Visualizing the African Internet Topology from Measurement Data

Blessed Chitamba*
chtble001@myuct.ac.za
University of Cape Town

ABSTRACT

This paper looks at internet topology visualization in the African context, by looking at the different steps involved from raw data collected from internet probing tools to having an entire internet topology that can then be used in visualization tools. We start by analysing what has been done by other researchers, comparing and contrasting the different literature in order to find unanswered questions and research gaps that still need attention. Next, we will then look at one chosen unanswered question to identify how it can be answered.

KEYWORDS

internet topology, traceroute, graphs, alias resolution, Autonomous System

ACM Reference Format:
Blessed Chitamba. 2020. Visualizing the African Internet Topology from Measurement Data. In *Proceedings of ACM Conference (Conference'17).* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

The world now heavily relies on the internet, and it has become the biggest network of interconnected entities. Its continued growth brings new challenges and problems that deserve research attention. One of such problems is properly mapping and visualizing the topology of the internet at any current point. This problem becomes more interesting when we focus on mapping the topology in Africa, which also happens to have the lowest internet penetration in the world [6]. Firstly, the African internet topology is characterised by heavily circuitous paths that go via other continents, which compromises internet user experience due to increased latency. The reason behind is lack of sufficient ISP peering amongst African local internet service providers, as previous similar research by Gupta, A. *et al* has shown [10]. Peering of ISPs depends on economic and political factors within each Autonomous System. Secondly, the data caches that host the popular sites like Facebook and Google are distant from most end users, sometimes even being in a different continent [6] [3]. Hence, such factors have to be considered in the data collection. Our research is important in revealing some of these unique characteristics in the topology so that network providers can make better decisions to improve their internet offering to customers.

In order to map internet topologies, researchers currently have to rely on a combination of known hacks and tricks to come up with a representation of the Internet's topology. These include the use of tools that make use of the *traceroute* utility (active techniques) and inferring from BGP tables (passive techniques) [3].

Employing the above techniques yields raw measurement data that simply gives the IP paths from each measuring source to target destination. Further processing needs to be done to this data in order to make sense of it. Issues such as alias resolution and anonymous router identification need to be employed in order to map IP addresses to routers, and further heuristics then need to be done in order to group routers according to their respective Autonomous Systems (ASes). The end goal is to represent the internet topology as a graph data structure, with nodes representing individual routers (or ASes depending on the resolution) and links as the logical connections between them. Annotations then need to be added to the links and nodes. This literature review goes into detail analysing the different processes and methods used to transform raw measurement data into a visualizable topology.

2 INTERNET MEASURING PROCESSES

Firstly, we will briefly look into the internet measurement process that enables us to get the data with which we represent the topology. Reproducing the topology of internet entities (routers, interfaces, or AS'es) is an interesting and non-trivial task that requires much thought, and hence much

Conference'17, July 2017, Washington, DC, USA
© 2020 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

research has been done on this subject. Researchers have always been trying to come up with techniques and tools that enable for different parts of the task to be done easier. Because there is no one centralized source that lists all the different routers in the world and their connections, such information needs to be inferred from other sources. Secondly, the internet landscape is rapidly changing and more and more routers and links are being added. This makes it hard for one to reproduce an accurate map at any given time without having to continuously update it frequently. From the literature research we have conducted, we have identified three main ways through which internet measurement data is obtained.

The first, less common one is through routing registry tables [5]. There exists Regional Internet Registries whose job is to allocate Autonomous System numbers and IP addresses to ASes, all of which are accessible through the WHOIS protocol. This is information such as which IP addresses belong to which AS-es. One advantage of using such registry databases as a source of information is that time is saved since one no longer has to manually probe the internet with *traceroute* themselves. However, the registries themselves rely on data submitted to them by the ISPs that own the routers and hence it probably will be incomplete; ISPs will decide which information to give out for privacy and policy reasons. Furthermore, this data is static and thus does not reflect the current state of the internet topology. This is a crucial factor because the topology is always changing.

The second method is by consulting BGP tables like RouteViews, PCH, and Hurricane Electric from which peering relationships between different AS-es can be inferred [10]. The Border Gateway Protocol (BGP) is used to share routing information between ASes. BGP routers connect to one another to exchange routing information [17]. When forwarding internet packets to a destination in another AS, a source AS will look up the BGP tables to see which AS-es it can forward the packets to. The BGP allows each AS to choose its own administrative policy in making inter-AS routing decisions. Hence one of the most important factors in determining routing policy decisions is the commercial contractual relationships between administrative domains [14]. To use BGP data to infer an AS level topology, researchers can consult these databases to check which AS-es each AS forwards packets to in the table entries and use this information to draw links between the AS-es. However, reproducing a reliable and complete AS level topology can only be done to a limited extent if only BGP data is relied upon. Unfortunately, BGP databases will not give the actual peering relationships between different AS-es, and hence the third method of measurement has to be employed [5]. Furthermore, Shavitt *et al* [18] mentioned in their paper that since this method relies on links published in BGP, as we gear more towards private peering between ASes, there is an increasing need to use active measurement techniques. Andersen *et al* [1] proposes a better method of using BGP tables that makes use of the BGP update messages that ASes send to each other regularly to infer an AS topology. They took advantage of the fact that ISPs aggregate prefixes into supemets, hiding many internal details. They then grouped IP address prefixes based upon how frequently they observed BGP update messages for both prefixes within the same time frame. Individual ASes would then be inferred by applying clustering algorithms on this data. The method proved to yield better accuracy than pure BGP tables.

The final method is to use the *traceroute* utility to continuously probe the internet from a source(s) to selected IP destination addresses with IP packets so as to collect a series of traces that show the routers along that path. *traceroute* is a utility that sends continuous IP packets of data to a chosen destination address, each successive packet with an incrementing *Time To Live (TTL)* value [11]. The aim is that along the path from the host to the destination, each of the routers in between receives a packet with *TTL=0* and sends back an ICMP error message that contains the address or name of the router. Hence, the output of running *traceroute* to one destination is a series of traces that show the different router addresses taken to reach the destination. It should be noted, however, that path routing of packets is asymmetric, meaning an ICMP message packet might not necessarily take the exact route that the outgoing *traceroute* packet took [3]. Hence, by running the *traceroute* utility on a list of destinations and using several heuristics and techniques to combine the data from the traces, a reasonable map of the topology will be obtained. Different researchers have developed different tools and methods to generate this list of destination IP addresses. Most use pre-existing databases from routing registries [7], and a few employ random address probing such as the Mercator tool developed by Govindan, R. *et al* in [8]. To perform the measurements, a source monitor is given these addresses and runs some sort of script to continuously send packets to multiple addresses in parallel and store the returned traces in a database. In their paper [15], Magoni, D. *et al* used *MySQLite* for storing their traces, and made use of a geolocation database to try and look up every IP address' geographic location and append the value to each trace entry. This enabled them to make their visualizations more meaningful. A few researchers have endeavoured to map the African internet topology using the above methods. Doing the same task in Africa poses a few extra challenges for researchers, and will reveal some interesting insights too. Gilmore J.S. *et al* [7] undertook a study to reproduce the African internet topology using the same techniques. They ran *traceroute* from a host in South Africa and they noticed that some routers do not respond to ICMP messages because they are configured to block or limit forwarding them. They also cited, as a future improvement, that future researchers should consider doing the *traceroute* measurements from multiple hosts that are spread evenly over the continent so as to have a more comprehensive map of the continent's topology [18]. Chavula *et al* also did a similar study of the African internet topology and their study revealed more insights [6]. They, however, performed their *traceroute* measurements from multiple vantage points all across Africa and had a more accurate representation of the topology.

Internet probing, however, produces a router level topology that needs to be processed further in order to convert it to an AS level map. IP addresses that belong to the same routers have to be identified and it is no trivial process; there exists tools like the CAIDA *iffinder* tool that help with this alias resolution, as it is called. Next, IP addresses have to be mapped to the AS-es they belong to [4, 12, 15].

3 TRANSFORMING RAW DATA INTO TOPOLOGIES

Once raw data has been collected from *traceroute*, it needs to be transformed into a graph data structure with nodes and links. This can be done at a router level and at an AS level, depending on the resolution desired. As highlighted earlier on, *traceroute* data helps mainly in reproducing a router level map. In order to do this visualization step, a few steps have to be taken as part of the process.

Alias resolution

The first important step in cleaning up trace data is to resolve aliases of IP addresses that belong to the same routers. A single router will have multiple interfaces through which it receives and sends out packets, and hence identifying which belong to the same router is important. To do this task, researchers have developed different heuristics and techniques.

Claffy *et al* in [4] developed their own internet measurement infrastructure called Archipelago (Ark) designed to offer researchers flexibility in taking coordinated measurements and collecting *traceroute* data. For alias resolution, Ark uses CAIDA's *iffinder* tool and the Analytical and Probe based Alias Resolver (APAR). The benefit then of using Ark is that it enables researchers to use these tools on the same monitor while concurrently collecting *traceroute* traces. This then produces the router level map.

Keys [13] gave a detailed description of the different classes of alias resolution techniques. They classify the techniques into two classes: fingerprint techniques and analytical techniques. Fingerprint techniques involve analysing trace results and looking for similarities that might indicate which responses came from the same routers. The benefit of this technique is that it is more accurate because the similarities are visible. However, since not all routers respond to probes, the method will only give limited alias resolution and hence affect the map's accuracy. Analytical techniques, on the other hand, work by using graph methods to analyse the IP address graph. However, since they are based on many assumptions, they are usually less accurate. Mercator [8] and CAIDA's *iffinder* [4] tool are both based on fingerprint techniques. [13] goes on to give details on examples of methods under each class. For instance, Ally and RadarGun are two methods that make use of the common IP ID counter. For analytical techniques, which involve analysing IP graphs, the two tools compared are Ark's APAR and *kapar*, where *kapar* is an improvement to the former one. They compared all the techniques on a large dataset of collected traces from a IP address list, and *kapar* combined with the *iffinder* tool gives the best alias resolution output. The two tools compliment each other in their weaknesses and strengths; *kapar* is an analytical technique tool hence it relies on no probes, but produces more false positives than *iffinder*, which, despite having to rely on ICMP messages returned by routers, has more accuracy.

Gilmore *et al* in [7] decided to add geolocation to their visualization so as to create a map that can be overlain on the true map. To achieve this, they decided to consult the GeoLite City database by MaxMind, a database that contains geographical locations of all the IP addresses that they had in their database. However, they did mention that the accuracy of this geo-location database is only known in the USA and not in Africa, hence more accurate database could have been used. Candela, M. *et al* also points out that geolocation data from such databases is often wrong or missing, and anycast addresses-addresses assigned to more than one physical device-cannot be mapped to a single location. For IP addresses that did not return geographical position when queried in the database, the researchers resorted to using other online sources like the *whois* register.

Govindan *et al* [8] decided to approach the alias resolution problem by using a fingerprint technique approach of sending a probe to a particular interface of a router and checking if the source address of the returned ICMP message is different from the IP address they sent to. If so, it is then concluded that the two interfaces are on the same router. However, they made some minor modifications to that approach to ensure they cater for some edge cases; instead of sending a single probe they sent multiple probes to the same interface and analysed the responses ' source interfaces. In addition, they also highlight that such probing might not resolve all aliases and they will have to turn to a more computationally intensive approach of using source-routing enabled routers that might be within their topology. [15] used the same fingerprint heuristics to resolve IP addresses that belonged to the same routers. They had a set of four criteria by which they determined if two IP addresses were aliases of each other. Apart from analysing response messages to probes, they also used DNS suffix matching and also took advantage of the fact that their *traceroute* measurements were done from multiple different *traceroute* vantage points. This helped them because aliases are easier to discover when traces are taken from multiple points. The difference between the two studies is that in Mercator, the alias resolution is done during probing, while in the latter study, it is only done once trace collection is complete.

After successful alias resolution, a router level map can be represented by a graph data structure in the form of adjacency lists.

Inferring an Autonomous System (AS) level map

In many studies, researchers are more interested in generating an AS level map of the internet topology. Autonomous Systems are basically clusters of routers all belonging to one central organization within that location. These AS-es may be internet service providers, universities, and other large corporate networks [14] which have a set of IP prefixes allocated to them. In our study, AS level maps would reveal more insights and would be more helpful in understanding how internet traffic in Africa flows and what can be done to improve it. Some research studies like [4] attempted to create a dual AS router topology that shows both the AS and router level maps on one. The links connecting the AS-es would be annotated with the router names that connect them while the individual routers would be annotated with the AS numbers to which they belong. The two maps have to be generated from two completely different techniques hence proper relationships between them cannot be directly inferred.

In their study of the African internet topology, Gupta *et al* [10] wanted to analyse certain unique characteristics about the internet topology that required visualization at an AS level. They wanted to find out what was the reason behind the high levels of internet latency in Africa. To construct their AS level map, they ran *traceroute* from BISmark routers in South Africa in order to discover the router level map as well as BGP information from sources like RouteViews, Packet Clearing House (PCH) to infer relationships from the BGP databases. BGP tables show which ASes packets can be forwarded to given a source AS. Information from these databases alone will not give a comprehensive view of the AS topology, hence the use of *traceroute* to complete the map. In their paper [3], Chavula *et al* used the same topology inferring approach of combining *traceroute* measurements with data from BGP tables; they were also interested mainly in the AS level topology of the internet. However, they improved the quality of the measurement process by running *traceroute* from different vantage points that are located in different countries in Africa. They also combined the techniques followed in [7] of using a geolocation database to look up the geographical locations of IP addresses in their target addresses. [8] followed the same procedure as well to infer the AS level topology but with a few further changes to improve the accuracy of the map. After producing their router level map with *traceroute* techniques, they then used the *traceroute* data together with information from BGP tables to produce an IP to AS map. After discarding distorted and ambiguous links (about 5% of all collected traces), the map produced is a simple undirected and unweighted graph which still needs annotations indicating business relationships between AS-es using techniques developed at CAIDA based on multiobjective optimization. This is done to add meaning to the map.

Possible sources of error that may affect the map's accuracy

This section focuses on some of the unique challenges faced when trying to convert raw internet measurement data into a meaningful topology of the internet. These are:

(1) Dealing with anonymous routers in *traceroute* measurements
(2) Dealing with wrongly named DNS routers from DNS databases.

Gunes *et al* [9] looks at how *traceroute* traces sometimes return routers that do not respond to *traceroute* commands and are marked by an asterisk ('*') on the trace, and the impact they can have on the accuracy of the map. Reasons as to why some routers are configured not to share their information mostly include security and policy. Some routers are simply configured to either limit or ignore ICMP messages passing through them. They propose a graph-based induction technique to resolve such anonymous routers. It works by them creating multiple topology maps with the anonymous routers in them and then visually analysing them in order to come up with algorithms that can detect them. They tested their approach against the previously used neighbor matching approach. However, they did not test it again other methods such as dimensionality reduction and graph minimization approach [19].

Zhang *et al* [20] deals with the problem of wrongly named routers. This research study looks into how researchers, while looking up the names of routers from geolocation databases, often come across incorrectly named routers, which would consequently reduce the accuracy of the inferred topology. Although misnamed routers constitute a very small percentage (about 0.5%), they can lead to a topology map with over 10% of its links being false. One of the techniques to resolve this is to run *traceroute* measurements from multiple vantage points.

4 DISCUSSION

The literature reviewed above gave us a robust understanding of the previous works done in this research topic and the answered and unanswered questions. We can see, firstly, that mapping the internet topology is the first step in many internet network studies hence much work has been done on it. It is also clear all the papers reviewed agree on one general approach of mapping the internet topology. The two go to methods for collecting data on and inferring the topology are traceroute measurements for router discovery and BGP databases for linking AS-es that connect as well as mapping routers to AS-es. Therefore it is imperative that using a blend of both traceroute and BGP information to do the topology mapping is the best way to go about it. Most of the further efforts researchers then do is to devise tools and techniques that simply make these two existing processes more efficient. We have seen this in tools like Mercator, Ark, and nec. The researchers in these papers have identified different aspects on which to improve on the current tools and methods, each improvement with its own strengths and weaknesses. Mercator, for instance, uses a heuristic that does not rely on input IP addresses when running its traceroute, which makes it an easier to use tool. On the other hand, Ark has developed the iffinder and kapar tools that are excellent at alias resolution and have been used by other researchers too.

From the papers analysed above too, taking traceroute measurements from multiple vantage points is a recurring recommendation in order to yield more accurate maps. This can be achieved by having them physically or relying on source routing enabled routers in the topology. Even though these only constitute 8% of all routers, [7] shows that such a number is enough to capture 90% of links that would not have been discovered by one traceroute. Running from multiple vantage points is also necessary because we have seen that not all routers in the topology respond to probe packets and hence they can only be discovered when they appear in multiple paths from different sources. Shavitt *et al* [18] did an extensive and detailed study on quantifying the benefits of vantage point distribution when taking measurements for map visualization, and they revealed how measuring from a set of well spread vantage points that perform measurements for a prolonged period of time will result in good network coverage even outside the network being studied. The quality of the topology map realized is also heavily reliant on other aspects such as the accuracy of the BGP database tables consulted, presence of anonymous or misnamed routers, and also the duration within which measurements are taken since the topology is never static.

We have also seen how alias resolution is a very important step of the topology visualization process and has to be made as accurate as possible. Any inaccurate resolution will lead to false links which will distort the topology. We have seen how one can make use of fingerprint techniques or analytical techniques to achieve this, and how fingerprint techniques tend to be more common and accurate.

Lastly, we have seen how internet topology studies in Africa have revealed interesting insights about the African topology. They have shown how there is lack of sufficient ISP peering among the major players, leading to packets often being routed via intercontinental routes before reaching their local destinations. This affects the internet experience of end users as there is now more latency introduced into the routing of packets. Furthermore, it makes it harder for companies which wish to install local data caches into the continent to do so since they might not be fully utilized by the local users. There is still more study to be done on the African internet topology in order to reveal more of these unique characteristics.

5 CONCLUSIONS

The literature we have reviewed has given us a clear roadmap of how we can proceed with our research study in terms of which tools and techniques are the best. Since our research focuses on visualizing the African topology, we are focusing more on the AS level visualizing processes. So we will heavily look into BGP databases to infer the AS level topology, and then also make use of traceroute data since BGP information is not enough to generate a complete topology. Our traceroute measurements will be performed from multiple vantage points spread across the continent as most of the papers have recommended to be done in future studies.

For alias resolution, we will make use of fingerprint techniques particularly using tools such as iffinder and kapar since these have proven to offer the best results. To add more meaning to the topology map, we will add geolocation of routers and hence the AS-es themselves by using appropriate geolocation databases. Also, research by Candela, M. *et al* on their study of geo-locating IP infrastructure using RIPE Atlas, they discovered that geo-location is more accurate when locating routers than edge infrastructure [2]. In their paper [16], however, Motamedi *et al* pointed out that geolocation for an AS level map is a bit more subtle and different from a router or interface level map, mainly because AS'es usually cover whole geographic regions which might overlap, and then factoring in the presence of Internet Exchange Points makes the task more complicated. Hence they recommend that an AS level map with geography added to it should be viewed as a hyper-graph as shown below. Vertical lines show connectivity between the respective AS'es.

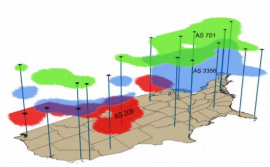

Figure 1: A sample image showing how AS'es may be represented with geolocation

REFERENCES

[1] Andersen, G. David, N. Feamster, S. Bauer, and H. Balakrishnan. 2002. Topology Inference from BGP Routing Dynamics. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurment* (Marseille, France) *(IMW '02)*. Association for Computing Machinery, New York, NY, USA, 243–248. https://doi.org/10.1145/637201.637239

[2] M. CANDELA, E. GREGORI, V. VALERIO LUCONI, and A. VECCHIO. 2019. Using RIPE Atlas for Geolocating IP Infrastructure. (English). 7 (2019), 48816–48830. https://doi.org/10.1109/ACCESS.2019.2909691

[3] J. Chavula, N. Feamster, A. Bagula, and H. Suleman. 2015. Quantifying the Effects of Circuitous Routes on the Latency of Intra-Africa Internet Traffic: A Study of Research and Education Networks. 147 (2015), 64–73. https://doi.org/10.1007/978-3-319-16886-9_7

[4] K. Claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov. 2009. Internet Mapping: From Art to Science. In *2009 Cybersecurity Applications Technology Conference for Homeland Security*. 205–211. https://doi.org/10.1109/CATCH.2009.38

[5] B. Donnet and T. Friedman. 2007. Internet topology discovery: a survey. *IEEE Communications Surveys Tutorials* 9, 4 (2007), 56–69. https://doi.org/10.1109/COMST.2007.4444750

[6] A. Formoso, J. Chavula, A. Phokeer, A. Sathiaseelan, and G. Tyson. 2018. Deep Diving into Africa's Inter-Country Latencies. (English). (2018). https://doi.org/10.1109/INFOCOM.2018.8486024

[7] J.S. Gilmore, N.F. Huysamen, P. Cronje, M.C. de Klerk, and A.E Krzesinski. [n.d.]. Mapping the African Internet. (English) [On recreating and mapping the African network]. ([n. d.]). https://doi.org/10.1.1.584.1292

[8] R. Govindan and H. Tangmunarunkit. 2000. Heuristics for Internet map discovery. In *Proceedings IEEE INFOCOM 2000. Conference on Computer Communications. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies (Cat. No.00CH37064)*, Vol. 3. 1371–1380 vol.3. https://doi.org/10.1109/INFCOM.2000.832534

[9] M. H. Gunes and K. Sarac. 2008. Resolving Anonymous Routers in Internet Topology Measurement Studies. In *IEEE INFOCOM 2008 - The 27th Conference on Computer Communications*. 1076–1084. https://doi.org/10.1109/INFOCOM.2008.162

[10] A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Calandro, and E. Katz-Bassett. 2014. Peering at the Internet's Frontier: A First Look at ISP Interconnectivity in Africa. 8362 (2014). https://doi.org/10.1007/978-3-319-04918-2_20

[11] B. Huffaker, D. Plummer, D. Moore, and K. Claffy. 2002. Topology discovery by active probing. In *Proceedings 2002 Symposium on Applications and the Internet (SAINT) Workshops*. 90–96. https://doi.org/10.1109/SAINTW.2002.994558

[12] C. Hyunseok, J. Sugih, and W. Walter. 2001. Inferring AS-level Internet Topology from Router-Level Path Traces. *Scalability and Traffic Control in IP Networks* (2001). https://doi.org/10.1117/12.434395

[13] Keys K. 2010. Internet-Scale IP Alias Resolution Techniques. (English) [On Internet topology mapping alias resolution techniques]. *ACM SIGCOMM Computer Communication Review* 40, 1 (2010), 50–55. https://doi.org/10.1145/1672308.1672318

[14] Lixin Gao. 2001. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking* 9, 6 (2001), 733–745. https://doi.org/10.1109/90.974527

[15] D. Magoni and M. Hoerdt. 2005. Internet core topology mapping and analysis. *Computer Communications* 28, 5 (2005), 494–506. https://doi.org/10.1016/j.comcom.2004.09.002

[16] R. Motamedi, R. Rejaie, and W. Willinger. 2015. A Survey of Techniques for Internet Topology Discovery. *IEEE Communications Surveys Tutorials* 17, 2 (2015), 1044–1065. https://doi.org/10.1109/COMST.2014.2376520

[17] M.O. Nicholes and B. Mukherjee. 2009. A survey of security techniques for the border gateway protocol (BGP). *IEEE Communications Surveys Tutorials* 11, 1 (2009), 52–65. https://doi.org/10.1109/SURV.2009.090105

[18] Y. Shavitt and U. Weinsberg. 2011. Quantifying the Importance of Vantage Point Distribution in Internet Topology Mapping (Extended Version). *IEEE Journal on Selected Areas in Communications* 29, 9 (2011), 1837–1847. https://doi.org/10.1109/JSAC.2011.111008

[19] B. Yao, Ramesh Viswanathan, F. Chang, and D. Waddington. 2003. Topology inference in the presence of anonymous routers. In *IEEE INFOCOM 2003. Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies (IEEE Cat. No.03CH37428)*, Vol. 1. 353–363 vol.1. https://doi.org/10.1109/INFCOM.2003.1208687

[20] M. Zhang, Y. Ruan, V. Pai, and J. Rexford. 2006. How DNS Misnaming Distorts Internet Topology Mapping. (2006), 369–374.