Analysis of Autonomous System Level Internet Topology Graphs and Multigraphs

Abdullah Yasin Nur
Department of Computer Science
University of New Orleans
New Orleans, LA 70148
Email: ayn@cs.uno.edu

***Abstract*—The Internet is a network of networks consisting of thousands of Autonomous Systems (ASes). ASes in the Internet are organizations that connect with each other in various forms to make global Internet communication possible. The Internet is one of the biggest human-engineered systems, and understanding the topology can help network engineers and researchers. In this work, we analyze the structure of the Internet by comparing AS level graphs and AS level multigraphs. We use several real-life datasets, including traceroute, DNS, BGP, commercial geolocation databases, to construct Cross-AS (X-AS) level topology maps. Then we retrieve AS level graphs and multigraphs from X-AS level topology maps. Comparisons between the multigraphs and graphs allow us to study the impact of parallel connections on clustering and ranking ASes in the Internet.**

Index terms— Autonomous Systems, Internet topology, AS, Cross-AS, X-AS

I. Introduction

The Internet is one of the largest human-engineered systems, and the recent advances in research elevate the Internet as the primary communication tool. The usage of the Internet is increasing every day, and by the time this paper is written, it is reported that more than 4.8 billion users are using the Internet, which is more than half of the world population [26]. The Internet is defined as a network of networks formed by thousands of Autonomous Systems (AS). An AS in the Internet is a group of networks administered by one or more network operators. Each AS is uniquely identified by an Autonomous System Number (ASN), which is given by the Internet Assigned Numbers Authority (IANA). ASes connect with each other in various forms and make global Internet communication possible.

ASes usually build their physical network in a hierarchical manner. They divide their network into sub-networks, which is called Point of Presence (PoP). A PoP contains multiple routers and network devices located in the same facility. PoPs create AS'es backbone, and the connection between two ASes is handling by their PoPs. A link between two PoPs indicates a physical connection.

Internet topology mapping is significantly helpful for network researchers, network engineers, and operators [2]. With an accurate map, researchers can develop more efficient protocols or services and evaluate their performances under better knowledge. Also, Internet Service Providers (ISP), Internet Content Providers, and other ASes can use the map to increase their backbone efficiency [1]. From the security perspective, an accurate map can help to detect possible attack points and assess vulnerabilities [8]. Additionally, it can be used to mitigate the effects of cyber-attacks [14].

Discovering the topology of the Internet is a challenging problem because of lack of measurement tools, lack of available sources, not enough support from ASes, and the current infrastructure of the Internet. However, the importance of the map is attracted to researchers, and they have extensively studied about creating an accurate topology map [2, 3]. Although their design purposes are different, Ping, Traceroute, and Border Gateway Protocol (BGP) are mainly used tools for discovering topology maps. Ping is designed to help network admins to check the IP address is alive or not. Traceroute is the tool that collects a list of IP addresses on the path between two hosts. BGP is a gateway protocol that has been using for exchanging reachability information of the ASes.

Internet topology mapping is generally described in four levels; interface, router, PoP, and AS level topology maps. A router is a network device that forwards packets toward their destinations in packet switching networks. Routers are connected to multiple networks through various interfaces which are assigned unique IP addresses. Interface level topology maps discover the connectivity between IP interfaces. Although these types of maps are easy to construct, they do not have much use. Router level maps group IP interfaces into corresponding routers and infer connections between routers. Each node presents a router in the map, and each link presents a connection between routers. The biggest challenge of these maps is discovering the IP addresses assigned to the same router, which is called the IP alias problem [7]. Even though researchers have suggested several solutions [5, 6], it is still an open problem due to the limited network support, rate limiting practices, and scalability issues [4].

PoP level Internet topology maps cluster the interfaces or routers located in the same facility for each ASes. The main focuses of the maps are the physical locations of the facilities and the connections between these facilities. The techniques for discovering the maps typically employ path traces to find the IP interfaces and use several geolocation techniques to assign those interfaces to their geographical locations. These types of maps provide the physical infrastructure information of the backbone networks instead of a simple abstraction. Therefore, it

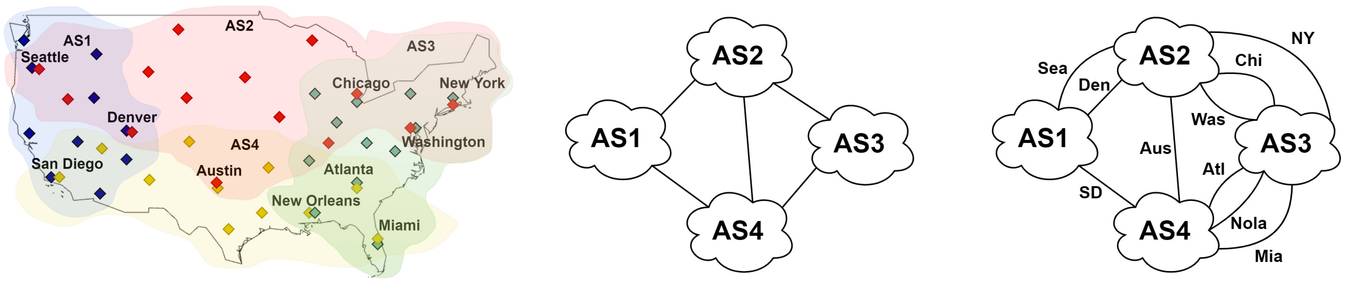

(a) Detailed Topology (b) AS Topology Graphs (c) AS Topology MultiGraphs

Fig. 1: An Example Topology Map of the Internet in AS Level and AS Multilevel

provides precious information to analyze the Internet structure. However, capturing the backbone connections within an AS is a challenging task because (i) IP protocol does not provide geolocation information; (ii) it requires carefully crafted traceroute queries per AS to increase coverage while reducing the probing overhead [10]; and (iii) probing the backbone topology of ISPs is more prone to packet filtering [9].

Finally, AS level Internet topology maps abstract the topology by using logical links where nodes present ASes and links present the business relations between ASes as customer-to-provider (c2p), peer-to-peer (p2p), and sibling-to-sibling (s2s) [12]. Abstracting ASes without any internal structure is an oversimplification since the ASes in the Internet span over various geographic regions and often cover the same regions in part or whole [11]. Therefore, it is inadequate to use AS level topology maps to analyze the routing, resilience, and robustness of the Internet [1]. On the other hand, AS level maps are suitable for BGP path analysis, analyzing Internet economics, and relatively easy to construct an accurate map.

It is clear that there is a significant gap between PoP level and AS level maps. AS level maps are easy to generate but abstract all critical information about ASes. On the other hand, PoP level maps are suitable to analyze the resilience, robustness, and efficiency of the Internet, but it is hard to generate because of ISP's security and privacy concerns. Mapping cross-AS connections is particularly important because it allows us to detect congestion points, analyze the resilience and robustness of the Internet, mitigate the impact of Denial of Service attacks, and optimize server deployment in content delivery networks [1]. In our previous work, we introduced Cross-AS (X-AS) Internet Topology Maps to address this problem [1]. X-AS maps capture both ASes and cross-connections between ASes, which allow us to go beyond the simple AS level graphs and abstract the topology of the Internet as a multigraph supporting multiple connections among the ASes. X-AS level topology maps use a set of techniques that exploit multiple data sources including, traceroute data, BGP advertisements, geolocation databases, and DNS datasets.

To illustrate, Figure 1 shows four ISPs, AS1, AS2, AS3, and AS4, providing Internet access service in the US. Figure 1b presents AS level abstract of the topology presented in Figure 1a. AS level Internet topology graph, $G = (V, E)$, abstract the topology by using logical links where the nodes present ASes and the links present the business relations between ASes. This abstraction miss the critical connection information between ASes because of the oversimplification. As an example, AS2 and AS3 have connections in New York, Chicago, and Washington DC. However, AS level graphs show only one logical connection between AS2 and AS3. A proper abstraction of the Internet topology would be a multigraph $G = (V, E, f)$ where the vertex set, V, corresponds to the ASes, the edge multiset, E, represents the cross-connections between the ASes and $f : E \rightarrow \{(v_i, v_j) : v_i, v_j \in V, v_i \neq v_j\}$ is a function returning the endpoints of the edges to support parallel edges between two ASes [1]. Figure 1c shows the multigraph version of the example topology, which is derived from X-AS topology maps.

In this work, we analyze the structure of the Internet by comparing AS level graphs and AS level multigraphs. We use several real-life datasets including traceroute, DNS, BGP, commercial geolocation databases, to construct X-AS level topology maps [1]. Then we retrieve AS level graphs and multigraphs from X-AS level topology maps. Comparisons between the multigraphs and graphs allow us to study the impact of parallel connections on AS clustering in the Internet.

The rest of the paper is organized as follows. In Section II, we present the related work. We introduce the details of our approach in Section III. Section IV demonstrates our experimental results and comparisons. Finally, we conclude the paper in Section V.

II. Related Works

Many tools and approaches have been suggested to derive complete and accurate Internet topology maps at router, PoP, and AS levels and these methods have been discussed at [2, 3]. Researchers have used traceroute-like tools to create router-level topology maps [13, 15]. However, routers may contain more than one IP address, and finding IP addresses in the same router is a challenging problem [7]. Mercator [13] is a probe-based IP alias resolver that depends on the similarity in IP addresses of the returned probe responses. Ally [10] extends Mercator by utilizing the IP Identification field values of the returning response packets to decide on aliases. Radargun [16] employs a velocity modeling scheme to reduce Ally's quadratic probing complexity. APAR [5] is an inference-based alias

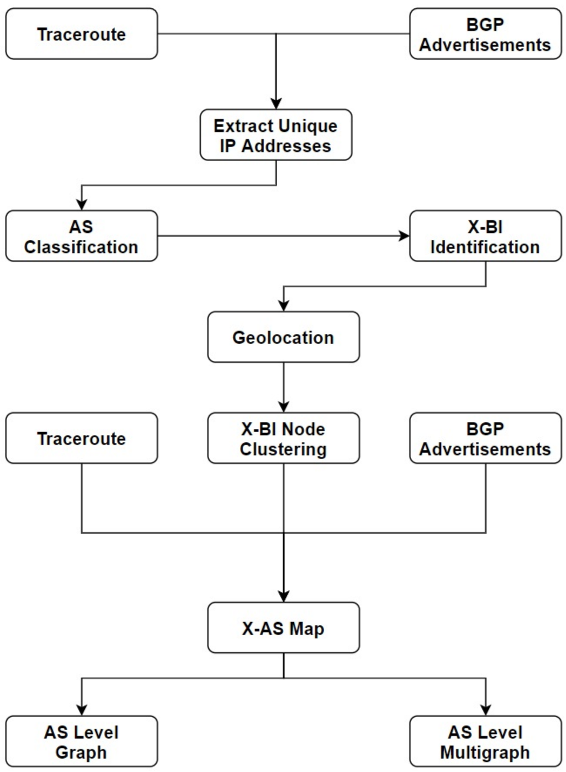

Fig. 2: X-AS Topology Map Generation Flow

resolver to resolve aliases among IP addresses collected via traceroute during a topology mapping study.

Evaluating each router on the Internet is a complicated job, and it produces a high error rate. Instead of mapping IP addresses to routers, clustering IP addresses into their geolocation and discovering connections between these facilities is called Point-of-Presence(PoP) level mapping. On the other hand, PoP level maps suffer from IP geolocation because IP protocol does not provide any geolocation information. Researchers suggest delay-based geolocation [17, 18] and topology-based geolocation [19]. Also, several commercial geolocation databases in the market use a couple of methods to increase their accuracy and sell their databases.

Several PoP level mapping techniques have been suggested in the literature. In addition to the router level topology, the Rocketfuel project aims to discover PoP level maps of the Internet [10]. They use DNS information to infer locations of IP addresses. Madhyastha et al. [20] use multiple vantage points from PlanetLab nodes and freely available Looking Glass servers to collect traceroute results from the Internet. They improved Rocketfuel's DNS project and applied it in their work. On the other hand, DNS naming is volunteer-based, and not all ASes provide IP address geolocation information with DNS. Although large ISPs generally use specific DNS naming conventions that give information about locations of IP addresses, less support from ISPs makes it impossible to use alone for the overall Internet. Feldman et al. [21] use common network motifs in IP interconnections to identify PoPs and use several commercial geolocation services to discover the location of PoPs.

TABLE I: Datasets used in this work

Dataset	Source	Dataset name	Reference
Traceroute	Caida	The IPv4 Routed /24 Topology Dataset	[29]
BGP	RIPE NCC	Routing Information Service (RIS)	[33]
BGP	RouteViews	The Route Views Project	[34]
IP2AS	Caida	IPv4 Routeviews Prefix to AS Mappings	[28]
DNS	Rocketfuel	UNDNS	[10]
Geolocation Database	DB-IP	IP Address to Location	[30]
Geolocation Database	Maxmind	GeoLite2 City	[31]
Geolocation Database	IP2Location	DB5 Lite	[32]

Generating AS level topology maps is relatively simple compared to router level and PoP level maps. Chang et al. [22] discover AS level topology by inferring individual connections from the router level path traces of the Internet. Mao et al. [23] improve the accuracy of IP-to-AS mappings by using BGP tables and traceroute paths collected from multiple vantage points. Mahadevan et al. [24] combine traceroute, BGP, and WHOIS measurements to create more accurate AS level maps. Gao [12] classified business relations between ASes into three groups, customer-to-provider (c2p), peer-to-peer (p2p), and sibling-to-sibling (s2s), based on the assumption that AS level paths are valley-free. Giotsas et al. [25] suggested a new algorithm to infer hybrid relations between ASes, where two ASes have different relationships at various locations.

Abstracting ASes without any internal structure is an over-simplification since the ASes in the Internet span over various geographic regions and often cover the same regions in part or whole [11]. Moreover, they physically connect at multiple colocation centers or Internet eXchange Points (IXPs) to exchange traffic and routing information. In one of our previous works [1], we introduced Cross-AS (X-AS) topology maps that capture both ASes and the parallel cross-AS connections observed at the network layer in the Internet. X-AS Internet topology maps allow us to abstract the Internet's AS level topology as a multigraph rather than a simplified graph.

III. Methodology

In this section, we present our methodology to discover and analyze AS level graphs and multigraphs.

Cross-AS (X-AS) level topology maps [1] use a set of techniques that exploit multiple data sources presented in Table I. Figure 2 presents the flow chart of our topology generation process. We use traceroute datasets, BGP advertisements, and IP address to AS mapping tools to extract IP addresses that appear in path traces where the paths switch from one AS to another. We identify those IP addresses as cross-border interfaces (X-BIs). Then, we apply a set of techniques based on DNS names, geolocation databases, and topological geolocation to accurately cluster X-BIs into their geolocations, which are

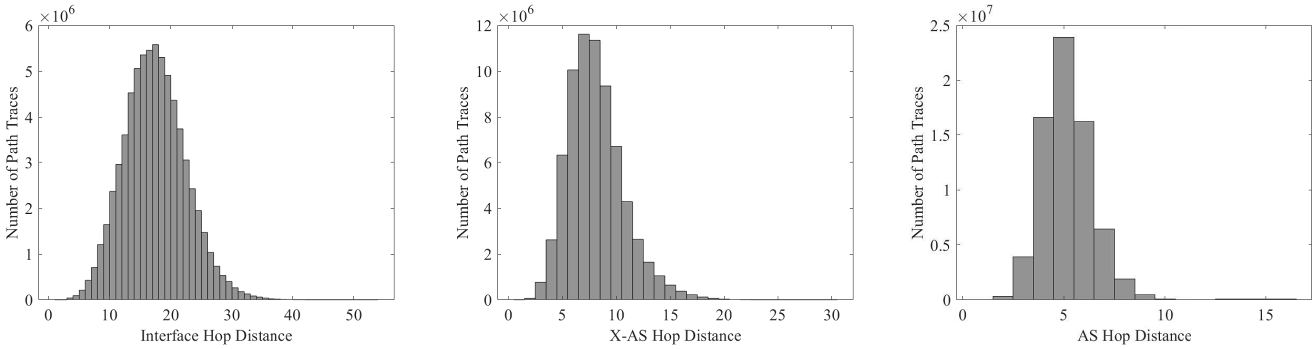

(a) Interface Level Hop Distance Distribution (b) X-AS Level Hop Distance Distribution (c) AS Level Hop Distance Distribution

Fig. 3: Hop distance distribution of the Internet

TABLE II: Summary statistics for Degree Distributions

$Level$	Q_0	Q_1	Q_2	Q_3	Q_4	$Mean$	$StdDev$
Interface	1	13	17	20	54	16.94	5.03
X-AS	1	6	8	10	30	8.10	2.58
AS	1	4	5	6	16	5.17	1.21

called X-BI nodes. Lastly, we exploit traceroute and BGP datasets to discover the cross-connections between X-BI nodes. The final X-AS map, $X = (N, C)$, consists of a set of X-BI nodes, N, and a multiset of X-BI connections, C. We define AS level graph G = (V; E) where the vertex set, V, corresponds to the ASes and the edge set, E, represents the logical relations between the ASes. Also, we define AS level multigraph $G = (V, E, f)$ where the vertex set, V, corresponds to the ASes, the edge multiset, E, represents the cross connections between the ASes and $f : E \to \{(v_i, v_j) : v_i, v_j \in V, v_i \neq v_j\}$ is a function returning the endpoints of the edges to support parallel edges between two ASes. For more details about X-AS topology maps, please check the original paper [1].

IV. Experimental Results and Comparisons

In this section, we present our comparison results of AS level graphs and multigraphs.

A. Dataset Analysis

We used the CAIDA IPv4 Prefix-Probing Traceroute Dataset [29] consisting of more than 69 million (69,942,041) path traces. The minimum and maximum Interface level hop lengths in our dataset are 1 and 54, respectively. The average hop length is 16.94, and the hop length distribution is symmetric-like as shown in Figure 3a. When we further analyze the hop distances, we observe that only 3,021 path traces out of 69,942,041 have more than 40 hops. Additionally, only 171 path traces have more than 45 hops, and only 11 path traces have 50 hop distance between two end hosts. Figure 3b presents the X-AS level hop distance between a source and a destination in the Internet. The minimum and maximum X-AS level hop lengths in our dataset are 1 and 30, respectively. The average X-AS level hop length is 8.10. We observe that only 97 path traces have more than 25 X-AS between two end hosts. Figure 3c presents AS level hop distance between a source and a destination in the Internet. The minimum and maximum AS level hop lengths in our dataset are 1 and 16, respectively. The average AS level hop length is 5.17. We observe that only 13 path traces have more than 13 AS, and only 2 path traces have more than 14 AS between two end hosts. Finally, Table II shows the minimum, first quartile, second quartile (median), third quartile, maximum, mean, and standard deviation for each level.

B. X-BI Node Analysis

Figure 4 shows X-BI node count by countries in the world as a heatmap. The United States leads the league with 14,443 nodes. The nearest competitor is Brazil with 4,688 nodes. The main reason behind the observation is the US contains more ASes than other countries. Also, more ISPs give service in the US. Therefore, the number of interconnection between ASes in the US is quite larger than in other countries. Interestingly, we observe that China has 678 X-BI nodes and ranked 19 overall, lower than many smaller countries. When we analyze the continents, we observe 17,087 X-BI nodes in North America, 19,814 in Europe, and 9,934 in Asia. The lowest count is Africa with 1,344 nodes.

Figure 5 shows X-BI node count by States as a heatmap. California is leading with 1,696 X-BI node in California, whereas Texas is in second place with 1,144. The minimum numbers of X-BI nodes are 33 in Rhode Island and 45 in Hawaii. We observe 283.19 nodes on the average in all states. Some of the states contain more nodes than some of the entire countries in the world. For example, California would be ranked 7 in the world if it was a country. Additionally, the average number 283.19 would be ranked in 32.

Figure 6 shows X-BI node count by countries in Europe as a heatmap. Russia has the most X-BI nodes in Europe with 3,676 nodes. The average number of nodes per country is 412.46. Interestingly, Poland ranked 4 with 1,636 nodes, passing Italy, Spain, and France on the list.

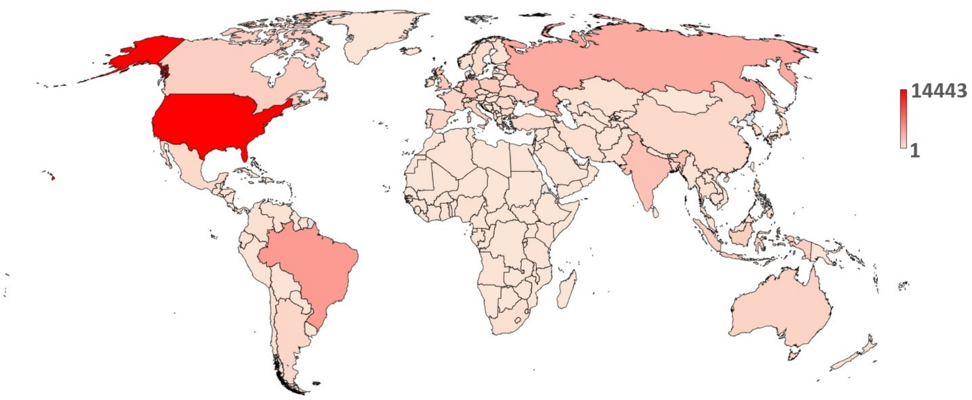

Top 10 Countries

Country	X-AS Node Count
United States	14443
Brazil	4688
Russia	3676
India	2229
Germany	1943
United Kingdom	1710
Poland	1636
Indonesia	1306
Canada	1306
Ukraine	1092

Fig. 4: X-BI Node Count by Country in the World

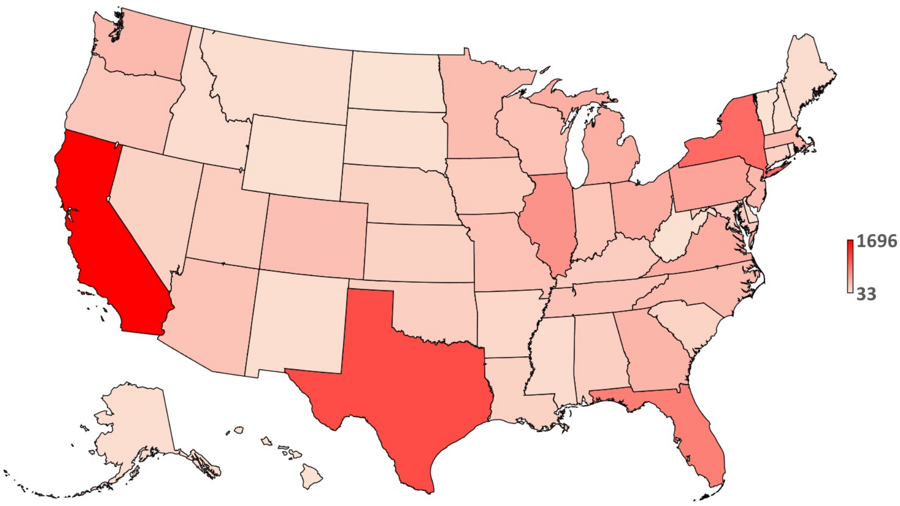

Top 10 States

States	X-AS Node Count
California	1696
Texas	1144
New York	922
Florida	769
Illinois	627
Pennsylvania	505
Ohio	440
Virginia	435
Michigan	427
Georgia	380

Fig. 5: X-BI Node Count by States

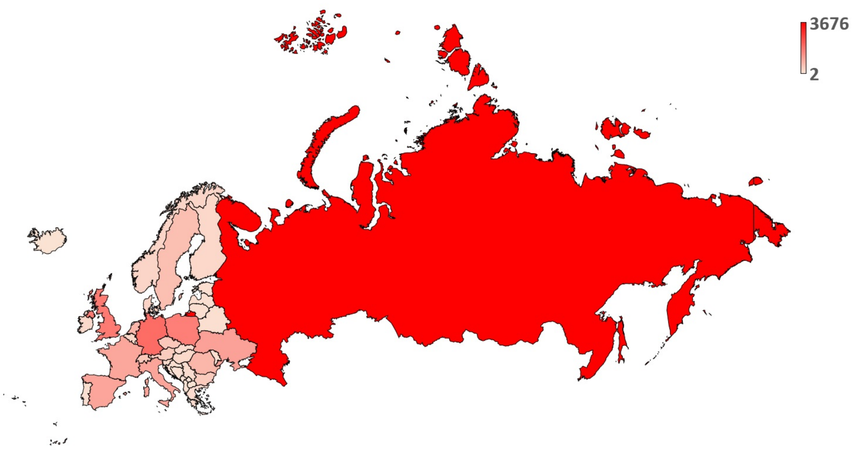

Top 10 Countries in Europe

Country	X-AS Node Count
Russia	3676
Germany	1943
United Kingdom	1710
Poland	1636
Ukraine	1092
Italy	1013
Spain	1008
France	1001
Netherlands	751
Romania	733

Fig. 6: X-BI Node Count by Country in Europe

TABLE III: Top-15 ASes Detailed Comparison (AS Multigraph Degree Sorted)

ASN	Company	Country	AS Multigraph Degree	AS Graph Degree	Average Number of IP Address in X-BI Nodes	Provider	Peer	Customer
1299	TeliaNet	Sweden	22528	6793	144.12	0	46	2063
174	Cogent	US	20849	7761	55.71	0	105	6172
3356	Level3	US	17128	6515	70.01	0	56	5783
6939	Hurricane	US	13165	5493	30.39	1	7064	1802
6461	Zayo	US	8524	4169	47.62	0	276	1832
3257	GTT	US	8410	3038	48.45	0	42	1965
2914	NTT	US	7633	2637	105.91	0	71	1625
6453	Tata	US	5491	1566	32.67	0	56	642
3549	Level3	US	4327	2085	24.08	2	199	5532
6762	Telecom Italia	Italy	3764	1435	76.48	0	65	523
7018	ATT	US	3391	1792	49.67	0	43	2497
3491	PCCW	US	2961	846	43.39	0	147	575
9498	Bharti Airtel	India	2875	1123	35.36	10	228	910
12389	Rostelecom	Russia	2830	1309	19.51	11	177	1015
9002	RETN	UK	2777	1786	81.25	2	540	946

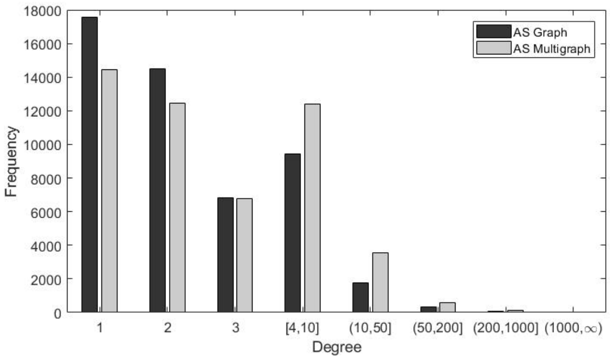

Fig. 7: AS Level Graph and Multigraph Degree Distribution

C. Degree Distribution Analysis

We define the degree of an AS as the number of connections it has to other ASes. Figure 7 presents the degree distribution of AS level graphs and multigraphs. It is clear that the majority of the ASes are virtually at the edge of the Internet without providing any internet access to other ASes. These ASes are called stub-ASes in the traditional tier classification. 95.7% of the ASes have a degree less than 10 in AS level whereas 91.4% in AS multilevel. Only 16 ASes have more than 1,000 connections with other ASes in AS level, where the maximum degree is 7,761 (ASN174 - Cogent). On the other hand, 29 ASes have more than 1,000 connections with other ASes in AS multilevel, where the maximum degree is 22,528 (ASN1299 - TeliaNet).

D. Detailed Analysis of Top-15 ISPs

In the traditional tier classification of ASes, tier-1 ASes are defined as provider-free ASes that peer with every other provider-free AS to ensure reachability to all destinations without purchasing IP transit. Even though we have less than 20 tier-1 ISPs in the Internet, they carry the majority of the traffic. To understand the main core of the Internet, we analyze the top-15 ASes with respect to their AS level multigraph degree. Since we use multigraph, we also consider parallel links between ASes. To provide more detail, we use the number of providers, peers, and customers of each ASes [27].

In our analysis, we observe a couple of interesting outcomes for some top-level ISPs. Our first observation is TeliaNet (ASN1299). TeliaNet has 0 providers, 46 peers, and 2063 customers. Even though some of the top ISPs have more peers and customers than TeliaNet, they ranked number one with respect to AS Multigraph degree. TeliaNet makes several connections in various locations with many ISPs comparing to other ISPs to increase redundancy. The additional connections improve the resilience and robustness of their network.

Next, Hurricane Electric (ASN6939) is ranked as number four. However, they have one provider (c2p relation with TeliaNet - ASN1299), which makes them tier-2 AS in the traditional tier system. They have the highest peering count with 7064 in the Internet. Additionally, India-based ISP Bharti Airtel (ASN9498) is ranked 13, and Russia-based ISP Rostelecom (ASN12389) is ranked 14. Even though they have 10 and 11 providers, they ranked higher than several tier-1 ISPs. Therefore, these observations prove that ranking an AS via the traditional tier system is neither adequate nor correct.

V. Conclusions

Technology era makes the Internet as the main communication tool. Understanding the structure of the Internet is significantly important to improve the quality of service of the Internet. In this work, we analyze the structure of the Internet by comparing Autonomous System (AS) level graphs and AS level multigraphs. We use several real-life datasets including traceroute, DNS, BGP, commercial geolocation databases, to construct Cross-AS (X-AS) level topology maps. Then we retrieve AS level graphs and multigraphs from X-AS level topology maps. Comparisons between the multigraphs and graphs allow us to study the impact of parallel connections on AS clustering in the Internet.

Acknowledgements

We are grateful to DB-IP [30] for providing their commercial geolocation database. We are also thankful to CAIDA, Maxmind and IP2Location projects for publicly sharing their datasets.

References

[1] A. Y. Nur and M. E. Tozal, ”Cross-AS (X-AS) Internet Topology Mapping”, Computer Networks, Volume 132, pp. 53-67, 2018
[2] R. Motamedi, R. Rejaie, and W. Willinger, ”A Survey of Techniques for Internet Topology Discovery”, IEEE Communications Surveys and Tutorials, 2014.
[3] W. Willinger and M. Roughan, ”Internet Topology Research Redux”, ACM SIGCOMM eBook: Recent Advances in Networking, 2013.
[4] K. Keys, Internet-scale ip alias resolution techniques, ACM SIGCOMM Computer Communication Review, 2010
[5] M. H. Gunes and K. Sarac, ”Resolving IP aliases in building traceroute-based internet maps”, IEEE/ACM Transactions on Networking, December 2009.
[6] M. E. Tozal and K. Sarac, ”PalmTree: An IP Alias Resolution Algorithm with Linear Probing Complexity”, Computer Communications Volume 34, Issue 5, pp. 658-669, April 2011
[7] K. Keys, ”Internet-Scale IP Alias Resolution Techniques”, ACM SIGCOMM Computer Communication Review, January 2010.
[8] A. Y. Nur and M. E. Tozal, ”Identifying Critical Autonomous Systems in the Internet”, The Journal of Supercomputing 74.10 (2018): 4965-4985
[9] K. Yoshida, Y. Kikuchi, M. Yamamoto, Y. Fujii, K. Nagami, I. Nakagawa, and H. Esaki, ”Inferring PoP-level ISP Topology through End-to-End Delay Measurement, Springer International Conference on Passive and Active Network Measurement, 2009
[10] N. Spring, R. Mahajan, and D. Wetherall, ”Measuring ISP Topologies with Rocketfuel”, ACM SIGCOMM 2002.
[11] M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush, ”10 Lessons from 10 Years of Measuring and Modeling the Internet’s Autonomous Systems”, JSAC, 2011
[12] L. Gao, On Inferring Autonomous System Relationships in the Internet, IEEE/ACM Transactions on Networking, 2001
[13] R. Govindan and H. Tangmunarunkit, ”Heuristics for Internet Map Discovery”, IEEE INFOCOM, March 2000
[14] A. Y. Nur and M. E. Tozal, ”Record Route IP Traceback: Combating DoS Attacks and the Variants”, Computers & Security 72 (2018): 13-25
[15] P. Merindol, B. Donnet, J. Pansiot, M. Luckie, and Y. Hyun, ”Merlin: Measure the Router Level of the Internet”, Next Generation Internet (NGI), 2011
[16] A. Bender, R. Sherwood, and N. Spring, ”Fixing Ally’s Growing Pains with Velocity Modeling”, ACM SIGCOMM conference on Internet measurement, 2008
[17] V. Padmanabhan and L. Subramanian, ”An Investigation of Geographic Mapping Techniques for Internet Hosts”, ACM SIGCOMM, 2001
[18] B. Gueye, A. Ziviani, M. Crovella, and S. Fdida, ”Constraint-based Geolocation of Internet Hosts”, IEEE/ACM Transactions on Networking, December 2006
[19] E. Katz-Bassett, J. P. John, A. Krishnamurthy, D. Wetherall, T. Anderson, and Y. Chawathe, ”Towards IP Geolocation Using Delay and Topology Measurements”, IMC, 2006
[20] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, ”iPlane: An Information Plane for Distributed Services”, OSDI, 2006
[21] D. Feldman, Y. Shavitt, and N. Zilberman, ”A Structural Approach for PoP Geo-location”, Computer Networks, 2012
[22] H. Chang, S. Jamin, and W. Willinger, ”Inferring AS-Level Internet Topology from Router-Level Path Traces”, Scalability and Traffic Control in IP Networks (SPIE), 2001
[23] Z. Mao, J. Rexford, J. Wang, and R. Katz, ”Towards an Accurate AS-Level Traceroute Tool”, ACM SIGCOMM, 2003
[24] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, k. claffy, and A. Vahdat, ”The Internet AS-Level Topology: Three Data Sources and One Definitive Metric”, ACM SIGCOMM Computer Communication Review, 2006
[25] V. Giotsas, M. Luckie, B. Huffaker, and k. claffy, ”Inferring Complex AS Relationships”, Internet Measurement Conference (IMC), Nov 2014
[26] Usage of the Internet http://www.internetlivestats.com/internet-users/
[27] Caida AS Ranking - https://asrank.caida.org/
[28] Routeviews Prefix to AS mappings Dataset for IPv4 and IPv6 - http://www.caida.org/data/routing/routeviews-prefix2as.xml
[29] CAIDA Prefix-Probing Traceroute Dataset - 2021/01/01, http://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml
[30] DB-IP Geolocation Database - https://www.db-ip.com/
[31] GeoLite2 Geolocation Database - www.maxmind.com
[32] LITE Geolocation Database - http://lite.ip2location.com/
[33] RIS RIPE - 2021/01/01, https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data
[34] The Route Views Project - 2021/01/01, http://archive.routeviews.org/