	.104
Germany	.005	.050	.032	.048	.008
Great Britain	.024	.140	.204	.500	.006
Ireland	.028	.106	.031	.133	.006
Netherlands	.019	1.00	.121	.253	.031
Spain	.176	.004	-	-	-
Kenya	-	-	-	1.00	-
Mauritius	-	-	-	.322	-
South Africa	-	-	-	.334	-
United Arab Emirates	-	-	-	.152	-
India	-	-	1.00	.058	-
Singapore	-	.002	.270	.040	.003

Table 2: Fraction of paths that a country transits.





Figure 4: Fraction of country code top-level domains that are hosted locally.

Hosting Diversity. Hosting diversity reflects how many unique countries host a domain. The more countries host a domain, the greater the likelihood that a client can find a path to that site that avoids a certain country. As a separate measurement experiment, we performed DNS queries for the sites we measure from 26 vantage points around the world, in geographically diverse locations. We then mapped the IP addresses in the DNS responses to countries to determine how many unique countries host a domain. We found about half of the top domains in each of the five countries studied are hosted in a single country and the other half are located in two or more different countries; this represents two cases: (1) CDNs and (2) a single hosting country. This shows that many domains are hosted in a single unique country, which leads us to our next analysis—where are these websites hosted, and which countries are traversed on the way to reach these locations.

Domain Hosting. Table 1 shows the fraction of paths that are hosted in various countries. Despite the extent of country-level hosting diversity, the majority of paths from all of the countries we studied terminate in a single country; 77%, 45%, 63%, 44%, and 97% of paths originating in Brazil, Netherlands, India, Kenya, and the United States, respectively, are currently reaching content located in the United States. Our results also show the Netherlands is a common hosting location for paths originating in the Netherlands, India, and Kenya.



(a) Brazil. (b) Kenya.

Figure 5: Countries that tromboning paths transit.

Domestic Traffic. All of the countries we studied (except for the U.S.) host content for a small percentage of the paths that originate in their own country; they also host a small percentage of their respective country-code top-level domains. Only 17% of paths that originate in Brazil also end there, and only 5% and 2% of Indian and Kenyan paths, respectively, end in the originating country. For Kenya, 24 out of the Top 100 Domains are .ke domains, but only 5 of the 24 are hosted within Kenya. 29 out of 40 .nl domains are hosted in the Netherlands; four of 13 .in domains are hosted in India; 18 of 39 .br domains are hosted in Brazil. Figure 4 shows these results. As one might expect, all .gov domains were hosted in their respective country.

Transit Traffic. The United States and Great Britain are on more paths than any other (foreign) country. 84% of Brazilian paths traverse the United States, despite Brazil’s strong efforts to avoid United States surveillance [6–11]. Although India and Kenya are both geographically far from the United States, 72% and 62% of their paths nonetheless transit the United States.

Great Britain and the Netherlands are on many of the paths from Kenya and India: 50% and 20% of paths that originate in Kenya and India, respectively, transit Great Britain. Many paths likely traverse Great Britain and the Netherlands due to the presence of large Internet Exchange Points (*i.e.*, LINX, AMS-IX). Mauritius, South Africa, and the United Arab Emirates transit 32%, 33%, and 15% of paths from Kenya. There are direct underwater cables from Kenya to Mauritius, and from Mauritius to South Africa [48].

Tromboning Traffic. Brazilian and the Netherlands paths often trombone to the U.S., despite the prevalence of IXPs in both countries. Figure 5 shows the fraction of paths that trombone to different countries for Brazil and Kenya. 24% of all paths originating in the Netherlands (62% of domestic paths) trombone to a foreign country before returning to the Netherlands. Despite Brazil’s strong efforts in building IXPs to keep local traffic local, their paths still trombone to the U.S. This is due to IXPs being seen as a threat by competing commercial providers; providers are sometimes concerned that interconnection will result in making business cheaper for competitors and stealing of customers [43].

Brazilian ISPs have often viewed one another as competitors and therefore as a threat at IXPs; this policy artifact causes Brazilian ISPs to peer with international providers instead of other local providers [43]. Additionally, we see Brazilian paths trombone to Spain and Italy. We see Italy often in tromboning paths because Telecom Italia Sparkle is one of the top global Internet providers [3]. MaxMind’s geolocation sometimes mislabels IP addresses to be

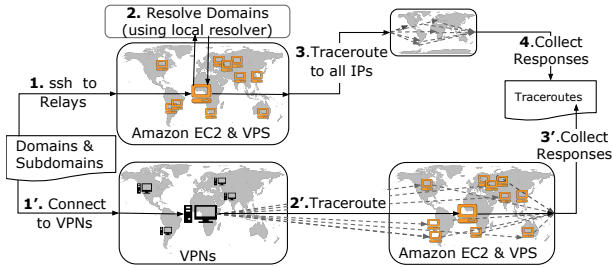


Figure 6: Measurement approach for country avoidance.

in Spain when they are actually located in Portugal. Despite our inability to disambiguate Spain and Portugal, some issues associated with tromboning, such as performance, are still pertinent. We are not aware of specific laws in either of these countries that would make this distinction important from a legal aspect, either.

Tromboning paths that originate in Kenya most commonly traverse Mauritius, which is expected considering the submarine cables between Kenya and Mauritius. Additionally, a cable from Mombasa, Kenya to Fujairah, United Arab Emirates likely explains why many paths include these countries.

Most U.S. Content Stays Local. Brazilian, Dutch, Indian, and Kenyan paths often transit the U.S. The results from studying paths that originate in the United States are drastically different from those of the other four countries. The majority of locally popular content in these countries is hosted outside of the respective country, which is shown in Table 1; in contrast, the United States hosts 97% of the content that is accessed from within the country. Only 13 unique countries are ever on a path from the U.S. to a webpage in our dataset, whereas 30, 30, 25, and 38 unique countries are seen on the paths originating in Brazil, Netherlands, India, and Kenya, respectively.

4 FEASIBILITY OF REGION-AWARE NETWORKING

We now explore the extent to which overlay networks can improve path diversity and help clients route around specific countries. We develop an avoidance metric and algorithm, and evaluate the effectiveness of overlay nodes to avoid specific countries.

4.1 Measurement Approach

An overlay network of relay nodes can help clients route around countries or access content that is hosted in a different country; this section performs measurements to evaluate the feasibility of such an approach. Figure 6 shows the steps in our measurement experiment. After selecting potential relay nodes, we perform traceroute measurements from the country of origin to each relay (1',2'), and from each relay to the set of top 100 domains in the original country (1,2,3). We then analyze these traceroutes using the approach shown in Figure 2 to determine the resulting country-level paths.

We use eight EC2 instances, one in each region (United States, Ireland, Germany, Singapore, South Korea, Japan, Australia, Brazil), as well as four Virtual Private Server (VPS) machines (France, Spain, Brazil, Singapore), which are virtual machines. Combining these two sets of machines allows us to evaluate country avoidance with

a diverse set of relays. For our measurements, we required root access to the servers that run as relays; for this reason, we used EC2 instances and VPS machines. This requirement also restricted the number of relays we could set up to measure country avoidance. Our results reflect this number of relays and their respective locations, but more relays in more diverse locations will provide for more country avoidability.

4.2 Avoidability Metrics

We introduce a new metric, avoidability, to measure how often a client in one country can avoid another specific country. Using the proposed metric and algorithm, we can compare how well the different methods achieve country avoidance for any (X, Y) pair.

Avoidability metric. We introduce an avoidability metric to quantify how often traffic can avoid Country Y when it originates in Country X. Avoidability reflects the fraction of paths that originate in Country X and do not transit Country Y. We calculate this value by dividing the number of paths from Country X to domains that do not traverse Country Y by the total number of paths from Country X. The resulting value is in the range [0,1], where 0 means the country is unavoidable for all of the domains in our study, and 1 means the client can avoid Country Y for all domains in our study. For example, there are three paths originating in Brazil: (1) $BR \rightarrow US$, (2) $BR \rightarrow CO \rightarrow None$, (3) $BR \rightarrow *** \rightarrow BR$. After processing the paths as described in Section 3.1.4, the resulting paths are: (1) $BR \rightarrow US$, (2) $BR \rightarrow CO$, (3) $BR \rightarrow BR$. The avoidance value for avoiding the United States would be $2/3$ because two out of the three paths do not traverse the United States. This metric represents a lower bound, because it is possible that the third path timed out ($***$) because it traversed the United States, which would make the third path: $BR \rightarrow US \rightarrow BR$, and would cause the avoidance metric to drop to $1/3$.

Avoidability algorithm with relays. Measuring the avoidability of Country Y from a client in Country X using relays entails two components: (1) Is Country Y on the path from the client in Country X to the relay? (2) Is Country Y on the path from the relay to the domain? For every domain, our algorithm checks if there exists at least one path from the client in Country X through any relay and on to the domain, and does not transit Country Y. The algorithm produces a value in the range [0,1] that can be compared to the output of the avoidability metric.

Upper bound on avoidability. Although the avoidability metric provides a way to quantify how avoidable Country Y is for a client in Country X, some domains may be hosted only in Country Y, so the avoidance value would never reach 1.0. For this reason, we measured the *upper bound* on avoidance for a given pair of (Country X, Country Y) that represents the best case value for avoidance. This algorithm analyzes the destinations of all domains from all relays and if there exists at least one destination for a domain that is not in Country Y, then this increases the upper bound value. An upper bound of 1.0 means that every domain that we measured is hosted (or has a replica) outside of Country Y.

4.3 Results

We examine the effectiveness of relays for country avoidance, as well as for keeping local traffic local. Table 3 shows avoidance

values; the top row shows the countries we studied and the left column shows the country that the client aims to avoid. Table 3 shows two trends: (1) the ability for a client to avoid a given Country Y increases with the use of relays; and (2) certain countries such as the United States, the United Kingdom, and other countries that are known to perform interference on traffic are also often the most difficult countries to avoid.

Relay Effectiveness. For 84% of the (Country X, Country Y) pairs shown in Table 3 the avoidance with relays reaches the upper bound on avoidance. In almost every (Country X, Country Y) pair, where Country X is the client’s country (Brazil, Netherlands, India, Kenya, or the United States) and Country Y is the country to avoid, the use of an overlay network makes Country Y more avoidable than the default routes. The one exception we encountered is when a client is located in Kenya and wants to avoid South Africa, where, as mentioned, all paths through our relays exit Kenya via South Africa.

Relays Achieve Upper Bound. Clients in the U.S. can achieve the upper bound of avoidance for all countries—relays help clients in the U.S. avoid all other Country Y unless the domain is hosted in Country Y. On the other hand, it is more rare for (Kenya, Country Y) pairs to avoid a given country. Relays can still be effective for clients in Kenya: for example, the default routes to the top 100 domains for Kenyans avoid Great Britain 50% of the time, but with relays this percentage increases to 97% of the time, and the upper bound is 98%.

U.S. is Least Avoidable. Despite increasing the ability to avoid the U.S., relays are less effective at avoiding the U.S. compared to all other Country Y. Clients in India can avoid the U.S. more often than clients in Brazil, Netherlands, and Kenya, by avoiding the U.S. for 65% of paths. Even using relays, Kenyan clients can only avoid the U.S. 40% of the time.

Keeping Local Traffic Local. Where there were relays located in one of the five countries that we studied, we evaluated how well the relays kept local traffic local. This evaluation was possible for the U.S. and Brazil. Tromboning Brazilian paths decreased from 13.2% without relays to 9.7% with relays; when relays are used, all tromboning paths go only to the U.S. With the relays, we see only 1.3% tromboning paths for a U.S. client, compared to 11.2% without relays. The 1.2% of paths that trombones from the U.S. traverse Ireland.

5 REGION-AWARE NETWORKING

Based on our measurement study, we design the first system to route traffic around a given country *without* the help of either ISPs or content providers.

Design Goals. Our measurement results motivate the design and implementation of a relay-based avoidance system, RAN, with the following design goals.

- **Country Avoidance.** The primary goal of RAN is to avoid a given country when accessing web content. RAN should provide clients a way to route around a specified country when accessing a domain. This calls for the role of measurement in the system design and systematizing the measurement methods discussed earlier in the paper.




Figure 7: RAN architecture.

- **Usability.** RAN should require as little effort as possible from clients. Clients should not have to download or install software, collect any measurements, or understand how the system works. This requires a way for clients to automatically and seamlessly multiplex between relays (proxies) based on different destinations. RAN uses a Proxy Auto-configuration (PAC) file to support this function. PAC files are supported on many types of devices, including mobile devices (smartphones, tablets, etc.). Additionally, this is a mechanism that is already being used in systems and tools. Many Internet users that use a VPN have *already* used a PAC file; when a user establishes a VPN connection, his device’s proxy settings are modified to point to a PAC file.
- **Scalability.** This country avoidance system should be able to scale to many users. Therefore, RAN should be able to handle the addition of relays, as well as be cost-effective in terms of resources required. This requires clever measurement vantage points, such that each vantage point is representative of more than one client. The PAC file allows RAN to grow with the number of clients and also supports incremental deployment.
- **Non-goals.** There are some challenges that RAN does not attempt to solve; in particular, it does not provide anonymity; it routes around countries, but it does not attempt to keep users anonymous in the event that traffic can be observed. RAN also does not address domestic interference or surveillance. For example, a client in the U.S. cannot use RAN to avoid network interference by the United States.

Design Overview. RAN comprises (1) an overlay network of relays; and (2) an oracle that directs clients to the appropriate relays, as shown in Figure 7. RAN’s relays are TCP proxy servers that allow clients to access web content without installing custom software. RAN uses the measurement methods described in Section 4 to learn paths between clients, relays, and domains; these results are stored at the oracle, which uses the data to decide which relay a client in some location should use for accessing a certain domain while avoiding a certain country. The oracle periodically computes paths for many combinations of client AS, destination, and country. A client can then query the oracle to determine the appropriate relay to use to avoid a certain country en route to a particular destination.

Country to Avoid	Brazil		Netherlands		India		Kenya		United States	
	No Relay	Relays	No Relay	Relays	No Relay	Relays	No Relay	Relays	No Relay	Relays
Brazil	0.00	0.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Canada	.98	1.00	.99	1.00	.98	.98	.99	.99	.92	1.00
United States	.15	.62	.41	.63	.28	.65	.38	.40	0.00	0.00
France	.94	1.00	.89	.99	.89	1.00	.77	.98	.89	.99
Germany	.99	1.00	.95	.99	.96	.99	.95	1.00	.99	1.00
Great Britain	.97	1.00	.86	.99	.79	1.00	.50	.97	.99	1.00
Ireland	.97	.99	.89	.99	.96	.99	.86	.99	.99	.99
Netherlands	.98	.99	0.00	0.00	.87	.99	.74	.99	.97	.99
Spain	.82	1.00	.99	.99	1.00	1.00	1.00	1.00	1.00	1.00
Kenya	1.00	1.00	1.00	1.00	1.00	1.00	0.00	0.00	1.00	1.00
Mauritius	1.00	1.00	1.00	1.00	1.00	1.00	.67	.99	1.00	1.00
South Africa	1.00	1.00	1.00	1.00	1.00	1.00	.66	.66	1.00	1.00
United Arab Emirates	1.00	1.00	1.00	1.00	1.00	1.00	.84	.99	1.00	1.00
India	1.00	1.00	.99	1.00	0.00	0.00	.94	1.00	.99	1.00
Singapore	.99	1.00	.99	1.00	.73	.94	.96	1.00	.99	1.00

Table 3: Avoidance values for using overlay network relays to avoid different countries. The upper bound on avoidance is 1.0 in most cases, but not all. It is common for some European countries to host a domain, and therefore the upper bound is slightly lower than 1.0. The upper bound on avoidance of the U.S. is significantly lower than for any other country; .886, .790, .844, and .765 are the upper bounds on avoidance of the U.S. for paths originating in Brazil, Netherlands, India, and Kenya, respectively.

5.1 Measuring Paths

RAN measures all paths using traceroute, which is then mapped to the country level using the same methods as described in Section 3 and shown in Figure 2. The paths we measure are the: forward paths from the client to each relay; forward paths from each relay to each domain; forward paths from the client to each domain; and reverse paths from each relay to the client. The portion of the reverse path from the domains to the relays is challenging to measure due to a lack of vantage points in ASes of common destinations. As discussed in Section 3.1, we found that the forward and reverse paths are asymmetric at the country level, and therefore RAN cannot make any guarantees about which countries are on the path between domains and relays even though it has calculated the paths from relays to domains. Despite the lack of knowledge about this part of the reverse path, we can reason about possible scenarios. If the client’s traffic is encrypted, then a country on this part of the reverse path that the client wishes to avoid cannot perform any traffic correlation attacks or website fingerprinting attacks, as the country cannot see who the client is (necessary for website fingerprinting) and does not have access to more than one part of the path (necessary for traffic correlation attacks).

Client-to-Relay Paths. To avoid requiring the client to install custom software, RAN measures client-to-relay paths from RIPE Atlas probes that serve as vantage points for the ASes where RAN clients might be. RAN selects probes that are geographically close to the client (e.g., in the same country). The oracle triggers the probe to run traceroutes to each relay. After collecting the responses, the oracle maps the IP-level paths to country-level paths and stores the results.

Relay-to-Client Paths. The RAN relays perform traceroutes to the IP addresses of RIPE Atlas probes, which represent client ASes. They then derive country-level paths; the oracle learns these paths from each relay.

Relay-to-Server Paths. Relays perform traceroutes to each domain. As with paths to clients, relays derive country-level paths and send them to the oracle.

Client-to-Server Paths. In case a path from a client to a domain does not pass through the country specified to avoid by default, then none of the proxies should be used. These paths are measured using the RIPE Atlas probes in similar locations as the clients, and the oracle triggers traceroutes from each of them to each of the domains. Corresponding country-level paths are stored at the oracle.

RAN must recompute these paths as they change. We measured the country-level paths from a RIPE Atlas probe to the Alexa Top 100 domains once per day for a month to see how stable country-level paths are. Across the measured domains, we found the average time between path changes to be about five days. Therefore, RAN recomputes the paths every five days to incorporate the most recent country-level paths.

5.2 Computing and Selecting Paths

The oracle follows four steps to decide which relay a client should use to access a specific domain: (1) If the default path from the client to the domain does not pass through the specified country, then do not use any of the relays. (2) Otherwise, for all the paths from the client to the relays, select suitable relays, which are relays where the country to avoid is not on the forward or reverse path between the client and relay. (3) From this set, if there is a path from a suitable relay to the domain that does not include the specified country, then use that relay for that domain. (4) If there is no path from the client through any of the relays to the domain that does not pass through the specified country, then select the relay that provides the most avoidance (measured by how many other domains that avoid the specified country).

The oracle applies this decision process to each domain, which results in a mapping of domains to relays that can be used to


```
function FindProxyForURL(url, host){
  if ((ShExpMatch(host, "*.google.com")))
    return "PROXY 1.2.3.4:3128";
  if ((ShExpMatch(host, "*.twitter.com")))
    return "PROXY 5.6.7.8:3128";
  return "DIRECT";
}
```

Figure 8: Example PAC file.

avoid the given country. To automate multiplexing between relays, RAN utilizes Proxy Autoconfiguration (PAC) files, which define how browsers should choose a proxy when fetching a URL. In the example PAC file in Figure 8, proxy 1.2.3.4:3128 should be used when accessing `www.google.com`, but proxy 5.6.7.8:3128 should be used when accessing `www.twitter.com`. The oracle uses the mapping of domains to relays to generate a PAC file, which specifies which domains should be accessed through which proxy. The PAC file is published online to a URL of the format `<client_country>_<country_to_avoid>_pac.pac`. The client uses this URL to specify their proxy configuration. Paths, and subsequently PAC files, are re-computed every five days.

5.3 Limitations and Extensions

From our experience conducting measurement studies of Internet paths, we have identified several limitations and obstacles. We can surmount these obstacles with the cooperation of content providers. To address the issue of path asymmetry, the reverse path could be measured from within the provider and used to determine if a country is on the reverse path; this could be used in conjunction with our measurements of the forward path. In addition, content providers could strategically publish DNS records such that when a client receives a DNS response, it is for a content replica that allows her to avoid a given country. A content provider could also replicate content in specific regions to allow clients to access replicas without traversing a specific country.

The current implementation of RAN does not include support by content providers, the design itself does not require any changes if providers were ultimately to cooperate with deployment. Cooperating ISPs and CDNs would collect and share traceroute data from their locations to different client and proxy locations and provide those measurements to the RAN oracle; RAN would then convert the traceroute data to country-level paths and incorporate them into the calculation of the PAC files. In certain cases, we could measure these paths without the cooperation of content providers. For example, for content hosted in public clouds, we could set up a virtual machine in those same data centers and have RAN collect the reverse path traceroute data to use when creating the PAC files.

6 IMPLEMENTATION & DEPLOYMENT

Our implementation of RAN includes relays, an oracle, and a client. RAN is open source and written in Python; the oracle is written in just 175 lines of code and the relay is written in just under 200 lines of code. RAN is currently deployed globally, and any user may use it today. We have released an anonymized source code repository, complete with usage instructions [44].

We assume that users and machines are trustworthy, and therefore the system runs securely. This implementation of RAN allows

a client to avoid a single country at a time; attacks on RAN, such as Denial of Service attacks and targeted surveillance of the relays, are outside the scope of the paper.

Relays. The current deployment has ten relays, one in each of the following countries: Brazil, Germany, Singapore, Japan, Australia, France, United States, United Kingdom, Netherlands, and Canada. These relays operate as Ubuntu Virtual Private Servers (VPSSes) with Squid as the proxy server and the RAN Relay software.

Oracle. The oracle software runs on a Fujitsu RX200 S8 server with dual, eight-core 2.8 GHz Intel Xeon E5 2680 v2 processors with 256GB RAM running RedHat Linux.

Client. To evaluate the RAN deployment, we set up a client machine in the Netherlands, which simply accesses web content and uses the PAC file generated by the oracle.

6.1 Other Considerations

Adding relays to RAN is straightforward. Additionally, RAN is resilient to failures of system components.

Adding relays and oracles. To add a relay, the system operator must set up a machine as a proxy server, install the relay software, and update the oracle’s list of relays. Adding an oracle requires installing the oracle software on a different machine, and specifying the client locations handled by that oracle (e.g., one oracle handles clients in North America and Europe, and another handles clients elsewhere). Both oracles will publish the PAC files to the same server, which causes no changes for the client.

Failed relays and oracles. Unresponsive relays are handled by the PAC file. The PAC file allows the oracle to specify multiple proxies in a sequential order, such that if the first proxy fails, then the client uses the second proxy (and so on). This feature can be used to specify all of the relays that have a path to the domain. And future work can include relay replicas that can be used in the case that a relay crashes.

Scaling RAN. As described above, adding relays and oracles to RAN is straightforward, and allows the system to support more clients. The addition of new relays and oracles is also essential as the number of clients grows to prevent the existing relays and oracles from being overloaded. Relays can be replicated in the same locations — especially in locations that are frequently used to avoid a given country — and relays can be set up in many more locations to provide avoidability of other countries and potentially provide alternate paths to the same destination. The ability to add relays and recover from relay failures allows RAN to scale with the number of clients, reducing the possibility of relays becoming bottlenecks in the system.

6.2 Evaluation

6.2.1 Country Avoidability. We measured RAN’s effectiveness in achieving country avoidance. We did so by first calculating the number of *default* paths that avoid a given country. Then we added a single relay, and calculated how many domains the client could access without traversing the given country. We repeated this approach for the remaining relays. We conducted the evaluation under the condition that the client wished to avoid different countries when accessing the Netherlands top 100 domains; Figure 9 shows these results. Each line represents the fraction of domains accessible

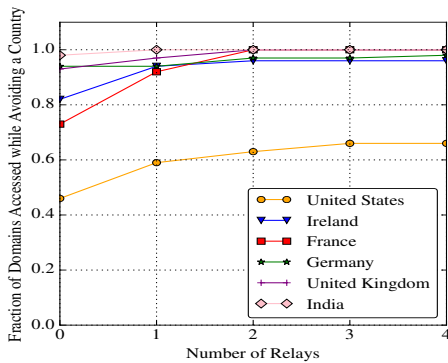


Figure 9: The effect of the number of relays on avoidance, for a client in the Netherlands.

while avoiding the country that the line represents. For example, 46% of domains are accessible without traversing the U.S. when RAN is not being used (zero relays), and if RAN is used, then 63% of domains are accessible without traversing the U.S.

RAN helps a client avoid a foreign country, as the fraction of domains accessible without traversing the specified country without RAN is lower than with RAN. Additionally, adding the first relay provides the greatest benefit, while subsequent relays offer diminishing returns. Figure 9 clearly shows that avoiding the U.S. is much more difficult (or impossible) than any other country. Only 63% of domains can be accessed while avoiding the U.S., whereas almost all domains can be accessed while avoiding any other given country.

As mentioned in Section 4.1, these results reflect how avoidable a country is based on the relays we were able to establish; results would show either the same or strictly better country avoidability if there were additional relays set up in other regions.

It is important to note that RAN cannot guarantee that a country is avoided because for some domains, the path must go through a certain country, as evidenced by our results for avoiding the United States. Despite this lack of guarantees, the system reduces the number of requests that transit the unfavorable country.

6.2.2 Performance. We measure both the throughput and latency of RAN and compare these results both to the default path and to RON [1], a comparable overlay system that focuses on improving performance and reliability. To measure throughput, we ran `wget` for each of the top 100 domains from the client machine in the Netherlands using an oracle-generated PAC file. Because different relays could have been used to avoid a single domain, the oracle selected a random relay from those that would allow the client to avoid the country. The oracle generated ten PAC files for a client in the Netherlands who wishes to avoid the U.S., randomly selecting a relay for domains that could have used different relays, and `wget` was used for the top 100 domains for each PAC file generated. Based on the `wget` output, we calculate the number of seconds to access content using our system and take the average across the ten experiments.

Figure 10 shows a CDF of the ratio of RAN throughput to direct throughput. The throughput of RAN is not significantly worse than

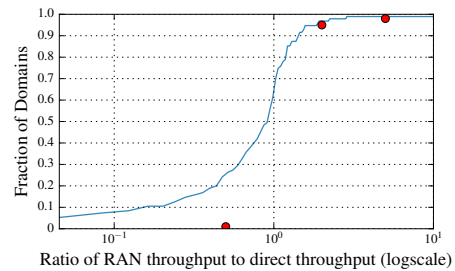


Figure 10: The ratio of RAN throughput to direct throughput.

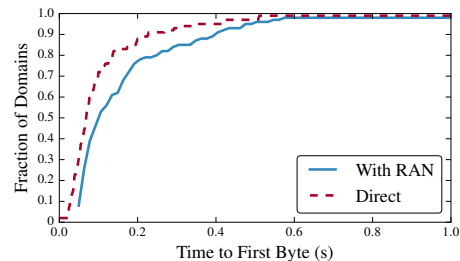


Figure 11: Time to first byte for RAN and direct paths.

that of default paths. In some cases the performance of RAN is *better* than that of default paths. Such improvements could be a result of the relays keeping local traffic local, or due to a closer content replica being selected. These results show that RAN’s performance is comparable to the performance of accessing domains without RAN. Figure 10 also compares RAN’s throughput to RON’s throughput, illustrated with the red dots; these data points are taken directly from the RON paper [1]. RAN performs worse than RON ($x < 1$), which is expected, as the detours that RAN introduces inherently inflate paths. Interestingly, both RON and RAN improve throughput for a similar fraction of samples ($x > 1$).

To measure the latency of RAN, we ran `curl` to each of the top 100 domains from the client in the Netherlands, using the ten oracle-generated PAC files. This experiment allowed us to measure the time to first byte (TTFB) for web downloads; we found the average TTFB when accessing content using RAN and found the TTFB when using direct paths; Figure 11 shows these results. The median TTFB for direct paths is 68.5 ms; for RAN paths the median is 100.8 ms; 90th percentile TTFB is 22.5 ms and 40.4 ms, respectively.

7 CONCLUSION

We have characterized routing detours through foreign countries, showing that underserved regions often depend on the U.S./Europe to access popular content; this can cause performance degradation, increased costs, and gives more power to dominant countries to perform surveillance and censorship. As a first step towards a remedy, we have designed, implemented, and deployed RAN, which employs overlay network relays to route traffic around a given country.

REFERENCES

- [1] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilient Overlay Networks. In *ACM Symposium on Operating Systems Principles (SOSP)*, volume 35. ACM, 2001.
- [2] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding Traceroute Anomalies with Paris Traceroute. In *ACM Internet Measurement Conference*, pages 153–158. ACM, 2006.
- [3] A Baker's Dozen, 2015 Edition. <http://research.dyn.com/2016/04/a-bakers-dozen-2015-edition/>.
- [4] S. Banerjee, T. G. Griffin, and M. Pias. The Interdomain Connectivity of PlanetLab Nodes. In *Passive and Active Network Measurement*, pages 73–82. Springer, 2004.
- [5] Z. S. Bischof, J. P. Rula, and F. E. Bustamante. In and Out of Cuba: Characterizing Cuba's Connectivity. In *The 2015 ACM Internet Measurement Conference*, pages 487–493. ACM, 2015.
- [6] Brazil Builds Internet Cable To Portugal To Avoid NSA Surveillance. <http://www.ibtimes.com/brazil-builds-internet-cable-portugal-avoid-nsa-surveillance-1717417>.
- [7] Brazil Conference will Plot Internet's Future Post NSA Spying. <http://www.reuters.com/article/us-internet-conference-idUSBREA3L10J20140422>.
- [8] Brazil Looks to Break from US Centric Internet. <http://news.yahoo.com/brazil-looks-break-us-centric-internet-040702309.html>.
- [9] Brazil to Host Global Internet Summit in Ongoing Fight Against NSA Surveillance. <https://www.rt.com/news/brazil-internet-summit-fight-nsa-006/>.
- [10] Brazil's President Tells U.N. That NSA Spying Violates Human Rights. <http://www.usnews.com/news/articles/2013/09/24/brazils-president-tells-un-that-nsa-spying-violates-human-rights>.
- [11] Brazil to Press for Local Internet Data Storage After NSA Spying. <https://www.rt.com/news/brazil-brics-internet-nsa-895/>, 2013.
- [12] CAIDA: Center for Applied Internet Data Analysis. <http://www.caida.org/home/>.
- [13] Data. <https://bitbucket.org/ransom,esearch/data/>.
- [14] Digital Colonialism & the Internet as a Tool of Cultural Hegemony. <http://www.knowledgecommons.in/brasil/en/whats-wrong-with-current-internet-governance/digital-colonialism-the-internet-as-a-tool-of-cultural-hegemony/>.
- [15] R. Dingleline, N. Mathewson, and P. Syverson. Tor: The Second-Generation Onion Router. Technical report, DTIC Document, 2004.
- [16] B. Eriksson, P. Barford, B. Maggs, and R. Nowak. Posit: a Lightweight Approach for IP Geolocation. *ACM SIGMETRICS Performance Evaluation Review*, 40(2):2–11, 2012.
- [17] B. Eriksson, P. Barford, J. Sommers, and R. Nowak. A Learning-Based Approach for IP Geolocation. In *Passive and Active Measurement*, pages 171–180. Springer, 2010.
- [18] EU Challenges US Hegemony in Global Internet Governance. <http://www.euractiv.com/section/digital/news/eu-challenges-us-hegemony-in-global-internet-governance/>.
- [19] R. Fanou, P. Francois, and E. Aben. On the Diversity of Interdomain Routing in Africa. In *International Conference on Passive and Active Network Measurement*, pages 41–54. Springer, 2015.
- [20] L. Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking (ToN)*, 9(6):733–745, 2001.
- [21] M. Gharaibeh, A. Shah, B. Huffaker, H. Zhang, R. Ensafi, and C. Papadopoulos. A Look at Router Geolocation in Public and Commercial Databases. 2017.
- [22] P. Gill, Y. Ganjali, B. Wong, and D. Lie. Dude, Where's That IP?: Circumventing Measurement-Based IP Geolocation. In *Proceedings of the 19th USENIX Conference on Security*, pages 16–16. USENIX Association, 2010.
- [23] C. Guo, Y. Liu, W. Shen, H. J. Wang, Q. Yu, and Y. Zhang. Mining the Web and the Internet for Accurate IP Address Geolocations. In *INFOCOM 2009, IEEE*, pages 2841–2845. IEEE, 2009.
- [24] A. Gupta, M. Calder, N. Fearnster, M. Chetty, E. Calandro, and E. Katz-Bassett. Peering at the Internet's Frontier: A First Look at ISP Interconnectivity in Africa. In *Passive and Active Measurement*, pages 204–213. Springer, 2014.
- [25] Y. He, M. Faloutsos, and S. Krishnamurthy. Quantifying Routing Asymmetry in the Internet at the AS Level. In *Global Telecommunications Conference*, volume 3, pages 1474–1479. IEEE, 2004.
- [26] Y. He, M. Faloutsos, S. Krishnamurthy, and B. Huffaker. On Routing Asymmetry in the Internet. In *Global Telecommunications Conference*, volume 2, pages 6–pp. IEEE, 2005.
- [27] How Brazil Crowdsourced a Landmark Law. <http://foreignpolicy.com/2016/01/19/how-brazil-crowdsourced-a-landmark-law/>, 2016.
- [28] Z. Hu, J. Heidemann, and Y. Pradkin. Towards Geolocation of Millions of IP Addresses. In *ACM Internet Measurement Conference*, pages 123–130. ACM, 2012.
- [29] B. Huffaker, M. Fomenkov, and K. Claffy. Geocompare: A Comparison of Public and Commercial Geolocation Databases. *Proc. NMMC*, pages 1–12, 2011.
- [30] Internet Hegemony and the Digital Divide. <http://www.economist.com/node/5165014>.
- [31] J. Karlin, S. Forrest, and J. Rexford. Nation-state Routing: Censorship, Wiretapping, and BGP. *arXiv preprint arXiv:0903.3218*, 2009.
- [32] 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 Internet Measurement Conference*, pages 71–84. ACM, 2006.
- [33] E. Katz-Bassett, H. V. Madhyastha, V. K. Adhikari, C. Scott, J. Sherry, P. Van Wesepe, T. E. Anderson, and A. Krishnamurthy. Reverse Traceroute. In *NSDI*, volume 10, pages 219–234, 2010.
- [34] D. Levin, Y. Lee, L. Valenta, Z. Li, V. Lai, C. Lumezanu, N. Spring, and B. Bhat-tacharjee. Alibi Routing. In *The 2015 ACM Conference on Special Interest Group on Data Communication*, pages 611–624. ACM, 2015.
- [35] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An Information Plane for Distributed Services. In *The 7th Symposium on Operating Systems Design and Implementation*, pages 367–380. USENIX Association, 2006.
- [36] MaxMind. <https://www.maxmind.com/en/home>.
- [37] No Internet Hegemony: Xi. <http://www.globaltimes.cn/content/958926.shtml>.
- [38] J. A. Obar and A. Clement. Internet Surveillance and Boomerang Routing: A Call for Canadian Network Sovereignty. In *TEM 2013: The Technology & Emerging Media Track-Annual Conference of the Canadian Communication Association (Victoria)*, 2012.
- [39] V. Paxson. End-to-end Routing Behavior in the Internet. *IEEE/ACM transactions on Networking*, 5(5):601–615, 1997.
- [40] S. Peter, U. Javed, Q. Zhang, D. Woos, T. Anderson, and A. Krishnamurthy. One Tunnel is (Often) Enough. *ACM SIGCOMM Computer Communication Review*, 44(4):99–110, 2015.
- [41] PlanetLab. <http://planet-lab.org/>.
- [42] I. Poesse, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye. IP Geolocation Databases: Unreliable? *ACM SIGCOMM Computer Communication Review*, 41(2):53–56, 2011.
- [43] Promoting the Use of Internet Exchange Points (IXPs): A Guide to Policy, Management and Technical Issues. <https://www.internetsociety.org/sites/default/files/Promoting%20the%20use%20of%20IXPs.pdf>, 2012.
- [44] RAN. <https://bitbucket.org/ransom,esearch/ran/>.
- [45] RIPE Atlas. <https://atlas.ripe.net/>.
- [46] H. Roberts, D. Larochelle, R. Faris, and J. Palfrey. Mapping Local Internet Control. In *Computer Communications Workshop (Hyannis, CA, 2011)*, IEEE, 2011.
- [47] A. Shah and C. Papadopoulos. Characterizing International BGP Detours. Technical Report CS-15-104, Colorado State University, 2015.
- [48] TeleGeography Submarine Cable Map. <http://www.submarinemap.com/>.
- [49] M. Wählisch, S. Meiling, and T. C. Schmidt. A Framework for Nation-Centric Classification and Observation of the Internet. In *The ACM CoNEXT Student Workshop*, page 15. ACM, 2010.
- [50] M. Wählisch, T. C. Schmidt, M. de Brün, and T. Häberlen. Exposing a Nation-centric View on the German Internet—A Change in Perspective on AS-level. In *Passive and Active Measurement*, pages 200–210. Springer, 2012.
- [51] D. L. Zhihao Li, Stephen Herwig, DeTor: Provably Avoiding Geographic Regions in Tor. In *USENIX Security 2017*, 2017.
- [52] S. Zhou, G.-Q. Zhang, and G.-Q. Zhang. Chinese Internet AS-level topology. *Communications, IET*, 1(2):209–214, 2007.