IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2017.DOI

DNSWeight: Quantifying Country-wise Importance of Domain Name System

DELIANG CHANG[1], SHANSHAN HAO[1], ZHOU LI[2], BAOJUN LIU[1], AND XING LI[1,3]

[1]Tsinghua University, Beijing 100084, China
[2]University of California, Irvine, Irvine, CA, USA
[3]Institute for Network Sciences and Cyberspace, Tsinghua University, Beijing, China
[3] CERNET center, Beijing, China

Corresponding author: Xing Li (e-mail: xing@cernet.edu.cn).

ABSTRACT DNS (Domain Name System) is one fundamental Internet infrastructure related to most network activities. As a feasible tool to govern the Internet, DNS's stability and interoperability will be impacted by the countries' policies or actions along the path. Especially now that many countries have stricter control over the Internet and even sometimes "unplug" it. But there was no study to quantify the countries' impact systematically. To fill this research gap, we present `DNSWeight`. This new data-driven approach utilizes a large-scale DNS dataset and BGP (Border Gateway Protocol) routing information to calculate the country-importance score so that a country's impact on DNS can be gauged. By applying `DNSWeight` on large-scale DNS and BGP datasets jointly, our study shows the importance among different countries is divided. A handful of countries show dominant significance to the current DNS ecosystem. Some countries with a history of Internet shutdowns are too influential to be ignored if they choose to break themselves from the Internet. We also examine the impact of IPv6 (IP Version 6) and reveal the "loop" phenomenon that occurs in some DNS queries. In conjunction with our findings, some discussion and suggestions are given. In summary, our study shows that DNS reliability needs to be reconsidered at the country's level.

INDEX TERMS BGP, Domain Name System, network measurement, network servers

I. INTRODUCTION

DNS translates the human-readable domains to network-layer IP addresses. Moreover, as one fundamental Internet infrastructure, it powers almost all Internet services like email and web. Thus its resilience is often highly concerned. One major threat to DNS resilience is due to its plain-text and connectionless nature: DNS is frequently under attacks like packet manipulation and eavesdropping, which fuels content censorship and access blocking [1]–[3]. We consider the issue from these two perspectives.

1) *Who*: Country policies and actions influence the DNS ecosystem. DNS is often used as tool by policymakers to censor/manage/monitor the Internet. Countries have the ability, motivation, and action to manipulate or eavesdrop on the system. In recent years, some countries even "unplug" their network from Internet [4]–[6], which may lead to more serious consequences.

2) *Where*: Countries on the resolution path are equally important. Many previous works only consider DNS servers and the network they are located, e.g., topological distribution of authoritative nameservers. It is insufficient since DNS manipulation is generally enacted using Man-In-The-Middle methods [2].

Until now, there was no consensus on how to quantify the real-world impact of a country on the DNS, though the answer is essential in guiding how the Internet should be advanced. Only a few previous studies looked into the geo-location distribution of root servers and TLD servers [7] and the influence of Autonomous System (AS) [8]. Still, the country-level impact cannot be derived from their result. Therefore, the main effort of this work is to collect relevant data and develop a new methodology to assess country-wise importance on DNS, under the consideration of path information.

Achieving such a goal is non-trivial, however. It is impossible to "turn off" and "turn on" a specific country net and learn its precise real-world impact. Although researchers have developed different platforms and client debugging applications for DNS measurement, we found that they cannot be directly applied to our problem (e.g., they are not scalable to analyze

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2021.3112926, IEEE Access

the whole DNS infrastructure). Besides, we do not find any evaluation metric about country-wise importance.

Our Approach. To address these challenges, we design a new approach named `DNSWeight` that can derive country-importance scores from large-scale DNS datasets and BGP routing information. We choose authoritative name servers as our measurement targets. Specifically, we first crawl recursively to fetch records of nameservers, given a list of domains. Then, we leverage the BGP routing dataset to discover the routes matching the addresses of authoritative name servers (abbreviated to ANS in our paper) and construct country paths from BGP AS paths. Notice that the path we are studying here is not the geolocation of the router but the registration country of the AS. Despite not being its geolocation, we argue that a router is managed by its network, and thus influenced by the country in which the network is registered. So analyzing AS paths allows us to study the impact of countries to some extent. Meanwhile, the geographical path of a route is also discussed later in the paper by introducing additional data. After that, inspired by the concept of betweenness centrality in graph theory, we propose three-level metrics to compute the country-wise importance.

Main Findings. We utilize `DNSWeight` to analyze a lists of over **1.3 million popular domains** and over **30 million** country paths (both IPv4 and IPv6 are included). Several new insights about the country's importance on DNS are obtained, and we highlight the major findings below. 1) we observe the importance between different countries is quite unbalanced. The United States plays a significantly dominant role in DNS infrastructure in every region. And the gap of country-importance is more evident in the IPv6 network, though IPv6 deployment is beneficial for enhancing diversity. 2) we show the evidence that the DNS infrastructure is highly connected, and countries with a history of Internet shut-downs, such as India and China, are controlling a large number of domain names and potentially impacting a lot more (about 50% domains in our list). 3) we found "country loops" of traffic forwarding are persistently observed on the routers that we surveyed. We conclude the network routes of DNS resolution should be meticulously optimized.

Contributions. The contributions of this paper are listed as follows:

- We propose a new approach `DNSWeight` to quantify the importance of a specific country/region to the entire DNS infrastructure.
- We use `DNSWeight` to perform a large-scale measurement study and obtained a suite of insights about DNS reliability in the lens of countries.
- We will release the source code of `DNSWeight` to help other researchers to study related issues.

II. BACKGROUND

A. DOMAIN NAME SYSTEM

The domain name system (DNS) is a distributed and hierarchical database. Resource Records (RR) of domain names

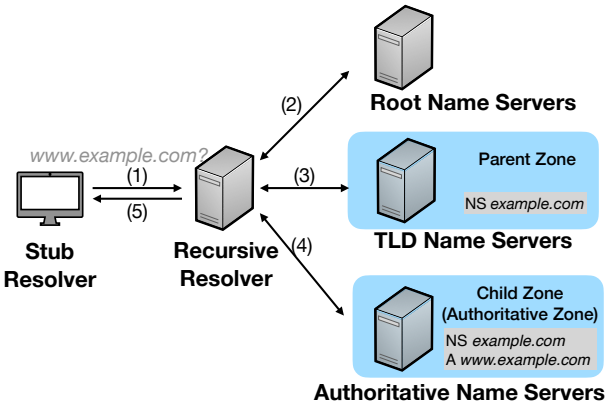

Figure 1: Standard DNS resolution process.

are stored in the ANS. Figure 1 shows the standard DNS resolution process. When a client requests the resolution of a domain name, the stub resolver will usually perform a DNS request to a recursive resolver (step 1). Then, the recursive resolver iteratively queries root, Top-Level Domain (TLD), Second-Level Domain (SLD), and deeper-level name servers (step 2-4). Finally, the DNS response will be returned to the client (step 5).

In order to make the DNS query function properly, some NS records of a domain name need to be stored in its parent zone. NS records contain domain names of name servers. To reach the name server, one needs to resolve domain names in NS records first. Sometimes, additional glued A or AAAA records are required. On the other hand, some NS records are stored in their authoritative zone. Figure 1 illustrates the two zones for `example.com`. In particular, `.com` has the parent zone, containing `NS` RRs and related glue records at its point of delegation. `example.com` has the child zone, containing the authoritative RRs of the name. In practice, the NS records extracted from the parent zone and the child zone might be inconsistent. Previous research [9], [10] has found that different recursive resolvers choose to use NS records in different zones. Considering the inconsistencies in NS records and resolver implementations, we need to measure the name servers contained in both zones as comprehensively as possible to provide a comprehensive picture of the domain name system. Our ANS collection process accommodates this parent/child zone setting and we elaborate the details in Section III-B.

B. BORDER GATEWAY PROTOCOL

The Internet is a decentralized network, consisting of more than 60,000 different interconnected network entities, which are named autonomous systems (ASes). The Border Gateway Protocol (BGP) is designed to broadcast routing and reachability information between ASes. The network routers that running BGP will propagate its IP prefixes and routing information to peers in an iterative way. Particularly, a router propagates and maintains the AS path for each IP prefixes, which represents a routing path from the router towards this

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2021.3112926, IEEE Access

prefix with AS information attached. In addition, the origin of this prefix is also included. As shown in Figure 2, for the routers in AS 500, the entry of IP prefix 3.2.1.0/24 has an AS path "`500 400 300 200`".

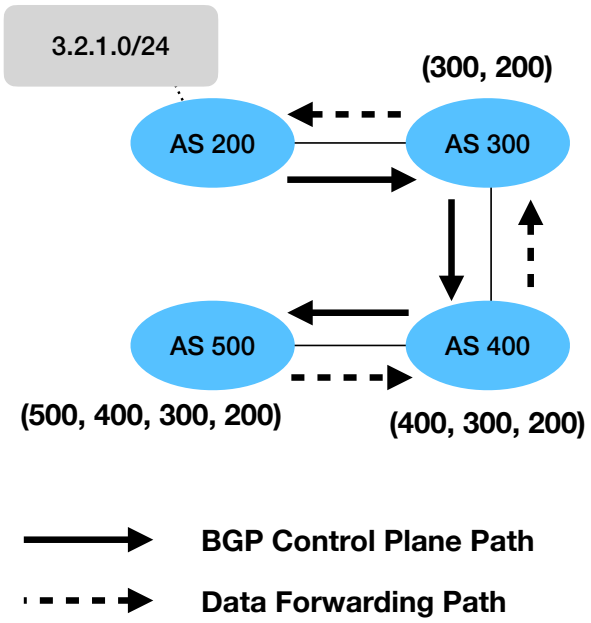

Figure 2: BGP AS Path

Used as a reference for route selection, an AS path also reveal the path information from an AS to a given IP prefix. We use the BGP data kept by the research projects [11]–[13] to construct the routing path from users to ANS. Section III-C elaborates how we process the BGP data.

C. NETWORK SOVEREIGNTY

Network sovereignty refers to the effort of an authority that creates boundaries on its governing network for the purposes like information control [14]. In this post-Snowden era, more countries are seeking network sovereignty [15]. Various countries from East [16]–[18] and West [19] censor web content as an approach towards network sovereignty. While those authorities often argue better network management can be achieved, network sovereignty raises the concerns because of not only the compromise of the open Internet, but also the potential collateral damage to other countries [20]. To the extreme, in 2019, Russia carried out an experiment which unplugs its network from the Internet [4].

While the issue is known, we argue its implication to DNS infrastructure needs to be better studied. As advocated by the Internet community, multiple ANS with broad geographical distribution should be installed by the domain owner. Techniques like IP Anycast automatically distribute users' DNS requests from one country to another. Taking the example of Russia, unplugging its network could adversely impact the DNS resolution of users outside Russia. There has been no study attempting to quantify a country's impact, not to mention a "what-if" analysis about the consequence when a country's network is drastically changed (e.g., what if China blocks all DNS requests from outside). We aim to fill this gap of understanding and present our measurement result in section IV-C.

III. METHODOLOGY

We design a system named `DNSWeight` to measure the importance of a country/region to the DNS infrastructure. Section III-A overviews the challenge to measure country-wise importance on DNS and the rationale of the design of `DNSWeight`. Section III-B, III-C and III-D elaborate the design and implementation details of each component composing `DNSWeight`.

A. OVERVIEW

Though we cannot "turn off" or "turn on" a network block to learn its precise impact, we can use the publicly available data from DNS and routers to estimate it.

Therefore, Internet Services Providers (ISPs) or even governments may have various potential impacts to manipulate the DNS traffic. To quantify the likelihood of this kind of threat, we want to measure the "country importance". It will help us to answer the question that, once a query of a random domain name is issued, what is the possibility that the resolution process could be influenced by a specified country/region? Intuitively, the higher probability indicates higher impact. To this end, a large-scale DNS dataset augmented with routing information has to be gathered and analyzed.

Firstly, we choose authoritative name servers as our measurement targets. We have several considerations over this choice. First, unlike recursive resolvers, authoritative name servers are owners of domain names. A stub resolver could connect to different recursive resolvers under different configurations, but the query will finally go to the authoritative name server of that domain name. Recursive resolvers are often located near stub resolvers because of performance considerations, while the authoritative name servers are distributed in different networks, which are governed by located countries. On the other hand, unlike DNS-over-HTTPS or DNS-over-TLS, the traffic between recursive resolvers and authoritative name servers is still mostly plaintext.

Secondly, in order to obtain a global landscape of country importance parameters, we want to build a comprehensive and representative domain name list and consider all of them in our measurement study. The volume of target domain names is on the order of millions, which may propose a challenge to existing measurement platforms.

Further, due to the country in the middle of the traveling network path of a DNS query may also have the ability to inspect or hijack DNS packets, it requires us not only to premeditate the country importance as the destination one, but also in the middle of the network path. Therefore, two types of importance scores are considered in our approach. One is "destination-wise importance", which indicates the possibility of a random query going towards a certain country/region.

The other is "path-wise importance", which measures the possibility of a random query going through a country/region. The detailed process is discussed in section III-D.

Note that the impact of recursive resolvers, especially public resolvers, on the DNS ecosystem cannot be ignored, especially from the user's perspective. A country is able to collect a large amount of user DNS queries from all over the world because it runs a popular and international recursive resolver. However, this paper discusses the impact from the perspective of the name owner, or name servers. The study of recursive resolver could be our future work.

While previous studies have introduced measurement platforms, client-side debugging tools, and data sources for DNS measurement, we found they cannot be easily applied to our setting. Measurement platforms like RIPE Atlas [10], [21] can be leveraged to probe any DNS entity with a selected vantage point, but using them to probe a large number of DNS entities is not scalable. DNS data sources like zone files [22] and passive DNS [23], though they offer a broad view of DNS ecosystem, are not cataloged by countries/regions and the routing information (i.e., the routers and nameservers passed by a DNS request) is not included. In addition, we are not aware of any metric that is tailored to measure country-wise importance on DNS and answer the question we raise at the beginning of this section.

`DNSWeight` tackles these issues with three components developed by us. They are illustrated in Figure 3 together with the workflow. The first component, *ANS crawler*, is designed to harvest the destinations of users' DNS requests by fetching the records related to ANS through active domain resolving. Due to the data inconsistency issue between parent zone and child zone, we develop new zone crawling algorithms that can recursively find the records of our interest. The second component, *country path finder*, is able to efficiently discover the routes matching our ANS dataset from the BGP data provided by public sources. In addition, it maps the BGP routes from AS to country to construct the country paths. The third component, *importance calculator*, leverages three metrics designed by us to assess the country-wise importance. The key building block of our metrics is *betweenness centrality*, a concept in graph theory. With these metrics, we can compute the destination-wise and path-wise importance for any country, on any network region, inside DNS infrastructure.

Our approach uses both DNS data and BGP data. On the one hand, we collect information about a large number of name servers through DNS active measurements. On the other hand, domain queries can be wiretapped or even tampered with by the network in the path, so attention needs to be paid to the path, too. As we discussed before, the existing approach cannot achieve large-scale path measurements, therefore we used data from BGP to measure the resolution path. In general, leveraging the hybrid model that combines active with passive measurements, we are able to obtain network path information of specified authoritative name servers from thousands of global distributed vantage points.

B. ANS CRAWLER

To quantify the country-wise importance with DNS, we take a collection of ANS associated with popular domains as the measurement target. We set the size of the domain list to be above **1 million**, so the majority of users' DNS requests can be covered. While one can obtain a larger list of domains by downloading TLD zone files (e.g., `.com` and `.net` from VeriSign [22]), we did not follow this approach based on two reasons: 1) some DNS zones are not accessible to researchers, such as many ccTLD (country code Top-Level Domain) zones except for a few, but neglecting domains based on TLDs will lead to biased measurement results; 2) according to the previous study [24], the majorly records in the zone files are not requested by any user, so including them in this study is unnecessary. We compile two lists of popular domains based on the *global* view and a *local* view, by collecting domain names from public sources and DNS servers.

Global List. We downloaded the domain names from Alexa Top Sites and Tranco Lists on March 26th, 2020, to fill the global list. Alexa ranks the popular domains according to the incoming traffic volume [25], and it has been widely used in previous DNS measurement studies [8], [26]–[28]. Tranco is another domain list proposed by Pochat et al. [29] which overcomes the issues like instability and ranking manipulation of the Alexa list. We merge the two lists, which have **1,415,146** unique domains in total.

Local List. The global list helps us understand a country's impact on the global DNS users. However, the impact could vary when the scope is reduced to a region, as the regional users have different preferences in visiting domains. As a comparison, we create a local list of popular domains based on the DNS requests to *a large China public resolver*. The estimated number of users using the resolver is 10 million by counting the source IP addresses. The list contains **1,048,575** unique domain names ordered by the count of DNS queries within one month of 2020[1]. Two lists have only 127,678 domains in common (9.0% in global list and 12.2% in local list).

Collecting Nameserver Records. We need to extract the IP addresses of *all* ANS associated with each domain from the domain lists for our measurement study. To this end, we fetch the Nameserver Resource Records (`NS` RRs) from the parent zone and the child zone of each domain. Section II-A describes their relations. As the nameservers in the parent zone might also be owned by the domain owner and the addresses could differ from the ones in the child zone, we combine the `NS` RRs in both zones and regard them as ANS in our study.

For each domain name in the list, we perform a recursive resolution starting from the root server. A chain of responses from different levels of authoritative servers (e.g., root, TLD,

[1]Lower-bound threshold of query count is 10,279.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2021.3112926, IEEE Access

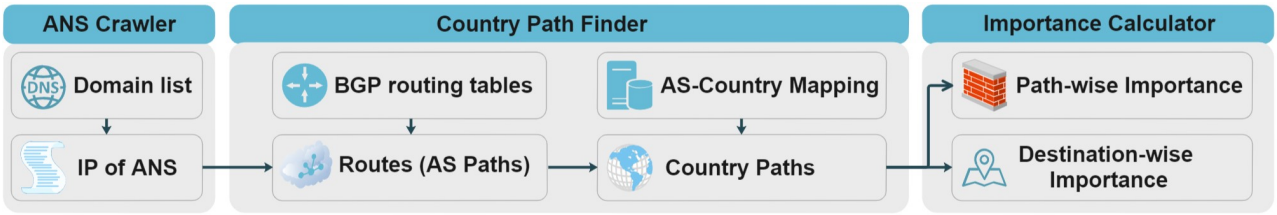

Figure 3: Overview of `DNSWeight`.

Algorithm 1 Parent Zone Crawling

Input: domain name d.

Output: `NS` (P_{NS}) and glued `A/AAAA` (P_{IP}) of d in the parent zone.

$\texttt{IP} \leftarrow$ Root server
$P_{NS}, P_{IP} \leftarrow \emptyset$
$\texttt{AA} \leftarrow$ **False**
$N \leftarrow 10$ // maximum iteration
while $N > 0$ **do**
  $\texttt{R} \leftarrow get\_dns\_response(d, \texttt{IP})$
  $\texttt{AA} \leftarrow get\_AA\_bit(\texttt{R})$
  if $\texttt{AA} = True$ **then**
    break;
  end if
  $P_{NS} \leftarrow get\_auth\_NS(\texttt{R})$
  $P_{IP} \leftarrow get\_glued\_IP(\texttt{R})$
  $\texttt{IP} \leftarrow get\_next\_nameserver\_IP(P_{NS}, P_{IP})$
  $N \leftarrow N - 1$
end while
return P_{NS}, P_{IP}

Algorithm 2 Child Zone Crawling

Input: domain name d, its parent `NS` P_{NS} and glued `A/AAAA` P_{IP}.

Output: `NS` (C_{NS}) and related `A/AAAA` (C_{IP}) of d in the child zone.

$C_{NS}, C_{IP} \leftarrow \emptyset$ // child zone
$Q_{NS}, Q_{IP} \leftarrow \emptyset$ // queried NS and IP
$U_{NS} \leftarrow P_{NS}$ // unqueried NS
$U_{IP} \leftarrow P_{IP}$ // unqueried IP
while True do
  $I \leftarrow get\_IP\_addr(U_{NS})$
  $Q_{NS} \leftarrow Q_{NS} \cup U_{NS}$
  $U_{IP} \leftarrow U_{IP} \cup (I \setminus Q_{IP})$ // set complement
  $C_{IP} \leftarrow C_{IP} \cup I$
  if $U_{IP} = \emptyset$ **then** // if there's no new address, quit the process.
    break;
  end if
  $T_{NS}, T_{IP} \leftarrow get\_ns\_record(d, U_{IP})$
  $Q_{IP} \leftarrow Q_{IP} \cup U_{IP}$
  $C_{NS} \leftarrow C_{NS} \cup T_{NS}$
  $C_{IP} \leftarrow C_{IP} \cup T_{IP}$
  $U_{NS} \leftarrow T_{NS} \setminus Q_{NS}$ // find unqueried NS
  $U_{IP} \leftarrow T_{IP} \setminus Q_{IP}$ // find unqueried IP
end while
return C_{NS}, C_{IP}

and SLD) will be returned and we extract the records related to ANS by looking for the response (termed R) with the `AA` flag set. We first extract the parent `NS` records as well as the glue records (including `A` and `AAAA`) from the responses issued before R. Algorithm 1 shows the process.

Following, we extract the child `NS` records. Algorithm 2 shows the process. For every domain d in the list, we maintain an unqueried NS set U_{NS} and an unqueried IP set U_{IP}. They are initialized by the data obtained in the parent zone crawling. At each step, we start with resolving IP addresses for all NS in U_{NS}, then add the new ones to U_{IP}. Next, we query `NS` records of d to all IP in set U_{IP}, to seek new `NS` and IP addresses to augment U_{NS} and U_{IP}. We will repeat this process until there are no new `NS/A/AAAA` records that could be found.

With this recursive searching process, all ANS with connections to a domain can be identified. Besides, this process might incur high overhead on our crawler and DNS servers. We optimize this process by caching the responses locally. Our key observation is that many domain names share a same group of ANS, usually operated by public name service provider like Cloudflare [8]. Thus, if two domains share a same name server, it doesn't need to be queried twice for its IP address. Therefore, in our approach, all query responses are cached locally in the process of measurement. Every request is sent only after failed attempt to retrieve the answer from the local cache. This optimization reduces the workload significantly of not only the recursive resolver, but also the target name servers.

We deployed our crawler on a VPS (Virtual private server) in Japan. It takes two days to crawl the global and local lists. For the global list, 1,380,395 (97.5%) names could be resolved. We found 314,270 distinct `NS` records associated with 154,333 IPv4 addresses and 13,126 IPv6 addresses in the parent zones. Besides, 357,278 distinct `NS` records associated with 223,154 IPv4 addresses and 25,682 IPv6 addresses are extracted from the child zone. The child zones have 45% more IPv4 addresses and 96% more IPv6 addresses than the parent zones. For the local list, 838,568 (80%) names could be resolved. While 62,422 distinct IPv4 and 11,160

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2021.3112926, IEEE Access

%	Global List	Local List
Unresolved	2.46	20.0
NS(P)=NS(C)	83.5	70.5
NS(P) $\subset$ NS(C)	4.71	2.65
NS(P) $\supset$ NS(C)	3.09	3.91
NS Other	6.27	2.92
IPv4(P) = IPv4(C) $\neq \emptyset$	38.3	28.3
IPv4(P) = IPv4(C) = $\emptyset$	1.15	1.64
IPv4(P) $\subset$ IPv4(C)	52.8	43.9
IPv4(P) $\supset$ IPv4(C)	1.88	3.10
IPv4 Other	3.46	3.06
IPv6(P) = IPv6(C) $\neq \emptyset$	21.9	20.4
IPv6(P)=IPv6(C) = $\emptyset$	41.0	30.7
IPv6(P) $\subset$ IPv6(C)	33.5	27.7
IPv6(P) $\supset$ IPv6(C)	0.57	0.60
IPv6 Other	0.61	0.63

Table 1: Ratios of ANS records related to parent zone (P) and child zone (C) in global and local list. NS(·), IPv4(·), IPv6(·) are the set of NS records, IPv4 addresses and IPv6 addresses of the parent and child zone.

IPv6 addresses could be found in the parent zone, 79,999 (28% more) IPv4 and 19,351 (74% more) IPv6 addresses could be found in the child zone. The ratios of different types of records are shown in Table 1. Note that the local list is a domain list ranked by query count from a public resolver. It does not check if the responses are valid. So more domains in the local list (20.0%) fail to resolve compared to the global list (2.46%).

When taking a closer look at our data, we found non-negligible inconsistency between the parent zone and the child zone, which illustrates the necessity of performing ANS crawling across different name servers. We found more than 43% collected domains have more IPv4 addresses in the child zone comparing to the parent zone. About 6% domains have IP addresses in parent zone not belonging to domains' child zone, which might be dangling records [30]. Our result about `NS` distribution is similar to a recent work [9] but the IP distribution is different. There could be two reasons: 1) We use popular domain list while previous work [9] uses zone files of 3 TLDs (`.com`, `.net`, and `.org`). 2) We crawl all ANS we find while previous work [9] randomly choose one.

Influence of vantage point (VP). Crawling from different parts of the world may lead to different `NS/A/AAAA` records. To quantify the influence of VP choice in the process of NS crawling, we acquire NS data of domains in the global list from two vantage points in Japan and Los Angeles in the same period of time and find that only less than 0.3% of successfully resolved domains have at least one different name server. We further look into the details and find that part of the differed records is not inclusion relationship, possibly due to CDNs. So in our study, we use IP addresses retrieved from one vantage point as the study target.

Influence of temporal IP churn. In addition to spatial VP choice, the influence of temporal IP churn is also considered.

Source	#IPv4 Routers	#IPv6 Routers	#ASN
RIPE RIS	685	518	493
RouteView	392	291	261
Isolario	195	129	133
Combined	654	477	679

Table 2: Statistics of BGP data. #ASN is number of distinct AS.

Crawling all domains in the list is still time-consuming and the results might vary at different time and by network conditions. The IP churn related to a domain could lead to inconsistency in the later stage of routes lookup. To estimate the likelihood of such IP churn, we crawl the domains again a month later and found only 1.3% domains have at least one IP address changed and all of the changes are IPv6 addresses. Therefore, we conclude our dataset is stable longitudinally during our experiment.

C. COUNTRY PATH FINDER

With the ANS collected, we obtain their associated country paths, which is defined as all countries that a request passes through from a vantage point to a destination IP address of ANS. To achieve this goal, We first collect the routing data and locate the AS paths related to the studied ANS. With the AS-to-country mapping, we convert each AS path to a country path. Below we elaborate on the three steps.

Step One: Collecting Routes. We download BGP tables from routers around the world and extract the routes. Three sources are leveraged: 1) RIPE Routing Information Service [11], 2) RouteViews [12], 3) Isolario Project [13]. The details of the data are listed in Table 2. The routes of the three sources are merged and the inconsistency between routers should be resolved. As such, when multiple routers are observed in an AS, we choose only one router with the largest number of routing entries. IPv4 routers and IPv6 routers are selected separately. We collected snapshots of route tables from the same period of time as ANS crawling. In the end, our dataset contains routes from 654 IPv4 routers and 477 IPv6 routers, covering 679 different ASes.

Step Two: Locating AS Path. A BGP route in the downloaded BGP data consists of a path of AS numbers (called AS path) and an entry IP prefix, which specifies how a packet should travel from one router to another. We locate the AS path between a vantage point router (V) and a nameserver (N), by querying the N's IP address on the routing table of V. We use the longest prefix match algorithm [31] to find the best route. The IP address of N could be converted to an IP prefix (/24 for IPv4 and /64 for IPv6) before query, as the prefixes longer than them are rarely seen in public routing tables [32]. For instance, assume `example.com.` has an ANS `199.43.135.53` and we choose a router in `AS 4777` as our vantage point. In its routing table, the most relevant entry is `199.43.135.0/24`, which has an AS path "`4777 2516 3257 40528 26710`". Then the path

is selected as the AS path from V to N.

Step Three: Country Path Mapping. After an AS path is located, we map each AS node to a country to construct a country path using an AS-to-Country database, ASRank [33]. The derived path could have identical country nodes consecutively, and we remove the duplicated nodes with a node collapsing process.

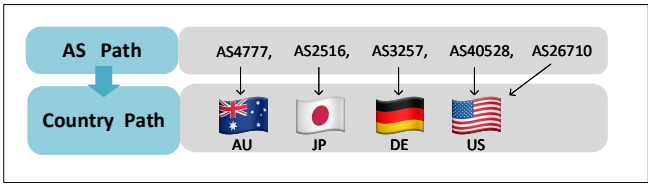

Figure 4: Converting AS path to country path.

As shown in Figure 4, AS path "`4777 2516 3257 40528 26710`" is converted to "`AU JP DE US`" as `AS 40528` and `AS 26710` share the same country. We also augment each country node with RIR (Regional Internet registry) using the Country-RIR mapping provided by RIPE [34], to study a country's impact on a bigger region.

Limitation of router selection. The vantage points used by our study are actually routers affiliated with the three data sources. Each router has a number of routes but the distribution is not even. In addition, the routers of some data sources are concentrated in one region. For example, the data we collected has far more routers in Europe (see Table 3). We divide vantage points by RIR in the later measurement study, which alleviates the issue of uneven distribution of routers and gives us a more accurate understanding of user perspectives on each continent. Also, the routing data is acquired from public shared routing tables, so private peering [35] is not considered in our study. We leave that for our future work.

Country-path & geographical path. We use the AS's owner country to construct a country path, instead of its routers' residential information. Sometimes, the residing country and owner country differ [36]. For instance, though managed by a German network provider, a router forwarding the traffic could reside in an IXP (Internet exchange point) in the US. We argue that, even if a router is not located in its registered country, it is still managed by the network to which it belongs, and thus influenced by the policies of the country to which the AS belongs. For instance, network surveillance or traffic hijacking could be deployed. In our study, we want to estimate the influence of the country on the DNS system, so we define the country path based on the aforementioned observation. On the other hand, the geographical path of a route is also discussed in the later section IV-E.

D. IMPORTANCE CALCULATOR

After we obtain the country path from all vantage points to each ANS, we calculate the importance scores for each country (or country-wise importance) within DNS. The country located at the destination and middle of the path will be computed separately.

If a country is important as a destination, it is supposed to have strong influences on the resolution result. To be more specific, if a country's destination-wise importance is 1, it means that all queries of all domains in the list will eventually go into this country.

On the other hand, if a country is important in the middle of the path, it is more likely to impact the resolution topology, e.g., deciding the next country to receive the DNS requests. If the country's path-wise importance is 1, it means that all DNS queries in our measurement will go through this country on the path.

The country-wise importance is computed separately for IPv4 and IPv6.

While the choices of importance metrics are abundant, we would like the metric to measure the *probability* that, on average, a DNS query passes through/to a country towards a domain name. Here, we use $t_{dst}(c_i)$ and $t_{path}(c_i)$ to represent the *destination-wise importance* and *path-wise importance* of a country c_i. We would like two properties to hold: 1) **linearity**, for country c_i and c_j, if $t_{dst}(c_i) = n \cdot t_{dst}(c_j)$, then c_i has n times influence as destination comparing to c_j. The same applies to path-wise importance; 2) **additivity**, t_{dst} and t_{path} can be added together to represent the overall importance of a country.

Based on the above requirement, we design three levels of importance. The higher-level importance is composed of the lower-level ones. A set of notations are defined here: we define the list of domains as $D = \{d_1, d_2, ..., d_n\}$. For each domain d_i, the ANS set is $N_i = \{A_1, A_2, ...\}$ and $A_j = \{IP_1, IP_2, ...\}$ are the associated IPv4 and IPv6 addresses. For an IP address IP_i, country paths going towards it composes a set $\sigma_{IP_i} = \{p_1, p_2, ...\}$, where p_i is a country path.

IP-level Importance: We use *betweenness centrality* [37], a concept in graph theory, to reflect the degree of interaction between one country node and others. It is built on the number of shortest paths passing through the node. This metric has also been used by routing studies [28], [38].

For IP_j, all country paths towards it is σ_{IP_j}. Then for country path p_k, we define $last(p_k)$ as the last node on the path. And define $path(p_k)$ as the nodes except the first and last nodes. For a country c_i and IP_j, the destination-wise importance is:

$$r_{dst}(c_i, IP_j) = \frac{|\{p_k | p_k \in \sigma_{IP_j}, c_i = last(p_k)\}|}{|\sigma_{IP_j}|}$$

Similarly, the path-wise importance could be defined as:

$$r_{path}(c_i, IP_j) = \frac{|\{p_k | p_k \in \sigma_{IP_j}, c_i \in path(p_k)\}|}{|\sigma_{IP_j}|}$$

To sum up, given a country c_i and an address IP_j, IP-level importance represents how likely a DNS packet goes

through/to country c_i when access IP_j.

Domain-level Importance: For a country c_i, domain d_k and its nameservers $A_j \in N_k$. Based on IP-level importance, we define domain-level destination-wise (s_{dst}) and path-wise (s_{path}) importance as follows:

$$s_{dst}(c_i, d_k) = \frac{1}{|N_k|} \sum_{A_j \in N_k} \frac{1}{|A_j|} \sum_{IP_m \in A_j} r_{dst}(c_i, IP_m)$$

$$s_{path}(c_i, d_k) = \frac{1}{|N_k|} \sum_{A_j \in N_k} \frac{1}{|A_j|} \sum_{IP_m \in A_j} r_{path}(c_i, IP_m)$$

Domain-level importance is the average of IP-level importance values calculated for each IP address of ANS. It represents how likely a DNS packet goes through/to country c_i when resolving a domain d_k. We consider the same weight for all ANS under a domain because previous research shows they are often accessed together by the recursive resolvers [39].

DNS-level Importance: For a country c_i, a domain list D, given domain-level importance, we define DNS-level destination-wise (t_{dst}) and path-wise (t_{path}) importance as follows:

$$t_{dst}(c_i) = \frac{1}{|D|} \sum_{d_k \in D} s_{dst}(c_i, d_k)$$

$$t_{path}(c_i) = \frac{1}{|D|} \sum_{d_k \in D} s_{path}(c_i, d_k)$$

It represents how likely a DNS packet goes through/to country c_i when all domains in D are considered. We do not multiply domain-level importance with a power-law coefficient [40], which reflects the distribution of Internet traffic to domains, because we want to measure a country's influence on DNS infrastructure rather than actual traffic. So in our model, an ANS that hosts multiple domains is more important than an ANS host one single popular domain. However, to provide a comprehensive picture of the process, we also study the importance score when considering domain popularity in section IV-F

Discussion. The BGP dump is the snapshot of the current routing table. If a path in it fails due to some reason, it is highly likely for the router to choose another viable route path. Therefore, instead of all possible routes, we only consider the current path in the snapshot in our approach to reduce the computation overhead and depict current importance. Still, the volume of paths (over tens of millions) presents a representative view of the country's importance in DNS.

IV. MEASUREMENT RESULT

With the routing data collected about domains in our lists, we assessed the country's importance in the DNS infrastructure and reported our findings in this section. We first look into the impact of different countries by IPv4 and IPv6 addresses and country's influence on domain names. Then, we switch the perspective from country to domain and measure the geographical patterns of ANS deployment. Next, we analyze the patterns of country paths and the loops in particular. Finally, we study how the local list will impact the measurement result.

A. IPV4 COUNTRY IMPORTANCE

Maintaining the ANS diversity is important for the robustness of DNS resolution, which is highly advocated by the Internet community. In particular, RFC 1034 requires at least 2 ANS to be maintained for each DNS zone [41] and RFC 2182 asks for geographical and topological diversity of ANS [42]. While the ANS diversity within the zone files has been measured [8], what is the role a country plays and how it impacts DNS resolution are not assessed. The data we collect allows us to answer these questions.

	IPv4		IPv6	
RIR	#VP	#Paths	#VP	#Paths
AfriNIC	36	1,454,679	23	195,473
APNIC	70	4,507,476	47	409,384
ARIN	115	7,905,671	72	676,426
LACNIC	63	3,913,610	53	469,937
RIPE NCC	370	23,285,636	282	2,654,343

Table 3: Statistics about vantage points and country paths. "VP" and "Paths" are vantage points and country paths.

Specifically, we choose the domain names in the global list (1,380,395) and their IPv4 ANS (154,333 in parent zone and 223,154 in child zone) to measure the country's importance in the IPv4 space. For the routes reaching the ANS, we separate them by *the routers' RIR* to more precisely assess the impact based on users' geo-locations, since the routing paths could differ based on where users initiate DNS request. The 5 RIRs are AfriNIC (Africa), APNIC (East Asia, Oceania, South Asia, and Southeast Asia), ARIN (Antarctica, Canada, parts of the Caribbean, and the United States), LACNIC (the Caribbean and all of Latin America) and RIPE NCC (Europe, Central Asia, Russia, and West Asia). Table 3 lists the statistics of our routers based on RIR. Though our collected paths are unbalanced among different RIRs, we are able to have sufficient paths for each one to obtain meaningful results.

We apply `DNSWeight` to compute t_{dst} and t_{path} for each country. Table 4 shows result divided by routers' RIRs. Figure 5 shows t_{path} and t_{dst} for each country/region. t_{path} is usually small for most countries (less than 10^{-2}) because a large number of routes connect one country to another without going through a third country.

Our first finding is **the United States (US) plays the dominant role in DNS**, which echoes with other studies about Internet infrastructures [43], [44]. Not only it serves most ANS (t_{dst} ranges from 0.471 to 0.541), but it also relates to most paths (t_{path} ranges from 0.087 to 0.252) for almost all RIRs. This can be explained partially by the US companies'

AfriNIC			APNIC			ARIN			LACNIC			RIPE NCC		
CC	t_{dst}	t_{path}	CC	t_{dst}	t_{path}	CC	t_{dst}	t_{path}	CC	t_{dst}	t_{path}	CC	t_{dst}	t_{path}
US	0.506	0.252	US	0.471	0.239	US	0.541	0.087	US	0.493	0.299	US	0.518	0.230
DE	0.127	0.019	DE	0.120	0.024	EU	0.007	0.182	DE	0.124	0.026	EU	0.014	0.245
EU	0.006	0.103	EU	0.006	0.135	DE	0.117	0.064	ES	0.020	0.123	DE	0.130	0.060
AO	0.000	0.090	HK	0.063	0.060	CN	0.061	0.001	EU	0.006	0.135	CN	0.061	0.003
ZA	0.060	0.002	CN	0.056	0.001	GB	0.025	0.025	IT	0.005	0.061	RU	0.041	0.007
CN	0.059	0.000	RU	0.038	0.006	RU	0.041	0.007	CN	0.058	0.006	GB	0.020	0.019
MU	0.000	0.051	SG	0.019	0.019	FR	0.030	0.001	CO	0.000	0.058	FR	0.030	0.006
IN	0.003	0.044	AU	0.005	0.031	HK	0.012	0.016	RU	0.039	0.006	CH	0.007	0.019
RU	0.040	0.007	JP	0.013	0.018	NL	0.013	0.006	BR	0.030	0.001	NL	0.017	0.008
GB	0.015	0.015	FR	0.027	0.001	TR	0.015	0.001	FR	0.029	0.001	AT	0.003	0.018

Table 4: Top 10 most important countries/regions of IPv4 DNS separated by RIRs. The countries are ranked by by overall importance ($t_{dst} + t_{path}$). CC is alpha-2 country code. t_{dst} is destination-wise importance. t_{path} is path-wise importance.

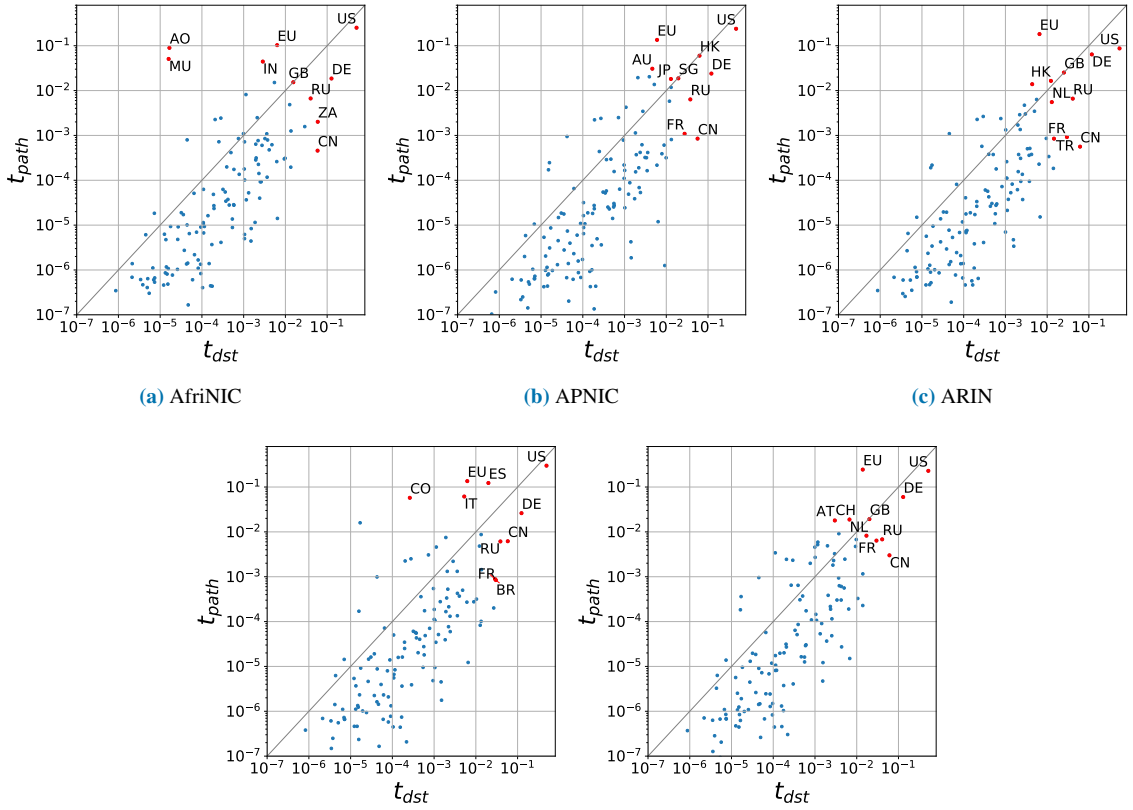

(a) AfriNIC **(b)** APNIC **(c)** ARIN

(d) LACNIC **(e)** RIPE NCC

Figure 5: t_{dst} and t_{path} of IPv4 importance of each country, divided by RIRs. The red points are the top 10 most important countries/regions by overall importance ($t_{dst} + t_{path}$). The x-axis and y-axis are in logarithmic scale.

dominance in providing domain hosting services and online content, and partially by its unique geographical and political location as the hub. The only exception is ARIN, where t_{path} of US is 0.087, which is less than 0.182 of European Union (EU)[2]. Surprising at first glance, the observation is reasonable because t_{path} considers the nodes except the start and the end, while a large portion of DNS requests in ARIN is originated or ended in the US without going to another country. Among all the RIRs, the destination-wise impact of the US is least in APNIC, which can be explained by that countries/regions like China (CN) and Hong Kong (HK) have a significant market share in Internet business locally. On the other hand, the US has a dominant impact on AfriNIC, with overall importance of 0.758 ($t_{dst} + t_{path}$), meaning that for

[2]Some ASes are affiliated with EU instead of individual countries.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2021.3112926, IEEE Access

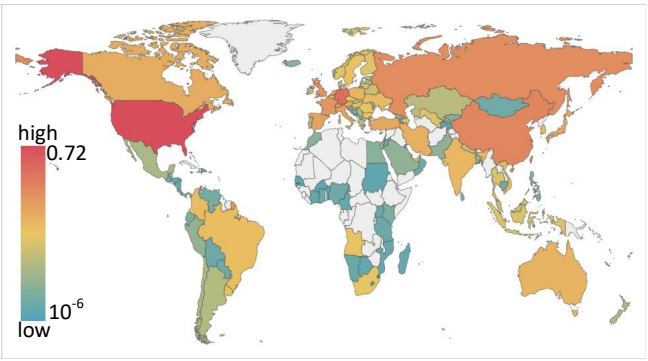

Figure 6: World HeatMap about country-wise importance (in logarithmic scale) aggregated across RIRs.

an African user, queries to 76% domains could be impacted by US. We speculate this is because 1) fewer local ANS are deployed in Africa, 2) African countries lack local IXP and have to rely on the US and 3) the small number of routers covered by our dataset. Followed by US, European countries like Germany (DE) also have strong impacts across RIRs. Among countries under AfriNIC, APNIC, and LACNIC, we found South Africa (ZA), Hong Kong (HK), and Brazil (BR) have strong local impacts.

Secondly, we find that **for the different regions, some countries have a greater local impact.** For instance, in APNIC, countries/regions like Hong Kong (HK, 0.123[3]), China (CN, 0.057) and Singapore (SG, 0.038) have a big share in importance score. While in LACNIC, Spain(ES, 0.143), Italy(IT, 0.066) and Brazil (BR, 0.031) play important roles. The reason for such diversity may be due to differences in user interests (large t_{dst} such as China), Internet infrastructure (large t_{path} such as Spain) or both.

Thirdly, as shown in Figure 5, **some countries such as China and Russia have a large t_{dst} but very small t_{path}, meaning that while they host a large number of ANS, they do not have significant impacts on DNS resolution paths.** Though, as described in Section II-C, these countries might enforce network sovereignty, our result indicates the resolution of domains outside those countries will not be largely impacted. Figure 6 also illustrates the overall importance of every country[4]. Generally speaking, developed countries are more "developed" even in DNS ecosystem with the exception of BRICS (Brazil, Russia, India, China and South Africa).

B. IPV6 COUNTRY IMPORTANCE

Though the push for broader IPv6 adoption is strong, section III-B shows yet the ANS support of IPv6 is disproportional to IPv4 (e.g., 11.7x and 8.7x IPv4 addresses associated with parent and child zones of the global list comparing to IPv6). On the other hand, the IPv6 coverage of domain resolution (57% for global list) is considerably better than that of user devices (24%) [45] and websites (20%) [46]. We are interested in how each country performs in the IPv6 space and we apply the same measurement method on the IPv6 data.

Table 5 lists the top 10 countries/regions in a similar way as Table 4. US still tops the overall impact ($t_{dst} + t_{path}$) in all RIRs, to more extend. The countries local to each RIR except US, DE, and EU have less impact in the DNS IPv6 space. Deployment of anycast, path diversity brought by direct peering [47], and relatively balanced ANS distribution could be the reasons why IPv4 shows a more diverse picture. We encourage more effort from the community to deploy IPv6 Anycast and more IPv6 ANS to enhance both the user's experience and the ANS's robustness.

Each country's impact on an RIR differs, with some countries having a much stronger presence than others, and we want to quantify the gap within RIR. To this end, we compute the *Gini coefficient* [48] among countries, which is the most widely used measure of inequality in economics. It has been leveraged to measure the inequality of social networks, e-commerce, and digital divide as well [49]–[51]. Assume $t_i, 1...n$ are the importance scores of n countries for users in an RIR, below is the equation we use to compute the Gini coefficient. Overall importance, t_{dest}, and t_{path} are computed separately.

$$G = \frac{\sum_{i=1}^{n}\sum_{i=1}^{n}|t_i - t_j|}{2n\sum_{i=1}^{n} t_i}$$

If the Gini coefficient equals 0, it means total equality among countries. On the contrary, if the Gini coefficient equals 1, it means one country has full control of the entire DNS in the region. As shown in Table 6, the Gini coefficients are all relatively high (over 0.9) for both IPv4 and IPv6 of all 5 RIRs, meaning the inequality gap is prominent. Comparing to IPv4, IPv6 has even larger Gini coefficients, which can be explained by the vastly different investment each country spends in IPv6 development. To make DNS infrastructure more robust, such inequality should be addressed with continuous efforts from the Internet community.

C. COUNTRY'S INFLUENCE ON DOMAINS

After examining a country's impact on the entire DNS, we drill down to the level of the individual domain. A domain could be influenced by a country on several levels. If a domain's ANS is located in a country, then all requests will be affected no matter where they come from. The resolution of this domain could be disrupted if the country chooses to cut off the links from the Internet. The country on the path can also eavesdrop or manipulate DNS packets when encryption is not enforced, which is still the dominant case [52]. Therefore, we break down a country's influence on domain names into four levels and measure them separately:

- `absolute`: all paths to all ANS of a domain are directed towards that country (this domain will not be resolvable after Internet cut-off).
- `semi-absolute`: excluding `absolute`, a country

[3]overall importance calculated by $t_{dst} + t_{path}$

[4]White countries are due to incompleteness of data sources.

AfriNIC			APNIC			ARIN			LACNIC			RIPE NCC		
CC	t_{dst}	t_{path}	CC	t_{dst}	t_{path}	CC	t_{dst}	t_{path}	CC	t_{dst}	t_{path}	CC	t_{dst}	t_{path}
US	0.616	0.316	US	0.572	0.322	US	0.599	0.110	US	0.590	0.370	US	0.600	0.305
DE	0.180	0.005	DE	0.184	0.009	DE	0.194	0.026	DE	0.179	0.007	EU	0.013	0.221
EU	0.010	0.056	EU	0.010	0.151	EU	0.010	0.172	EU	0.009	0.135	DE	0.188	0.030
GB	0.004	0.046	CN	0.045	0.000	CN	0.047	0.000	ES	0.012	0.111	CN	0.047	0.000
CN	0.047	0.000	RU	0.028	0.005	RU	0.029	0.004	CN	0.047	0.000	RU	0.029	0.003
RU	0.029	0.004	JP	0.006	0.020	FR	0.026	0.001	IT	0.001	0.044	FR	0.026	0.002
FR	0.026	0.001	FR	0.025	0.001	GB	0.004	0.018	RU	0.028	0.005	UA	0.005	0.017
HK	0.006	0.011	HK	0.006	0.013	HK	0.006	0.012	FR	0.026	0.001	CH	0.002	0.017
VN	0.012	0.000	RO	0.000	0.014	BR	0.003	0.014	HK	0.006	0.008	NL	0.011	0.005
NL	0.010	0.002	VN	0.012	0.000	AU	0.002	0.011	UY	0.013	0.000	AT	0.003	0.012

Table 5: Top 10 most important countries/regions of IPv6 DNS. The settings are the same as Table 4.

	IPv4			IPv6		
RIR	G_o	G_d	G_p	G_o	G_d	G_p
AfriNIC	0.916	0.931	0.954	0.964	0.965	0.977
APNIC	0.915	0.925	0.947	0.959	0.961	0.973
ARIN	0.920	0.931	0.952	0.955	0.963	0.964
LACNIC	0.917	0.924	0.957	0.955	0.959	0.971
RIPE NCC	0.918	0.927	0.947	0.956	0.962	0.965

Table 6: Global Gini coefficients of IPv4/IPv6 Internet separated by 5 RIRs. G_o, G_d and G_p are computed on overall importance, t_{dest} and t_{path} respectively.

appears in every path to all ANS of a domain[5] (the country has the ability to inspect and manipulate all queries about this domain).

- `influential`: excluding the above two cases, a country appears in at least one path to domain's ANS.
- `none`: a country does not appear in any path.

We combine all IPv4/IPv6 paths from all 5 RIRs to compute every country's influence level on every domain in the global list. We find `absolute` and `influential` are the main reasons when a domain is influenced by a country. Table 7 lists top countries ranked by `absolute` domain counts. US is leading with 681,969 out of 1.3 million domains, followed by Germany (59,956), China (54,377), Russia (50,618), France (29,873) and Japan (19,277). These countries have domestic Internet services well developed to serve their residents' needs. If many domains are hosted in a country or region, it will also be important.

Next, we want to answer this "what-if" question: when a country isolates itself from the Internet, how much impact will be introduced to the DNS ecosystem? We investigate a list of countries which once was found cutting off the Internet [4], [5], [53]–[55] and their influence levels, some of them are listed in Table 7, too. Among these countries with a history of "Internet cut-off", China (709,899), Russia (328,389), and India (767,454) are the top 3 in `influential`. Though not very important path-wise, countries like China, Russia and India has absolute influence on plenty of domains, and could influence even more. To notice, though the number of domains characterized as `influential` is large, many of the domains should still be accessible even when the top countries cut off the Internet, as they have multiple routes available. Iran, on the other hand, has more domains characterized as `absolute` than `influential`, indicating that Iran's network is relatively isolated from the world already. The influences of the surveyed countries vary. Even if some of the services might be viewed by local users mostly, the impact cannot be ignored if China, Russia, Iran, and India enforce network isolation.

Country	#absolute	#semi	#influential
United States	681,926	115	692,756
Germany	59,956	1,925	1,312,804
China	54,377	5	709,899
Russia	50,618	4	328,389
France	29,873	2	832,646
Japan	19,277	3	1,027,612
Turkey	18,881	15	6,972
Iran	12,376	140	6,428
Netherlands	11,074	227	989,791
Great Britain	9,987	7	1,324,062

Table 7: Country's influence on domains ranked by "absolute" domain count. The numbers are domain count.

[5]The first country in the path longer than 1 is not counted because it is the vantage point.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2021.3112926, IEEE Access

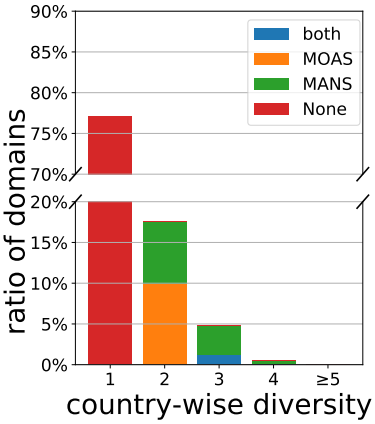

(a) Global list

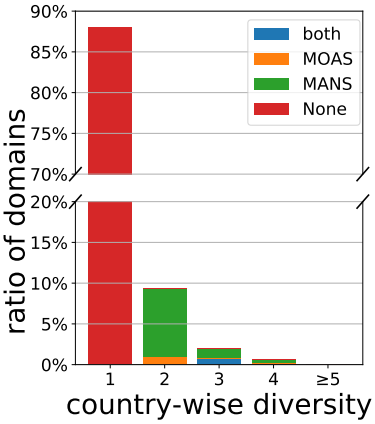

(b) Alexa top 1k

Figure 7: Country-wise diversity versus the ratio of domains. "MOAS", "MANS", "both" and "none" characterize the reason behind the country distribution.

D. ANS DEPLOYMENT PATTERNS

The previous measurement tasks investigate the impact of countries on the resolution paths. In this section, we change the view from country to domain and analyze the preferences of domain owners in installing ANS into the zone files. The deployment pattern of ANS has been measured by previous work [8] and we complement this work by adding another view about countries: we compute *country-wise diversity*, or the number of countries associated with all ANS, for each domain.

Figure 7a illustrates the country-wise diversity of domain names in the global list (1,380,395 domains) and we combine the IPv4 and IPv6 data. We found 77% domains placed all ANS in one country, which can be vulnerable when the country's Internet is disrupted.

We further characterize the reason behind the country distribution, which is also illustrated in Figure 7a. Two prominent reasons are identified: *Multiple Origin AS (MOAS) [56]* and *multiple ANS (MANS)*. MOAS happens when an IP of an ANS is announced by multiple ASes, which is usually caused by IP Anycast, load-balancing with multi-homing [57] or misconfiguration of routers [56]. MANS occurs when the domain owner intentionally installs multiple ANS in different countries. We classify a domain with multiple ANS into MOAS, MANS, and both, using BGP and AS-to-Country data. We found MANS is the dominant category when more than two countries are involved and it can be explained by the use of third-party DNS providers. For instance, many domain names are hosted by `DNS.COM`, a Chinese company. The address of its ANS `218.98.111.0/24` is published by at least two ASes, one is `AS 21859`, an American network, the other is `AS 133775`, a network in China. Thus requests around the world have the chance to be routed to either one of them. A similar phenomenon can be observed on `Godaddy`'s DNS. When two countries are involved, MOAS has a comparable ratio as MANS.

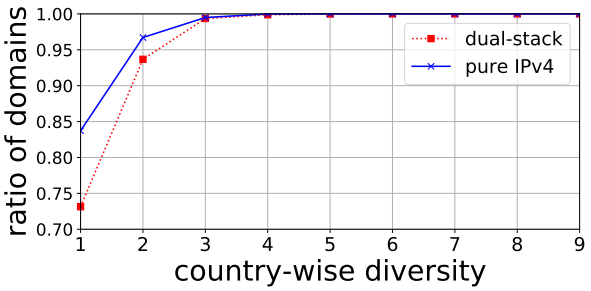

Figure 8: CDF of the country-wise diversity for domains supporting dual-stack/pure IPv4 resolution.

Then we look into the details about domain names with dual-stack name resolution (supporting both IPv4/IPv6) [58]. Out of 806,361 domains with at least one IPv6 ANS, we found the country-wise diversity for 27.2% domains is more than 1, and the ratio is higher than domains with only IPv4 ANS (16.4%). Still, the majority (72.8%) of dual-stack domain names have all of IPv4 and IPv6 ANS located in one country. On the other hand, for 17,114 domains, at least one new country not covered by the IPv4 ANS is introduced by its IPv6 ANS, suggesting they use a different set of ANS for IPv4 and IPv6 resolution. Yet, the number of IPv6 ANS is much fewer than IPv4 ANS and we recommend broadening the deployment of IPv6 ANS for more robust IPv6 and overall DNS resolution. Figure 8 compares the distribution of country-wise diversity between dual-stack and pure IPv4 resolution.

The prior measurement is carried out on the global list containing both popular domains and less popular domains. We are interested in whether similar country-wise diversity is observed in the very popular domains. To this end, we select a subset of top 1K domains from the Alexa list and the measurement result is illustrated in Figure 7b. It turns out the top domains are more likely to concentrate their ANS: about 88% domains use one country for ANS, comparing to 77% of all domains in the global list. A possible reason is

that these sites prefer self-hosting, or being served by DNS providers that do not support country diversity. When more than one country hosts ANS, MANS is the major reason.

E. COUNTRY PATH AND COUNTRY LOOP

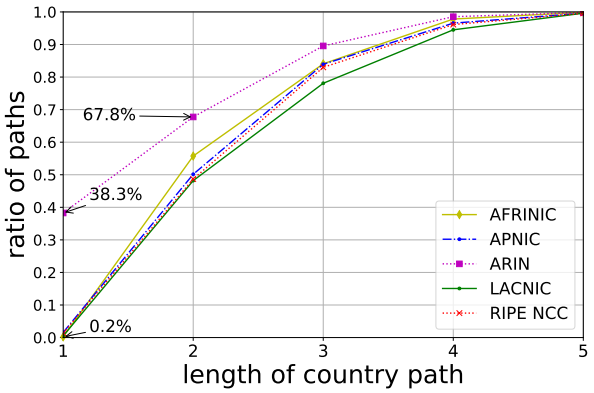

Figure 9: CDF of the length of country paths, separated by RIR.

With Country Path Finder of `DNSWeight`, we are able to reconstruct the DNS resolution routes with countries being the nodes. We are interested in the path statistics, particularly how many countries a DNS request has to come across utill reaching ANS. Figure 9 shows the distribution. The average country path length is 2.6, meaning that a DNS request needs to travel 2.6 countries on average to be responded to. It suggests a high degree of interconnection between countries. The distribution is quite similar for all RIRs except ARIN, of which 67.8% paths involve at most two countries and 38.3% paths involve one only country (we call them *domestic paths*). As a vast number of paths have to go through the US, shorter paths within ARIN are observed. This result also matches our findings about US's high country importance (Section IV-A and IV-B). On the contrary, only 0.2% domestic paths within the AfriNIC indicates its high reliance on other countries to fulfill DNS resolutions. To reduce the latency and improve the reliability of DNS resolution, these countries can deploy more local DNS services.

Among the country paths, *country loops*, which travel the same country/region twice, deserve more attention because they could introduce inefficient or privacy-violating routing [59]. Previous research uses BGP or traceroute data to identify the loops solely on the routing plane [14], [59]. Our study extends it to the DNS resolution setting. Specifically, we analyze all 78,904,677 country paths (IPv4 and IPv6 together) collected from 2,035 routers and identify the ones with repeated country/region nodes. In total, 2,997,486 paths are detected, constituting 4.6% of all paths. Then, we select the ones with the same starting and ending country/region node and obtain 1,372,231 paths from 1,274 routers, about 1.7% of all paths. They are considered as the country loops of our interest and we divide them based on the RIR of the country nodes. The numbers are shown in Table 8.

Though the ratio of country loops constitutes only a small portion of all country paths, they persist in 62.6% routers (1,274 out of 2,035). We find that 62.1% of country loops follow the path "ARIN -> RIPE NCC -> ARIN", suggesting the strong connection between Europe and North America. And some countries such as Poland, Japan, and Australia have over 90% of routers which keep at least one country loop.

Country Loop Type	Number
ARIN -> RIPE NCC -> ARIN	853,073
RIPE NCC -> ARIN -> RIPE NCC	156,273
Inside RIPE NCC	138,508
ARIN -> APNIC -> ARIN	73,144
ARIN -> LACNIC -> ARIN	31,741
Other	119,492

Table 8: Types and numbers of country loops.

To further analyze the reasons behind country loops, we cross-check some of the path data with PeeringDB [36], a database storing the geographic locations of facilities of exchange points. We selected the paths of which every hop has a match in PeeringDB. Then the geo-location of every hop could be inferred from the facility location of the exchange point. Our study find that many country loops (15,636 different paths in our data) are possibly located in a single country, since every hop as well as the start and the end of the path has at least one facility in the same country/region. However, this could still be an issue because, as we discuss in section III-C, a router is managed by its network owner and thus could be influenced by the policy of the country behind it, e.g., under network surveillance. On the other hand, a portion of loops goes across countries if no private peering is assumed. For instance, European countries are well connected to each other, and many loops (202 in our data) have emerged. Frequent data exchange between Hong Kong and Singapore sometimes leads to loops (28 in our data) as well. And we also find some "Canada-US-Canada" loops (157 in our data), which echo previous works [14].

F. LOCAL LIST

We use the global list for the prior measurement tasks. However, country-wise importance could differ when a different set of domains is inspected. To quantify the impact of domain selection, we switch from the global list to the local list described in Section III-B, which consists of 1,048,575 domains encountered by a public resolver in China. The routes to those domains are collected and we use `DNSWeight` to construct the country paths. We reduce the routers to the ones within APNIC to understand the local impact.

In Table 9 we show the top 10 countries based on t_{dst} and t_{path}. Unsurprisingly, the country has the maximum t_{dst} is changed from the US to China (CN), which scores 0.585 for IPv4 and 0.542 for IPv6. Other Asian countries/regions like Hong Kong (HK) and Singapore (SG) rise as well. The

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/ACCESS.2021.3112926, IEEE Access

IPv4				IPv6			
CC	t_{dst}	CC	t_{path}	CC	t_{dst}	CC	t_{path}
CN	0.585	US	0.265	CN	0.542	US	0.611
US	0.193	EU	0.139	US	0.286	EU	0.095
HK	0.068	HK	0.081	DE	0.054	HK	0.045
DE	0.034	SG	0.044	HK	0.041	JP	0.008
SG	0.030	AU	0.034	FR	0.006	SG	0.007
FR	0.006	KR	0.033	EU	0.005	DE	0.006
NL	0.006	CH	0.023	NL	0.003	BE	0.005
GB	0.005	JP	0.023	JP	0.002	RO	0.005
JP	0.004	DE	0.010	AT	0.002	CH	0.001
EU	0.003	IT	0.004	BG	0.001	NL	0.001

Table 9: Top 10 countries based on local list, ordered by t_{dst} and t_{path} separately.

IPv4				IPv6			
CC	t_{dst}	CC	t_{path}	CC	t_{dst}	CC	t_{path}
CN	0.690	US	0.318	CN	0.574	US	0.617
US	0.094	HK	0.126	US	0.224	SG	0.149
SG	0.078	EU	0.122	EU	0.092	HK	0.063
HK	0.075	SG	0.042	HK	0.050	EU	0.051
EU	0.023	AU	0.040	DE	0.011	JP	0.010
DE	0.003	KR	0.038	BG	0.002	DE	0.010
NL	0.001	JP	0.025	FR	0.001	RO	0.005
FR	0.001	CH	0.025	NL	0.001	CH	0.004
RU	0.001	CN	0.011	RU	0.001	PL	0.002
GB	0.001	DE	0.008	SE	0.001	GB	0.001

Table 10: Top 10 countries based on local list, considering domain popularity, ordered by t_{dst} and t_{path} separately.

result shows the popular regional domains prefer to use the ANS within the same region. On the other hand, this result should be taken a grain of salt. As China government actively censors Chinese users' Internet traffic [18], Internet services in China differ greatly from the rest of the world, which introduces vast differences between the global and local lists The conclusion could differ when other countries' list is used.

In contrast to t_{dst}, we found t_{path} is less influenced by the change of domain list. Specifically, for IPv4, the ten most important countries ordered by t_{path} in Table 9 are also the ten most important countries for global list. For IPv6, the overlap is nine. The US tops the ranking with 0.265 for IPv4 and 0.611 for IPv6. The result shows the routing decisions are similar across different sets of domains. Due to restrictions of space and data we choose not to perform measurements with vantage points and targets in more countries/regions. This is left for future studies.

In our previous experiments, we treated all domains in the target list equally in order to represent the impact of a region on the entire domain name system. In this setting, a region that can influence more domains receives a greater importance score. On the other hand, we can also consider the different levels of importance of each domain name. As introduced in section III-D, Previous research [40] showed that user access to domain names followed a power-law distribution. This implied that more users would visit the most popular domains. For the domains in the local list, we collected the total number of times they were queried. In the following experiment, we weight and normalize each domain by its query count and use this to calculate the importance score. A region that can influence more queries or more popular domains receives a greater importance score in this setting.

The result is shown in table 10. In general, the regions/countries that were previously more important in the algorithm remain essential, such as China (CN), United States (US), Hong Kong (HK) and Singapore (SG). Thus, it suggests that the regions that affect more domain queries generally also affect more domain names. Comparing to IPv6, IPv4 scores are more concentrated in China because many popular websites in the list are used by Chinese users. Their DNS is also deployed in China. Our further fine-grained analysis reveals that for IPv4 authoritative name servers, 97.6% of the importance score is related to just 1% of the IP addresses. This result is consistent with previous studies. [8] Overall, multiple methods of importance calculation are consistent in their description of reality.

V. RELATED WORKS

Our study measures the importance of a country in the lens of DNS. Below we review the measurement studies focusing on routing and DNS first and then the ones using both data.

A. MEASUREMENT OF ROUTING

To measure the dynamics of Internet routes, two data sources are mainly used [60], including BGP [59], [61]–[67] and traceroute [14], [28], [68]–[71]. Traceroute requires active probing, which is not applicable for our large-scale analysis [28], [71]. In addition, prior studies show that AS information derived from the traceroute data could be inaccurate [72], [73]. As such, we use BGP data, which has better coverage of routes and accuracy. On the other hand, neither data source is perfect [60], [74], [75], and inconsistencies have been observed [76]. We plan to investigate how to augment the BGP data with traceroute data to obtain more precise results in the future.

Measuring the geographic characteristics of routing is the focus of previous studies. [65] measured country-to-country importance in routing, that is, how likely a country is on the routes between any two other countries. [66] and [67] evaluated AS-to-AS and AS-to-country importance similarly. Some works estimated country-to-web [28], [71] or AS-to-web [64] importance, that is, how like a country or an AS is on the routes for visiting popular websites. Routing detour is measured in [14], [28], [59], [71] and [77] found anycast traffic sometimes are routed to out-of-country PoP (Point of Presence) even when an in-country PoP is available.

Compared to the prior studies, the main contribution of this work is to offer new insights about country-to-DNS importance, which has never been investigated a priori. In addition, the number of websites and countries we select to assess the importance is significantly larger than prior works

using hundreds of popular sites and few countries [28], [64], [71].

B. MEASUREMENT OF DNS

Numerous works have been done to measure the DNS infrastructure with passive or active data collection. The first approach [78], [79] obtains DNS logs from DNS servers [77] and historical database [23], or downloads zone files [8]. The second approach issues DNS queries from vantage points against DNS servers for performance measurement or anomaly identification. Platforms like RIPE Atlas [9], [10], [21], [80], [81], proxy networks [2], [82] and ad networks [83] were leveraged as crowd-sourced vantage points. In addition, researchers use open resolvers, which can be identified through scanning IPv4 address space [3], [39], [84], to forward DNS requests and conduct active measurement [26], [80], [85].

Our study attempts to assess country importance based on the distribution of nameservers. Previous works have extensively measured nameservers, focusing the aspects like performance [79], [81], [86], [87], security [30], [88], privacy [52], [89], configuration issues [8], [90] and record inconsistencies [9], [10], [21]. Though DNS is designed as a distributed system for reliability, recent studies revealed the trend of centralizing DNS services. For example, [8] showed that SLD names are increasingly sharing nameservers. [27], [91], [92] discovered that nameservers of popular websites run on a small number of hosting or cloud services. Interdependencies were identified between zone files [93], [94], which could damage the reliability of DNS potentially [95]. The results of our study complement prior works regarding the distribution of DNS services, showing that some countries have prominent impacts on the whole DNS infrastructure.

C. MEASUREMENT ON DNS AND ROUTING JOINTLY

To optimize the DNS resolution performance, the routes between users and the nameservers are heavily engineered. IP anycast is a technique leveraged to this end and it has been measured using passive or active analysis [77], [96], [96]–[99]. On the other hand, route hijack against DNS has been discovered and DNS and routing data are combined to detect such incidents [1], [98], [100]–[102]. Yet, no prior study has assessed the importance of certain parties related to the DNS infrastructure in the lens of routes, and we make the first attempt.

VI. CONCLUSION

To quantify the importance of a country on the entire DNS infrastructure, we present `DNSWeight`, which analyzes DNS records and BGP routing data and computes destination-wise and path-wise scores. Measurement tasks that compare countries' importance based on regions, IP protocols (IPv4 and IPv6), domains, and paths can be fulfilled. We will release the source code of `DNSWeight` to help other researchers study issues related to DNS and routing.

Here we revisit the measurement results and highlight the key insights. Firstly, the importance among countries is quite unbalanced. The US plays the dominant role in DNS infrastructure in every region, with over 0.75 on AfriNIC, LACNIC, and RIPE NCC, while the second country/region only has less than 0.25. European countries like Germany also have a strong influence across RIRs consistently. The gap is even enlarged when IPv6 is inspected, with the US being able to reach over 0.9 overall RIRs except for ARIN. Such observation could be unique to the DNS infrastructure, as diversifying nameservers are recommended by IETF RFCs and the network infrastructure of US and other European countries is more likely to be relied on. Secondly, countries with a history of network sovereignty have a significant impact on a large number of domains if they choose to isolate themselves from the Internet. Out of the 1.38 million domains we surveyed, China, Russia can achieve absolute control over the 50K domains, while China and India can influence the resolution of over 700K domains. This result shows the DNS world is highly connected, and DNS reliability needs to be reconsidered in the context of country politics. Thirdly, the routes of DNS resolution are far from being optimal, with the average length of a country path being 2.7 (except the US), and country loops are observed in 62.6% routers we surveyed.

While we are pleased to see the community's effort in making DNS more reliable (e.g., multiple nameservers for one domain), we are concerned about the inequality of investment into DNS infrastructure between countries and the impact of network sovereignty potentially by certain countries. The issues with DNS routing should also be addressed to improve users' DNS experiences. We believe country-wise importance should be considered an essential factor when structuring DNS infrastructure and new research is warranted.

References

[1] Ben Jones, Nick Feamster, Vern Paxson, Nicholas Weaver, and Mark Allman. Detecting dns root manipulation. In Thomas Karagiannis and Xenofontas Dimitropoulos, editors, Passive and Active Measurement, pages 276–288, Cham, 2016. Springer International Publishing.

[2] Baojun Liu, Chaoyi Lu, Haixin Duan, Ying Liu, Zhou Li, Shuang Hao, and Min Yang. Who is answering my queries: Understanding and characterizing interception of the dns resolution path. In Proceedings of the 27th USENIX Conference on Security Symposium, SEC'18, page 1113–1128, USA, 2018. USENIX Association.

[3] Paul Pearce, Ben Jones, Frank Li, Roya Ensafi, Nick Feamster, Nick Weaver, and Vern Paxson. Global measurement of dns manipulation. In Proceedings of the 26th USENIX Conference on Security Symposium, SEC'17, page 307–323, USA, 2017. USENIX Association.

[4] BBC. Russia 'successfully tests' its unplugged internet. https://www.bbc.com/news/technology-50902496, 2019.

[5] New York Times. After long ban, western china is back online. https://www.nytimes.com/2010/05/15/world/asia/15china.html, 2010.

[6] SFLC.IN. Internet shutdown tracker. https://internetshutdowns.in/, 2020.

[7] Steve Gibbard and P House. Geographic implications of dns infrastructure distribution. The Internet Protocol Journal, 10(1):12–24, 2007.

[8] Mark Allman. Comments on dns robustness. In Proceedings of the Internet Measurement Conference 2018, IMC '18, page 84–90, New York, NY, USA, 2018. Association for Computing Machinery.

[9] Raffaele Sommese, Giovane C. M. Moura, Mattijs Jonker, Roland van Rijswijk-Deij, Alberto Dainotti, K. C. Claffy, and Anna Sperotto. When

parents and children disagree: Diving into dns delegation inconsistency. In Anna Sperotto, Alberto Dainotti, and Burkhard Stiller, editors, Passive and Active Measurement, pages 175–189, Cham, 2020. Springer International Publishing.

[10] G. C. M. Moura, J. Heidemann, R. D. Schmidt, W. Hardaker, and Machinery Assoc Comp. Cache me if you can: Effects of dns time-to-live. In Imc'19: Proceedings of the 2019 Acm Internet Measurement Conference, pages 101–115, 2019.

[11] NCC RIPE. Ripe routing information service raw data. https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data, 2020.

[12] David Meyer. Routeviews project. http://www.routeviews.org/routeviews/, 2004.

[13] IIT-CNR. Isolario project. https://www.isolario.it/, 2020.

[14] Jonathan A. Obar and Andrew Clement. Internet surveillance and boomerang routing: A call for canadian network sovereignty. 2013.

[15] Laura DeNardis. The global war for internet governance. Yale University Press, 2014.

[16] Phillipa Gill, Masashi Crete-Nishihata, Jakub Dalek, Sharon Goldberg, Adam Senft, and Greg Wiseman. Characterizing web censorship worldwide: Another look at the opennet initiative data. ACM Transactions on the Web (TWEB), 9(1):1–29, 2015.

[17] Simurgh Aryan, Homa Aryan, and J Alex Halderman. Internet censorship in iran: A first look. In Presented as part of the 3rd {USENIX} Workshop on Free and Open Communications on the Internet, 2013.

[18] Xueyang Xu, Z Morley Mao, and J Alex Halderman. Internet censorship in china: Where does the filtering occur? In International Conference on Passive and Active Network Measurement, pages 133–142. Springer, 2011.

[19] Yaman Akdeniz. Internet content regulation: Uk government and the control of internet content. Computer Law & Security Review, 17(5):303–317, 2001.

[20] Philip Levis. The collateral damage of internet censorship by dns injection. ACM SIGCOMM CCR, 42(3), 2012.

[21] Moritz Müller, Giovane CM Moura, Ricardo de O Schmidt, and John Heidemann. Recursives in the wild: Engineering authoritative dns servers (extended). SIDN Labs June, 21, 2017.

[22] Verisign. Zone files for top-level domains (tlds). https://www.verisign.com/en_US/channel-resources/domain-registry-products/zone-file/index.xhtml, 2020.

[23] Farsight Security. Sie data sharing. https://www.farsightsecurity.com/community/data-sharing/, 2019.

[24] Baojun Liu, Chaoyi Lu, Zhou Li, Ying Liu, Hai-Xin Duan, Shuang Hao, and Zaifeng Zhang. A reexamination of internationalized domain names: The good, the bad and the ugly. In DSN, pages 654–665, 2018.

[25] Alexa. Top sites. https://www.alexa.com/topsites, 2020.

[26] A. Klein, H. Shulman, and M. Waidner. Counting in the dark: Dns caches discovery and enumeration in the internet. In 2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), pages 367–378, 2017.

[27] S. Hao, H. Wang, A. Stavrou, and E. Smirni. On the dns deployment of modern web services. In 2015 IEEE 23rd International Conference on Network Protocols (ICNP), pages 100–110, 2015.

[28] Anne Edmundson, Roya Ensafi, Nick Feamster, and Jennifer Rexford. Nation-state hegemony in internet routing. In Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies, COMPASS '18, New York, NY, USA, 2018. Association for Computing Machinery.

[29] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. Tranco: A research-oriented top sites ranking hardened against manipulation. arXiv preprint arXiv:1806.01156, 2018.

[30] Daiping Liu, Shuai Hao, and Haining Wang. All Your DNS Records Point to Us Understanding the Security Threats of Dangling DNS Records. In CCS'16: PROCEEDINGS OF THE 2016 ACM SIGSAC CONFERENCE ON COMPUTER AND COMMUNICATIONS SECURITY, pages 1414–1425, 2016. 23rd ACM Conference on Computer and Communications Security (CCS), Vienna, AUSTRIA, OCT 24-28, 2016.

[31] Willibald Doeringer, Günter Karjoth, and Mehdi Nassehi. Routing on longest-matching prefixes. IEEE/ACM transactions on networking, 4(1):86–97, 1996.

[32] APNIC. Visibility of ipv4 and ipv6 prefix lengths in 2019. https://blog.apnic.net/2019/04/19/visibility-of-ipv4-and-ipv6-prefix-lengths-in-2019/, 2019.

[33] CAIDA. As rank: A ranking of the largest autonomous systems (as) in the internet. http://as-rank.caida.org/, 2020.

[34] RIPE NCC. List of country codes and rirs. https://www.ripe.net/participate/member-support/list-of-members/list-of-country-codes-and-rirs.

[35] Narine Badasyan, Subhadip Chakrabarti, et al. Private peering among internet backbone providers. Economics Working Paper Archive, Series on Industrial Organization, Washington University in St. Louis (WUSTL), 2003.

[36] PeeringDB. Peeringdb. https://www.peeringdb.com/, 2020.

[37] Linton C Freeman. A set of measures of centrality based on betweenness. Sociometry, pages 35–41, 1977.

[38] Matthias Wählisch, Thomas C Schmidt, Markus de Brün, and Thomas Häberlen. Exposing a nation-centric view on the german internet–a change in perspective on as-level. In International Conference on Passive and Active Network Measurement, pages 200–210. Springer, 2012.

[39] Marc Kührer, Thomas Hupperich, Jonas Bushart, Christian Rossow, and Thorsten Holz. Going wild: Large-scale classification of open dns resolvers. In Proceedings of the 2015 Internet Measurement Conference, IMC '15, page 355–368, New York, NY, USA, 2015. Association for Computing Machinery.

[40] Lee Breslau, Pei Cao, Li Fan, Graham Phillips, and Scott Shenker. Web caching and zipf-like distributions: Evidence and implications. In IEEE INFOCOM'99. Conference on Computer Communications. Proceedings. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. The Future is Now (Cat. No. 99CH36320), volume 1, pages 126–134. IEEE, 1999.

[41] Paul V Mockapetris. Rfc1034: Domain names-concepts and facilities, 1987.

[42] R Elz, R Bush, S Bradner, and M Patton. Rfc2182: Selection and operation of secondary dns servers, 1997.

[43] Madeline Carr. US power and the internet in international relations: The irony of the information age. Springer, 2016.

[44] Dwayne Winseck. The geopolitical economy of the global internet infrastructure. Journal of Information Policy, 7:228–267, 2017.

[45] apnic. Ipv6 users by country. https://labs.apnic.net/dists/v6dcc.html, 2020.

[46] Mike Leber. Global ipv6 deployment progress report. https://bgp.he.net/ipv6-progress-report.cgi, 2020.

[47] Muhammad Arif Wicaksana. Ipv4 vs ipv6 anycast catchment: A root dns study. Master's thesis, University of Twente, 2016.

[48] Amartya Sen, Master Amartya Sen, Sen Amartya, James Eric Foster, James E. Foster, and The Radcliffe lectures, 1972. On Economic Inequality. Clarendon Press, 1997.

[49] Giseli Rabello Lopes, Roberto da Silva, and J Palazzo M de Oliveira. Applying gini coefficient to quantify scientific collaboration in researchers network. In Proceedings of the International Conference on Web Intelligence, Mining and Semantics, pages 1–6, 2011.

[50] Gal Oestreicher-Singer and Arun Sundararajan. Recommendation networks and the long tail of electronic commerce. Mis quarterly, pages 65–83, 2012.

[51] Jayajit Chakraborty and M Martin Bosman. Measuring the digital divide in the united states: Race, income, and personal computer ownership. The Professional Geographer, 57(3):395–410, 2005.

[52] Chaoyi Lu, Baojun Liu, Zhou Li, Shuang Hao, Haixin Duan, Mingming Zhang, Chunying Leng, Ying Liu, Zaifeng Zhang, and Jianping Wu. An end-to-end, large-scale measurement of dns-over-encryption: How far have we come? In Proceedings of the Internet Measurement Conference, IMC '19, page 22–35, New York, NY, USA, 2019. Association for Computing Machinery.

[53] Wired. How the iranian government shut off the internet. https://www.bbc.com/news/world-asia-india-50819905, 2019.

[54] BBC. Why india shuts down the internet more than any other democracy. https://www.bbc.com/news/world-asia-india-50819905, 2019.

[55] HIJA KAMRAN. An internet shutdown is keeping coronavirus information from millions in pakistan. https://slate.com/technology/2020/04/coronavirus-covid19-pakistan-internet-shutdown-fata.html, 2020.

[56] Xiaoliang Zhao, Dan Pei, Lan Wang, Dan Massey, Allison Mankin, S Felix Wu, and Lixia Zhang. An analysis of bgp multiple origin as (moas) conflicts. In Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement, pages 31–35, 2001.

[57] Fanglu Guo, Jiawu Chen, Wei Li, and Tzi-cker Chiueh. Experiences in building a multihoming load balancing system. In IEEE INFOCOM 2004, volume 2, pages 1241–1251. IEEE, 2004.

[58] Scott Hogg. Why you should dual-stack your dns nameservers. https://blogs.infoblox.com/ipv6-coe/why-you-should-dual-stack-your-dns-nameservers/, 2018.

[59] Anant Shah, Romain Fontugne, and Christos Papadopoulos. Towards characterizing international routing detours. In Proceedings of the 12th Asian Internet Engineering Conference, AINTEC '16, page 17–24, New York, NY, USA, 2016. Association for Computing Machinery.

[60] Reza Motamedi, Reza Rejaie, and Walter Willinger. A Survey of Techniques for Internet Topology Discovery. IEEE COMMUNICATIONS SURVEYS AND TUTORIALS, 17(2):1044–1065, 2015.

[61] R. Govindan and A. Reddy. An analysis of internet inter-domain topology and route stability. In Proceedings of INFOCOM '97, volume 2, pages 850–857 vol.2, 1997.

[62] Beichuan Zhang, Raymond Liu, Daniel Massey, and Lixia Zhang. Collecting the internet as-level topology. SIGCOMM Comput. Commun. Rev., 35(1):53–61, January 2005.

[63] Phillipa Gill, Michael Schapira, and Sharon Goldberg. Modeling on quicksand: Dealing with the scarcity of ground truth in interdomain routing data. SIGCOMM Comput. Commun. Rev., 42(1):40–46, January 2012.

[64] H. B. Acharya, S. Chakravarty, and D. Gosain. Few throats to choke: On the current structure of the internet. In 2017 IEEE 42nd Conference on Local Computer Networks (LCN), pages 339–346, 2017.

[65] Josh Karlin, Stephanie Forrest, and Jennifer Rexford. Nation-state routing: Censorship, wiretapping, and bgp. ArXiv, abs/0903.3218, 2009.

[66] Romain Fontugne, Anant Shah, and Emile Aben. The (thin) bridges of as connectivity: Measuring dependency using as hegemony. In Robert Beverly, Georgios Smaragdakis, and Anja Feldmann, editors, Passive and Active Measurement, pages 216–227, Cham, 2018. Springer International Publishing.

[67] Kirtus G. Leyba, Benjamin Edwards, Cynthia Freeman, Jedidiah R. Crandall, and Stephanie Forrest. Borders and gateways: Measuring and analyzing national as chokepoints. In Proceedings of the 2nd ACM SIGCAS Conference on Computing and Sustainable Societies, COMPASS '19, page 184–194, New York, NY, USA, 2019. Association for Computing Machinery.

[68] Kai Chen, David R. Choffnes, Rahul Potharaju, Yan Chen, Fabian E. Bustamante, Dan Pei, and Yao Zhao. Where the Sidewalk Ends: Extending the Internet AS Graph Using Traceroutes from P2P Users. IEEE TRANSACTIONS ON COMPUTERS, 63(4):1021–1036, APR 2014.

[69] Yuval Shavitt and Eran Shir. Dimes: Let the internet measure itself. SIGCOMM Comput. Commun. Rev., 35(5):71–74, October 2005.

[70] Enrico Gregori, Luciano Lenzini, and Valerio Luconi. AS-Level Topology Discovery: Measurement strategies tailored for crowdsourcing systems. COMPUTER COMMUNICATIONS, 112:47–57, NOV 1 2017.

[71] Anne Edmundson, Roya Ensafi, Nick Feamster, and Jennifer Rexford. A first look into transnational routing detours. In Proceedings of the 2016 ACM SIGCOMM Conference, SIGCOMM '16, page 567–568, New York, NY, USA, 2016. Association for Computing Machinery.

[72] Yu Zhang, Ricardo Oliveira, Yangyang Wang, Shen Su, Baobao Zhang, Jun Bi, Hongli Zhang, and Lixia Zhang. A Framework to Quantify the Pitfalls of Using Traceroute in AS-Level Topology Measurement. IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, 29(9):1822–1836, OCT 2011.

[73] Pietro Marchetta, Valerio Persico, Antonio Pescapè, and Ethan Katz-Bassett. Don't trust traceroute (completely). In Proceedings of the 2013 Workshop on Student Workhop, CoNEXT Student Workhop '13, page 5–8, New York, NY, USA, 2013. Association for Computing Machinery.

[74] Randy Bush, Olaf Maennel, Matthew Roughan, and Steve Uhlig. Internet optometry: Assessing the broken glasses in internet reachability. In Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement, IMC '09, page 242–253, New York, NY, USA, 2009. Association for Computing Machinery.

[75] R. Oliveira, D. Pei, W. Willinger, B. C. Zhang, and L. X. Zhang. The (in)completeness of the observed internet as-level structure. Ieee-Acm Transactions on Networking, 18(1):109–122, 2010.

[76] Young Hyun, Andre Broido, and Kc Claffy. Traceroute and bgp as path incongruities. 2003.

[77] W. B. de Vries, R. Van Rijswijk-Deij, P. de Boer, and A. Pras. Passive observations of a large dns service: 2.5 years in the life of google. In 2018 Network Traffic Measurement and Analysis Conference (TMA), pages 1–8, 2018.

[78] Hongyu Gao, Vinod Yegneswaran, Jian Jiang, Yan Chen, Phillip Porras, Shalini Ghosh, and Haixin Duan. Reexamining dns from a global recursive resolver perspective. Ieee-Acm Transactions on Networking, 24(1):43–57, 2016.

[79] Pawel Foremski, Oliver Gasser, and Giovane C. M. Moura. Dns observatory: The big picture of the dns. In Proceedings of the Internet Measurement Conference, IMC '19, page 87–100, New York, NY, USA, 2019. Association for Computing Machinery.

[80] W. B. Vries, Q. Scheitle, M. Muller, W. Toorop, R. Dolmans, and R. Rijswijk-Deij. A first look at qname minimization in the domain name system. In Passive and Active Measurement. 20th International Conference, PAM 2019. Proceedings: Lecture Notes in Computer Science, pages 147–60, 2019.

[81] Z. Y. Gao, A. Venkataramani, and Ieee. Measuring update performance and consistency anomalies in managed dns services. In Ieee Conference on Computer Communications, Ieee Infocom, pages 2206–2214, 2019.

[82] Taejoong Chung, Roland van Rijswijk-Deij, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson. A Longitudinal, End-to-End View of the DNSSEC Ecosystem. In PROCEEDINGS OF THE 26TH USENIX SECURITY SYMPOSIUM (USENIX SECURITY '17), pages 1307–1322, 2017. 26th USENIX Security Symposium, Vancouver, CANADA, AUG 16-18, 2017.

[83] P. Callejo, R. Cuevas, N. Vallina-Rodriguez, and Á. Cuevas Rumin. Measuring the global recursive dns infrastructure: A view from the edge. IEEE Access, 7:168020–168028, 2019.

[84] J. Park, A. Khormali, M. Mohaisen, and A. Mohaisen. Where are you taking me? behavioral analysis of open dns resolvers. In 2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), pages 493–504, 2019.

[85] Hitesh Ballani, Paul Francis, and Sylvia Ratnasamy. A measurement-based deployment proposal for ip anycast. In Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement, IMC '06, page 231–244, New York, NY, USA, 2006. Association for Computing Machinery.

[86] Jinjin Liang, Jian Jiang, Haixin Duan, Kang Li, and Jianping Wu. Measuring query latency of top level dns servers. In Matthew Roughan and Rocky Chang, editors, Passive and Active Measurement, pages 145–154, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[87] Zheng Wang. Understanding the Performance and Challenges of DNS Query Name Minimization. In 2018 17TH IEEE INTERNATIONAL CONFERENCE ON TRUST, SECURITY AND PRIVACY IN COMPUTING AND COMMUNICATIONS (IEEE TRUSTCOM) / 12TH IEEE INTERNATIONAL CONFERENCE ON BIG DATA SCIENCE AND ENGINEERING (IEEE BIGDATASE), IEEE Trustcom BigDataSE ISPA, pages 1115–1120, 2018. 17th IEEE International Conference on Trust, Security and Privacy in Computing and Communications (IEEE TrustCom) / 12th IEEE International Conference on Big Data Science and Engineering (IEEE BigDataSE), New York, NY, JUL 31-AUG 03, 2018.

[88] Thomas Vissers, Timothy Barron, Tom Van Goethem, Wouter Joosen, and Nick Nikiforakis. The Wolf of Name Street: Hijacking Domains Through Their Nameservers. In CCS'17: PROCEEDINGS OF THE 2017 ACM SIGSAC CONFERENCE ON COMPUTER AND COMMUNICATIONS SECURITY, pages 957–970, 2017. 24th ACM-SIGSAC Conference on Computer and Communications Security (ACM CCS), Dallas, TX, OCT 30-NOV 03, 2017.

[89] Casey Deccio and Jacob Davis. Dns privacy in practice and preparation. In Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies, CoNEXT '19, page 138–143, New York, NY, USA, 2019. Association for Computing Machinery.

[90] V. Pappas, D. Wessels, D. Massey, S. Lu, A. Terzis, and L. Zhang. Impact of configuration errors on dns robustness. IEEE Journal on Selected Areas in Communications, 27(3):275–290, 2009.

[91] Samantha Bates, John Bowers, Shane Greenstein, Jordi Weinstock, Yunhan Xu, and Jonathan Zittrain. Evidence of decreasing internet entropy: The lack of redundancy in dns resolution by major websites and services. SSRN Electronic Journal, 2018.

[92] Matteo Dell'Amico, Leyla Bilge, Ashwin Kayyoor, Petros Efstathopoulos, and Pierre-Antoine Vervier. Lean on me: Mining internet service dependencies from large-scale dns data. In Proceedings of the 33rd Annual Computer Security Applications Conference, page 449–460, 2017.

[93] Haiyan Xu, Zhaoxin Zhang, Jianen Yan, and Xin Ma. Evaluating the impact of name resolution dependence on the dns. Security and Communication Networks, 2019:1–12, 2019.

[94] Jian Jiang, Jia Zhang, Haixin Duan, Kang Li, Wu Liu, and Ieee. Analysis and measurement of zone dependency in the domain name system. In 2018 Ieee International Conference on Communications, IEEE International Conference on Communications, 2018.

[95] Venugopalan Ramasubramanian and Emin Gün Sirer. Perils of transitive trust in the domain name system. In Proceedings of the 5th ACM SIGCOMM conference on Internet measurement, page 35, 2005.

[96] Zhihao Li, Dave Levin, Neil Spring, and Bobby Bhattacharjee. Internet Anycast: Performance, Problems, & Potential. In PROCEEDINGS OF THE 2018 CONFERENCE OF THE ACM SPECIAL INTEREST GROUP ON DATA COMMUNICATION (SIGCOMM '18), pages 59–73, 2018. Conference of the ACM-Special-Interest-Group-on-Data-Communication (ACM SIGCOMM), Budapest, HUNGARY, AUG 20-25, 2018.

[97] Ricardo de Oliveira Schmidt, John Heidemann, and Jan Harm Kuipers. Anycast Latency: How Many Sites Are Enough? In Kaafar, MA and Uhlig, S and Amann, J, editor, PASSIVE AND ACTIVE MEASUREMENT (PAM 2017), volume 10176 of Lecture Notes in Computer Science, pages 188–200, 2017. 18th International Conference on Passive and Active Measurement (PAM), Sydney, AUSTRALIA, MAR 30-31, 2017.

[98] X. Fan, J. Heidemann, and R. Govindan. Evaluating anycast in the domain name system. In 2013 Proceedings IEEE INFOCOM, pages 1681–1689, 2013.

[99] Stephen McQuistin, Sree Priyanka Uppu, and Marcel Flores. Taming anycast in the wild internet. In Proceedings of the Internet Measurement Conference, IMC '19, page 165–178, New York, NY, USA, 2019. Association for Computing Machinery.

[100] Lan Wang, Xiaoliang Zhao, Dan Pei, R. Bush, D. Massey, and Lixia Zhang. Protecting bgp routes to top-level dns servers. IEEE Transactions on Parallel and Distributed Systems, 14(9):851–860, 2003.

[101] I. Avramopoulos and M. Suchara. Protecting the dns from routing attacks: Two alternative anycast implementations. IEEE Security Privacy, 7(5):14–20, 2009.

[102] Giovane C.M. Moura, Ricardo de O. Schmidt, John Heidemann, Wouter B. de Vries, Moritz Muller, Lan Wei, and Cristian Hesselman. Anycast vs. ddos: Evaluating the november 2015 root dns event. In Proceedings of the 2016 Internet Measurement Conference, IMC '16, page 255–270, New York, NY, USA, 2016. Association for Computing Machinery.

DELIANG CHANG received the master's degree in engineering from Tsinghua University, Beijing, China, in 2013, where he is currently pursuing the Ph.D. degree.

His research interests include network measurement and network security.

SHANSHAN HAO received the B.S. degree in microelectronic science and engineering from Tsinghua University, Beijing, China, in 2017, where she is currently pursuing the M.S. degree in information and communication engineering.

Her research interest includes improving domain name system security based on blockchain technology and IPv6 development.

ZHOU LI received PhD in computer science from Indiana University Bloomington.

He was a principal research scientist at RSA Labs from 2014-2018. He is currently an assistant professor at EECS department of University of California, Irvine. He has published more than 40 refereed research articles. His research interests include cyber-security, privacy and machine learning. He is an IEEE senior member.

BAOJUN LIU received the B.E. degree from Xidian University, China, in 2015, and received his PhD degree from the Department of Computer Science and Technology, Tsinghua University in 2020.

He is currently a postdoctoral researcher at the Institute for Network Science and Cyberspace, Tsinghua University. His research interests include DNS security and Internet measurement.

XING LI received the B.S. degree in radio electronics from Tsinghua University, Beijing, China, in 1982, and the M.S. and Ph.D. degrees in electrical engineering from Drexel University, USA, in 1985 and 1989, respectively.

He is currently a Professor with the Electronic Engineering Department, Tsinghua University. He is the Deputy Director of the China Education and Research Network (CERNET) Center. He published more than 300 articles in his research areas and the coauthor of 11 IETF RFCs. His research activities and interests include compute networks, multimedia communications, and statistical signal processing. He was a member of Internet Architecture Board (IAB), the Co-Chair of the Coordinating Committee for Intercontinental Research Networking (CCIRN), the Chair of the Asia pacific Network Group (APNG), a member of the APNIC Executive Council, and the Chair of Internet Sub Chapter of Computer Society of China (CCF).