Identifying Critical Autonomous Systems in the Internet

Abdullah Yasin Nur
Mehmet Engin Tozal

Abstract The Internet not only facilitates our daily activities, such as communication, entertainment and shopping but also serves as the enabling technology for many critical services, including finance, manufacturing, healthcare and transportation. On the other hand, a wide spectrum of attacks target its communication infrastructure to disable or disrupt the network connectivity and traffic flow until recovery processes take place. Attacking all Autonomous Systems (ASes) in the Internet is typically beyond the capability of an adversary. Therefore, targeting a small number of ASes which results in the highest impact is the best strategy for attackers. Similarly, it is important for network practitioners to identify, fortify and secure those critical ASes to mitigate the impact of the attacks. In this study we introduce an intuitive and effective measure, IP address spatial path stress centrality, to assess and identify the critical ASes in the Internet. We compare IP address spatial path stress centrality to the three well known and widely used centrality measures, namely customer-cone size, node degree and betweenness. We demonstrate that the proposed measure incorporates business relations and IP address spaces to achieve a better measure for identifying the critical ASes in the Internet.

Keywords Autonomous Systems · Internet Security · Complex Systems

1 Introduction

The Internet is not only a critical infrastructure but also an enabling technology for many other critical services. It is a highly engineered, large scale

A. Y. Nur and M. E. Tozal
School of Computing and Informatics
University of Louisiana, Lafayette
Lafayette, LA 70504 USA
E-mail: {ayasinnur, metozal}@louisiana.edu

Preprint, The Journal of Supercomputing, Springer
https://doi.org/10.1007/s11227-018-2336-3

complex system which has no central governance. The global communication infrastructure of the Internet is formed by thousands of autonomous networks connecting various organizations and individuals together. These autonomous networks are owned and operated by a diverse set of organizations including businesses, network service providers, cloud providers, web hosting companies, universities and government agencies all around the world.

A group of networks managed by one or more operators under a well defined routing policy is called an *Autonomous System (AS)* in the Internet [1]. Autonomous Systems (ASes) are identified by unique AS numbers and they connect to each other in different forms to enable the "global" Internet communication [2]. Individual users, small businesses and ASes located at the edge of the Internet participate in the global infrastructure by means of other ASes called *Internet Service Providers* (ISPs). Typically, ISPs are business entities providing Internet access service to their customers while getting the same service from one or more upstream ISPs. At the core of the Internet, a small number of ISPs peer with each other through settlement-free interconnections to attain the global communication infrastructure.

The majority of the ASes (around 85%) are located at the edge of the Internet and they are solely Internet access consumers. That is, they pay to ISPs to acquire global Internet access service. Note that these ASes may be content or service providers, yet they are consumers in terms of the Internet access service. The ASes forming the communication infrastructure in the center, on the contrary, are Internet access consumers and providers, simultaneously. They provide the Internet access service to each other and consume the service from each other. Internet access service is provided and consumed with respect to the business relations among ASes. That is, ASes connect to each other via business relations that define the characteristics of the Internet access service. More importantly, inter-AS traffic in the Internet is usually routed according to the business relations among the ASes [3].

Traditionally, business relations between ASes are categorized as customer-to-provider (c2p), peer-to-peer (p2p) and sibling-to-sibling (s2s) [4]. In a c2p relation, the provider AS provides global reachability to its customer AS. In return, the customer pays to the provider for the traffic exchanged between them. In a p2p relation, two peer ASes provide mutual reachability to each other and their customer ASes, recursively. Peer ASes typically engage in settlement-free business agreements which means that neither party pays to the other for the traffic exchanged. In the less frequently observed s2s relation, two ASes provide full reachability to each other because they are operated by the same or sibling organization(s). More complex relations such as hybrid relations and partial relations are also reported in the Internet [5]. However, c2p and p2p relations abstract the majority of the business agreements between ASes for practical purposes [3].

Figure 1 shows the AS-level topology graph of the Internet obtained from CAIDA [6].The topology graph consists of 54,140 ASes connected to each other through 466,190 relations (logical links). Among those ASes 45,796 (85%) are located at the edge of the Internet without having any customer ASes.

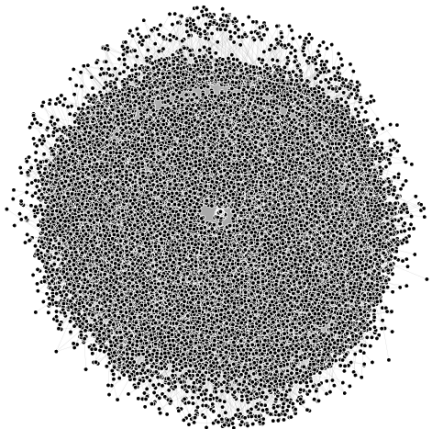

Fig. 1: Internet topology graph consisting of 54,140 ASes and 466,190 relations (drawn using the Kamada-Kawai layout algorithm).

Put in other words, 45,796 ASes are solely Internet access consumers and the remaining 8344 (15%) ASes provide Internet access service to organizations and individuals. Out of 466,190 relations among the ASes, 107,195 (23%) are c2p and 358,995 (77%) are p2p relations.

The security of the ASes forming the communication infrastructure is of the utmost importance because many critical services depend on the Internet as an enabling technology. Various types of attacks such as distributed denial of service [7], crossfire [8], link cut [9], coordinated cross plane session termination [10], unauthorized router access and session hijacking [11] target the ASes, especially the ISPs, in the Internet. The goal of those attacks is to disable or disrupt the network connectivity and traffic flow until recovery processes take place. Attacking all ASes in the Internet is typically beyond the capability of an adversary. Hence, it is important for an adversary to choose a small set of target ASes which results in the maximum traffic disruption in the Internet. Similarly, it is important for network practitioners such as chief information officers and IT managers to identify, fortify and secure those critical ASes to mitigate the impact of the attacks.

In our recent work we investigated AS rankings based on their topological characteristics including customer degree, provider degree, peer degree, customer-cone size, alpha centrality and betweenness centrality [12]. One major observation is that the centrality measures based on the structural characteristics of the Internet topology graph fall short to capture the importance of ASes under targeted attacks. In this study we introduce *IP (Internet Protocol) address spatial path stress* as a measure to identify and group the critical ASes under targeted attacks in the Internet. We define the criticality of an AS as the amount of potential traffic that it carries between the pairs of other ASes. Hence, the criticality of an AS is proportional to the number of the inter-AS paths passing through it as well as the amount of traffic carried via each path. Those ASes are good candidates for attackers because they allow

an adversary to disrupt a greater portion of the Internet traffic and negatively affect more users. To compute the paths between pairs of ASes we use the policy-preferred inter-AS path enumeration algorithm introduced in our earlier work [13, 14]. The policy-preferred paths are free from the artifacts of shortest paths in AS-level Internet graphs, such as inflated number of paths, policy inconsistent paths and undesirable paths. To approximate the potential traffic exchange between pairs of ASes we use the IP address spaces of ASes. Our observations show that large scale content consumers such as universities, government agencies and extensive businesses own larger IP address blocks translating into larger IP address spaces. Additionally, large scale content providers such as online social media, web hosting companies and content delivery networks own many IP address blocks summing into larger IP address spaces. Therefore, the IP address spaces of ASes can serve as a relative approximation of the potential traffic exchange between the ASes.

We experimentally compare the IP address spatial path stress centrality to the three well known and widely used centrality measures, namely customer-cone size, node degree and betweenness. Customer-cone size is widely used in ranking the ASes in the Internet [3]. Degree and betweenness are two common measures in assessing the importance of nodes in complex systems [12, 15, 16, 17]. We demonstrate that the proposed IP spatial stress centrality incorporates business relations and IP address spaces to achieve a better measure for identifying the critical ASes in the Internet. Our empirical results show that the most critical AS in the Internet as of this writing is $AS1299$, $AS174$ and $AS3356$ run by Telianet, Cogent Communications and Level 3 Communications, respectively.

The rest of the paper is organized as follows. We present the related work in the next section. Section 3 gives an overview of inter-AS traffic routing and policy-preferred AS paths in the Internet. We introduce our approach for assessing AS criticality levels in Section 4. In Section 5, we present our experimental results. We discuss threats and possible defense mechanisms in Section 6. Finally, Section 7 concludes the paper.

2 Related Work

Different measures have been introduced in the literature to rank, classify and cluster the autonomous systems in the Internet for various purposes.

A widely known measure to compare and rank the ASes in the Internet is AS customer-cone [3]. In its generic form, the customer-cone of an AS is the set of ASes consisting of the AS itself, its customer ASes and the customer-cones of those customer ASes. This measure reflects the position of an AS in the semi-hierarchical structure of the Internet as well as the routing influence of the AS in the Internet. However, multi-homing and peering practices in the Internet introduces multiple paths that bypass upstream providers. Therefore, the routing influence of an AS in the Internet might be different from the one reflected by its customer-cone.

Zimmerli et al., [18] suggested an AS rating approach based on the traceroute collected performance metrics. They ranked ASes based on the network performance within the ASes and the performance of their neighboring ASes. This ranking scheme is highly volatile because it is sensitive to the real time changes in the Internet. Besides, it is difficult to scale the technique to the entire Internet.

Clérot and Nguyen proposed an AS ranking heuristic based on the concept of alpha-centrality in social network analysis [19]. Their heuristic starts from an undirected graph of ASes and gradually introduces asymmetry by allowing directed edges reinforce the relationships between ASes. The rank score of an AS includes the centrality of the AS as well as the centralities inherited from the neighboring ASes. The authors show that the ranking results of their heuristic is quite close to the simple degree-based centrality. This method requires careful selection of parametric values and it may artificially rank the ASes with many neighbors higher.

In another study, Wagner et al., [20] proposed an AS ranking method for detecting the ASes which provide transit services to the ASes that host malicious software and services. The authors use existing AS scores reflecting the malware hosting capacity of ASes to annotate AS graphs and use PageRank to rank the ASes.

Finally, in [21] the authors classify ASes as large ISPs, small ISPs, customer ASes, university ASes, Internet exchange point ASes and network information center ASes using supervised learning.

In this study, we develop an intuitive and effective measure to assess the criticality levels of ASes from a targeted-attacks perspective. Our measure incorporates IP address spaces of ASes and policy-preferred paths between ASes together to evaluate ASes' impact on the overall Internet traffic under targeted attacks.

3 Background

Before introducing our approach to assess the criticality levels of ASes, we present a brief background on inter-AS traffic routing and policy-preferred AS paths in the Internet.

3.1 Inter-AS Traffic Routing

ASes in the Internet use the Border Gateway Protocol (BGP) [22] to exchange information about how to reach blocks of contiguous IP addresses (IP address prefixes). Essentially, the reachability information consists of an IP address prefix, one or more AS paths to reach the prefix and a set of AS path attributes. BGP supports a wide variety of AS path attributes and allows prefix withdrawals as well [22]. An AS willing to deliver traffic to the devices within an IP address prefix originates a BGP advertisement declaring the prefix and its AS number as the path to the prefix. This advertisement is sent to the neighboring ASes of the originating AS. The neighboring ASes independently

decide to employ, drop and/or re-advertise the new IP address prefix with or without modifying any AS path attributes. A neighboring AS willing to transit traffic destined to the new IP address prefix, re-advertises the prefix to its own neighbors by prepending its AS number in the path. The neighbors of a re-advertising AS repeat the same process. Hereby, multiple AS paths to an IP address prefix gets disseminated in the Internet through neighbor-to-neighbor BGP advertisements while each AS independently selects/employs a path(s) toward the prefix. The traffic however, follows the reverse AS path direction to reach from a source AS to the destination AS that originated the prefix.

BGP protocol allows a path toward a routing prefix to be incrementally disseminated in the Internet through neighbor-to-neighbor advertisements. However, the ASes do not have to re-advertise a prefix that they learn from a neighbor to their other neighbors. In fact, AS path advertisements are locally assessed according to the business relations among ASes and their neighbors. Typically, an AS receiving a prefix advertisement from one of its customers re-advertises the prefix to its providers, peers and other customers because it charges the advertising customer for the transit traffic. Similarly, an AS receiving a prefix advertisement from one of its peers re-advertises the prefix only to its customers because transiting traffic between a peer and a provider costs money and transiting traffic between two peers adds additional load on its network without any financial gain. An AS receiving an advertisement from multiple neighbors prefers the path from a customer over a peer and from a peer over a provider. Because, ASes charge their customers, do not pay to their peers and pay to their providers for the traffic exchanged between them, respectively. Finally, ASes prefer the shorter paths over the longer equal-cost paths. In summary, the existence of an AS path in an Internet topology graph does not necessarily mean that the path is promoted by BGP for utilization. The paths are utilized according to the business relations between the ASes.

3.2 Policy-Preferred AS Paths

In the previous part we outlined how AS path information for an IP address prefix propagates in the Internet. An AS path from a source AS to a destination AS in a topology graph reflects the path taken to reach the IP address prefixes originated by the destination AS. Hence, computing AS paths in a business relations annotated AS-level Internet topology graph helps us to sketch the inter-AS traffic routes in the Internet. However, reducing an AS-level Internet topology map into an undirected graph and computing the shortest paths between pairs of ASes do not reflect the actual paths employed. Simply, it ignores the business relations or policies between ASes. Therefore, it usually inflates the number of paths between ASes; introduces erroneous paths that do not conform to economic policies; and/or generates symmetric paths, which in reality is not a rule.

In our earlier work we introduced a single-destination, policy-preferred path enumeration algorithm which discovers policy consistent paths from all ASes to a given destination AS in an AS-level Internet topology graph [13, 14]. The

algorithm provides a holistic solution to the AS-level path enumeration problem by incorporating common practices and incentives in inter-AS routing, including first-hop-edge policy preferred paths, valley-free preferred paths, and shortest-distance, equal-cost preferred paths [13]. Given an AS-level Internet topology graph and a destination vertex, the algorithm starts from the destination vertex and incrementally builds AS paths in backwards from source vertices toward the destination vertex. At each iteration, a new vertex is joined to the subgraph of the established, policy-preferred paths toward the destination vertex via one or more edges. At the end, the algorithm returns a rooted, directed, acyclic subgraph (r-DAG) of the input graph, which is formed by policy-preferred paths from the source vertices toward the destination vertex. The time complexity of the algorithm is the same as Dijkstra's shortest path algorithm with a priority queue implementation.

The proposed IP address spatial path stress centrality in this study utilizes the policy-preferred AS paths, because the shortest paths algorithm on the undirected graph representation of the Internet introduces erroneous paths that violate policy consistency.

4 Methodology

In this section, we introduce autonomous system IP address spaces as a heuristic for the potential traffic intensity of AS paths in the Internet. Next, we develop IP address spatial path stress centrality as a measure to identify the critical ASes in the Internet.

4.1 AS IP Address Spaces

Enumerating policy-preferred paths in an AS-level Internet topology graph helps us to learn the route(s) from a source AS to a destination AS toward the IP address prefixes originated by the destination AS. However, it does not tell us anything regarding the potential traffic intensity between the two ASes.

In this study we define the *IP address scope* of an advertised routing prefix as the number of the assignable IP addresses of the prefix. The scope of an IPv4 routing prefix, p, of prefix length l is $2^{(32-l)}$ including the subnet and broadcast addresses. Since an AS can originate more than one routing prefix, we define the *IP address space* of an AS, $\mathcal{I}_{AS}$, as the sum of the IP address scopes of the originated routing prefixes.

$$\mathcal{I}_{AS} = \sum_{p_i \in AS} 2^{(32-l_i)} \tag{1}$$

where p_i is an IP address prefix originated by the AS and l_i is the corresponding prefix length.

We heuristically state that the paths between ASes having larger IP address spaces have more traffic intensity compared to the paths between ASes having smaller IP address spaces. Our heuristic is based on the following observations:

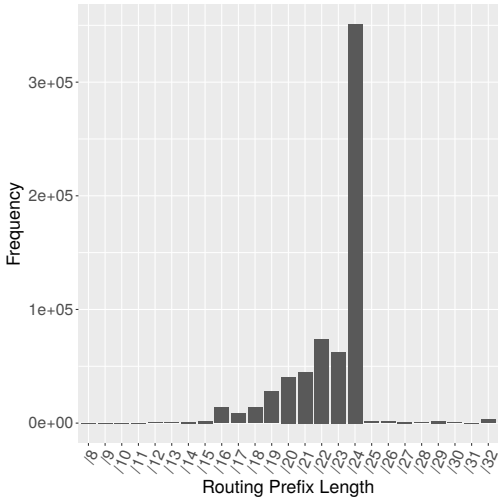

Fig. 2: Advertised prefix length distribution.

– ASes that belong to large scale content providers, web hosting companies and content delivery networks, e.g., Facebook, Godaddy and Akamai, advertise multiple routing prefixes that sum up to larger IP address spaces.
– ASes belonging to large scale private/public organizations, universities and government agencies, e.g., National Institute of Standards and Technology and University of Louisiana, advertise one or more large scope routing prefixes that translate into larger IP address spaces.
– ASes that belong to residential and mobile Internet access providers, e.g., Cox and Verizon, advertise multiple routing prefixes that sum up to larger IP address spaces.

The first and second observations imply that content provider networks have larger IP address spaces. The second and third observations imply that content consumer networks also have larger IP address spaces. As a practical heuristic, the paths between larger IP address space ASes potentially have higher traffic intensity compared to the paths between smaller IP address space ASes.

Note that the scope of an IP address prefix does not directly correspond to the in-use IP addresses, i.e., the prefix may be underutilized. An alternative method to estimate the IP address spaces of ASes is to actively probe all or a sample of their routing prefixes. This approach is costly because of the long probing duration and significant probing traffic overhead. Moreover, it introduces its own type of artifact due to private IP address deployments behind NAT (Network Address Translation) boxes, non-responsive host configurations and rate limiting practices by ISPs. On the other hand, the best approach which is having a global inter-AS traffic matrix compiled by ISPs is not available to the best of our knowledge. Typically, network operators do not share this information due to security and business concerns.

The ASes in our dataset originate more than 649,701 routing prefixes. Figure 2 shows the prefix length distribution of the advertised routing prefixes. In the Figure, 98% of the prefix lengths are between /16 and /24 such that 55% of them are /24s. An examination of Figure 2 suggests that the larger IP address scopes (smaller prefix length) on the left tail and the smaller scopes (larger prefix length) on the right tail are outliers. To reduce the impact of the outliers on IP address space estimations we replace the prefix lengths on the left tail by a /16 and the ones on the right tail by a /24. Note that our technique is similar to omitting the outliers on the tails of an empirical distribution. Instead of omitting the outliers we project them, because those outlying routing prefixes accommodate in-use IP addresses as well.

4.2 Identifying Critical ASes

In this part we develop *IP spatial path stress* centrality, $\mathcal{C}$, to identify the critical ASes in an AS-level topology graph of the Internet. Let $\mathcal{P}_{S,T}(R) = \{S, \ldots, R, \ldots, T\}$ be a sequence of ASes denoting a path between a source AS, S, and a destination AS, T, passing through an intermediate AS, R. Let $\mathcal{I}_S$ and $\mathcal{I}_T$ be the IP address spaces of the source AS, S, and the destination AS, T, respectively. Let $\mathcal{T}_{S,T}$ be the traffic intensity of the path $\mathcal{P}_{S,T}(.)$. Based on the discussions in Section 4.1, we define the traffic intensity of the path $\mathcal{P}_{S,T}(.)$ as

$$\begin{aligned} \mathcal{T}_{S,T} &= \mathcal{I}_S \mathcal{I}_T \\ &= \left(\sum_{p_i \in S} 2^{(32-l_i)} \right) \left(\sum_{p_j \in T} 2^{(32-l_j)} \right) \end{aligned} \tag{2}$$

In practice, Equation 2 can be normalized by either taking the logarithm of the traffic intensity or by dividing it by 2^{64}. We define the IP spatial path stress centrality, $\mathcal{C}_R$, of the intermediate AS, R, as

$$\mathcal{C}_R = \sum_{\forall \mathcal{P}_{S,T}(R)} \mathcal{T}_{S,T} \tag{3}$$

such that $S \neq T \neq R$. The IP spatial path stress of an AS is equal to the sum of the traffic intensities of the paths passing through it. The centrality measure not only reflects the number of paths passing through an AS but also the inferred intensity of the traffic transited by the AS in the Internet. Naturally, those ASes having higher IP spatial path stress values are good candidates for attacks because they allow an adversary to disrupt a greater portion of the Internet traffic and negatively affect more users.

5 Experimental Results

In this section we experimentally analyze the results of the IP spatial stress centrality as well as compare it to other centrality measures. In the first part,

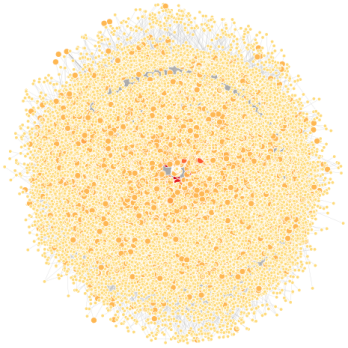

Table 1: AS frequency distribution by criticality levels

C-1	1	**C-8**	3	**C-15**	8	**C-22**	33
C-2	1	**C-9**	2	**C-16**	5	**C-23**	42
C-3	2	**C-10**	3	**C-17**	5	**C-24**	63
C-4	2	**C-11**	4	**C-18**	8	**C-25**	82
C-5	1	**C-12**	2	**C-19**	9	**C-26**	177
C-6	2	**C-13**	3	**C-20**	8	**C-27**	447
C-7	3	**C-14**	3	**C-21**	21	**C-28**	7121

Fig. 3: AS-level Internet topology map. ASes are clustered by their levels of criticality.

we compute the IP spatial stress centrality on the observed topology of the Internet and examine the most critical ASes in the Internet. In the second part, we compare our results to the three well known and widely used AS centrality measures, namely customer cone size, node degree and betweenness [3, 12, 15, 16, 17].

5.1 Empirical Analysis

In the following, we compute the IP spatial path stress centralities of the ASes in our dataset and cluster the ASes with respect to their criticality levels. We use ck-means algorithm [23] to cluster the ASes with respect to their IP spatial stress centralities. The ck-means algorithm uses a dynamic programming strategy to cluster univariate data by minimizing the total within-cluster sums of squares. Different from the classical k-means algorithm, ck-means finds a unique, optimal cluster separation for one dimensional data and guarantees reproducible results.

The ASes at the edge of the Internet do not transit any traffic belonging to other ASes. Therefore, the IP spatial stress centralities of those ASes are zero and they form their own cluster. Transit ASes forming the communication infrastructure of the Internet, however demonstrate 28 distinct clusters such that the top level critical AS cluster is **C-1** and the bottom level cluster is **C-28**. Figure 3 shows the critical ASes color and size differentiated where the IP spatial stress centrality increases from orange to red. The ASes having zero centrality are shown in yellow in the figure. Table 1 shows the AS frequency distribution by the levels of criticality. The table shows that the bottom eight levels of criticality (**C-21** thru **C-28**) accommodate majority of the transit ASes; 7,986 ASes in total. On the other hand, the top eight levels of criticality (**C-1** thru **C-8**) accommodate only 15 ASes.

Table 2 shows the top 15 ASes in the top eight levels of criticality as well as their organizations, the number of their customers, providers and peers. The top AS in the table is $AS1299$ run by Telia international carrier. Telianet

Table 2: Top-15 critical ASes in the Internet w.r.t IP spatial stress centrality

Rank	AS Number	Organization	Customers	Providers	Peers	Cluster
1	AS1299	Telianet	1257	0	174	C-1
2	AS174	Cogent Comm.	4881	0	225	C-2
3	AS3356	Level 3 Comm.	4368	0	62	C-3
4	AS6939	Hurricane Elect.	986	2	4596	C-3
5	AS3257	GTT Comm.	1376	0	155	C-4
6	AS2914	NTT Comm.	1364	0	100	C-4
7	AS3549	Level 3 Comm.	1045	3	2709	C-5
8	AS7018	AT&T Services	2360	0	57	C-6
9	AS6453	Tata Comm.	655	0	95	C-6
10	AS2516	KDDI Corp.	214	6	111	C-7
11	AS4809	China Telecom	127	15	43	C-7
12	AS701	MCI Comm.	1328	0	30	C-7
13	AS209	Qwest Comm.	1685	0	71	C-8
14	AS12989	Eweka Internet	91	13	1699	C-8
15	AS43531	IX Reach	232	4	1813	C-8

is based in Sweden and it is Europe's largest telephone and mobile network provider operating in Europe and Asia. $AS1299$ is followed by $AS174$ (Cogent Communications), and $AS3356$ (Level 3 Communications). Cogent Communications is a multinational ISP based in the US and it is specialized in providing high speed Internet access service all over the world. Level 3 is another US based telecommunications company providing Internet access service to medium sized ISPs in North America, Latin America and Europe.

In Table 2, nine ASes, i.e., $AS1299$, $AS174$, $AS3356$, $AS3257$, $AS2914$, $AS7018$, $AS6453$, $AS701$ and $AS209$, are tier-1 ASes that bind the Internet together only through peer relations. These ASes are located at the core of the Internet and they do not have any providers as shown in the "Providers" column. Moreover, these tier-1 ASes in Table 2 have relatively higher number of customers and lower number of peers. Further analysis shows that those ASes have very high number of c2p descendants (customer-cone). Since those tier-1 ASes undertake the role of bridging different parts of the Internet, they appear on many paths between pairs of ASes and have very high IP spatial stress centralities.

A more interesting observation in Table 2 is the six transit ASes, i.e., $AS6939$, $AS3549$, $AS2516$, $AS4809$, $AS12989$ and $AS43531$, that are not considered as tier-1. These ASes are not part of the tier-1 ASes because they do not participate in the largest, fully connected clique in the Internet, i.e. they attain the global Internet access through their providers and peers. Yet, they have very high IP spatial stress centralities, even higher than some of the tier-1 ASes. Among those ASes, $AS6939$, $AS3549$, $AS12989$ and $AS43531$ have very high number of peers. Peering allows those ASes to appear more frequently on the paths between their own descendants as well as their peer's descendants which also have larger IP address spaces. On the other hand, $AS2516$ and $AS4809$ have relatively small number of customers and peers. Further analyses show that these ASes have large number of descendants with large IP address spaces and they peer with similar ASes. Therefore, they appear on many paths having high inferred traffic intensity as well.

5.2 Comparitive Analysis

In this part, we compare IP spatial stress centrality to other three well known and widely used AS centrality measures, namely customer cone size, node degree and betweenness.

AS customer-cone size is a widely used measure to study the ASes in the Internet in terms of their routing capability [3]. In general, customer-cone of an AS is recursively defined as a set consisting of the AS itself along with its customers' customer-cones. That is, the customer-cone of an AS is a set consisting of the AS itself and its customer descendants. The customer-cone of an AS corresponds to a sub-topology where the connected component is formed through c2p relations. Customer-cone size, the number of ASes in the customer-cone of an AS, may show the importance of the AS regarding the global traffic routing in the Internet.

Degree is another measure widely used in complex systems domain to examine the "key" or "important" nodes in a graph [17]. It is defined as the number of edges of a given vertex in an undirected graph. For directed graphs, indegree and outdegree centralities specify the number of incoming and outgoing edges, respectively. To compute the degree, we transformed the AS topology map of the Internet into an undirected graph consisting of 54,140 nodes and 466,190 links. Since the ASes at the edge of the Internet do not transit any inter-AS traffic and collectively have zero IP spatial stress centrality, we use the ASes in the center of the Internet for the comparison.

Betweenness is a measure that quantifies the centrality of a vertex in terms of its involvement in connecting pairs of vertices in a graph [15, 24]. Formally, betweenness centrality of a vertex v_k is defined as $\beta(v_k) = \sum \sigma_{v_i v_j}(v_k)/\sigma_{v_i v_j}$ such that $\sigma_{v_i v_j}$ is the number of the shortest paths between vertices v_i and v_j and $\sigma_{v_i v_j}(v_k)$ is the number of those paths that pass through v_k where $v_i \neq v_j \neq v_k$. Betweenness is used to assess the load of nodes in telecommunication networks [15, 25]. Similar to the degree, we use the undirected graph representation of the Internet to compute betweenness and used only the transit ASes in the center for the comparison. The betweenness' of the edge ASes are zero, since they do not transit any traffic belonging to other ASes.

To motivate the reader, Table 3 shows the top-15 ASes ranked based on IP spatial stress, degree, customer-cone size and betweenness. IP spatial stress centrality (first column) and customer-cone size (third column) have ten ASes in common in their top-15 lists. However, none of those ASes are ranked at the same position in both ranking schemes. Please note that the the percentage of common ASes quickly decays for top-k lists where $k \leq 300$ (See Section 5.2.2). Similarly, IP spatial stress centrality shares eight ASes and ten ASes with degree (second column) and betweenness (fourth column) in their top-15 lists, respectively. However, only two of those ASes, AS174 and AS3356, appear at the same position between IP spatial stress and degree centralities as well as IP spatial stress and betweenness centralities. AS174, belonging to Cogent Communications, is a multinational, tier-1 ISP based in the US. It has 4,881 direct customers that primarily use AS174 for Internet access. It frequently

Table 3: Top-15 critical ASes in the Internet w.r.t different AS characteristics

	IP Spatial Stress	Degree	Customer Cone Size	Betweenness
1	AS1299	AS6939	AS3356	AS6939
2	AS174	AS174	AS1299	AS174
3	AS3356	AS3356	AS174	AS3356
4	AS6939	AS3549	AS3257	AS3549
5	AS3257	AS24482	AS2914	AS7018
6	AS2914	AS7018	AS6453	AS1299
7	AS3549	AS8220	AS4436	AS209
8	AS7018	AS43531	AS701	AS4323
9	AS6453	AS20485	AS6762	AS2914
10	AS2516	AS4323	AS7018	AS701
11	AS4809	AS36351	AS6939	AS3257
12	AS701	AS12989	AS209	AS6461
13	AS209	AS10026	AS3320	AS9498
14	AS12989	AS209	AS5511	AS20485
15	AS43531	AS34224	AS1239	AS2828

appears on the paths from those customers to other ASes in the Internet, which increases its rank in terms of both betweenness and IP spatial stress centrality. Similarly, AS3356, owned by Level 3 Communications, is a multinational, tier-1 ISP based in the US. Again, it has a high number of direct customers, 4,368, which improves its rank in terms of both betweenness and IP spatial stress centrality.

The rank discrepancies among top-15 lists of IP spatial stress centrality, customer-cone size, degree and betweenness do not solely demonstrate the overall discrepancy among those different ranking schemes. First of all, the amounts of concordance/discordance in top-15 lists are not representative for the whole dataset. Second, one may in general be interested in top-k lists such that k assumes any positive integer, e.g., 20, 50 or 100. In the following we first study the amount of concordance/discordance among rankings by different characteristics in the whole dataset. Next, we extend our analysis to the top-k lists where k is a positive integer less than the dataset size.

5.2.1 Complete Rank Correlations

Given a finite set of objects, $S = \{s_1, s_2, \ldots, s_{|\mathcal{S}|}\}$, *ranking* is a binary relation $\mathcal{R} = \{(s_i, s_j) \subset \mathcal{S} \times \mathcal{S}\}$ denoting the first element "precedes" (or "succeeds") the second element while satisfying irreflexivity, $(s_i, s_i) \notin \mathcal{R}$; antisymmetry, $(s_i, s_j) \in \mathcal{R} \Rightarrow (s_j, s_i) \notin \mathcal{R}$; and transitivity, $(s_i, s_j) \in \mathcal{R}, (s_j, s_k) \in \mathcal{R} \Rightarrow (s_i, s_k) \in \mathcal{R}$ [12]. We adopt the matrix notation introduced by Emond and Mason [26] to represent ranking relations among the objects of a set. An $|S| \times |S|$ rank matrix, $\boldsymbol{R}$, over a set S is formulated as follows:

$$\boldsymbol{R}[i,j] = \begin{cases} 1 & \text{if } i \text{ precedes or tied with } j \\ 0 & \text{if } i = j \\ -1 & \text{if } i \text{ succeeds } j \end{cases} \tag{4}$$

Equation 4 allows the distances between ranking schemes abide by Kemeny-Snell axioms ensuring non-negativity, symmetry, triangle inequality and rank consistency. A rank correlation coefficient is a statistic for measuring the

strength of ordinal association between two ranking schemes. Given two rank matrices $\boldsymbol{R}_A$ and $\boldsymbol{R}_B$ (Equation 4) over a set S, τ_x [26] is defined as:

$$\tau_x = \frac{\sum_{i=1}^{|S|}\sum_{j=1}^{|S|} \boldsymbol{R}_A[i,j]\boldsymbol{R}_B[i,j]}{|S|(|S|-1)} \tag{5}$$

where $|S|$ denotes the cardinality of S. The numerator of Equation 5 increases as the rank matrices $\boldsymbol{R}_A$ and $\boldsymbol{R}_B$ are concordant on the relative orderings of pairs of objects. Similarly, the numerator decreases for discordant object pairs under the two ranking schemes. τ_x takes values between -1 and 1 such that -1 denotes perfect disagreement and 1 denotes perfect agreement between the two ranking schemes.

Table 4 shows the rank correlations between IP spatial stress and customer-cone size, degree and betweenness. The table does not demonstrate any strong rank correlation between IP spatial stress and other AS characteristics.

Table 4: Rank Correlation Coefficient τ_x

	Degree	**Customer-Cone Size**	**Betweenness**
IP Spatial Stress	0.39	0.53	0.51

Degree is a measure to evaluate the criticality of a node under targeted attacks in complex systems, especially in scale-free graphs. We do not observe a strong correlation between degree and IP spatial stress in Table 4, because higher degree does not necessarily mean an AS is more critical in the Internet topology. First of all, ASes having many providers and less many customers do not carry inter-AS traffic for their providers. Yet, their providers contribute to their total degree. Second, 77% of the edges in our Internet topology graph are p2p links. Although, an AS may have many customers, those customers prefer to use their p2p links to route their traffic instead of using their providers via c2p links. IP spatial stress centrality, the measure introduced in this study, accounts for those paths employing p2p links while degree centrality simply fails to capture their impact on traffic routing.

We observe a moderate level of correlation, 0.53, between IP spatial stress and customer-cone size (Table 4). Customer-cone size has been an important metric that reflects the routing capability of an AS. However, the Internet has evolved from a semi-hierarchical topology to a flatter topology in the last decade [27] mostly due to the increasing number of p2p links. As a result, the descendants of an AS having a large customer cone, prefer routing the traffic through their peers instead of upstream provider(s) when possible. IP spatial stress centrality naturally accounts for those paths that do not utilize upstream providers in the semi-hierarchical topology.

Similarly, we observe a moderate level of correlation, 0.51, between IP spatial stress and betweenness (Table 4). Note that betweenness is usually computed on an undirected graph transformation of the Internet topology map.

Therefore, betweenness accounts for shortest paths between pairs of ASes regardless of the types of relations, i.e., c2p, p2p, among ASes. On the other hand inter-AS traffic in the Internet is routed according to the economic incentives which are reflected by relation types among ASes. Specifically, an AS prefers to use a longer path over the shortest path if the longer path is economically more advantageous [13, 14]. IP spatial stress centrality uses policy-preferred paths which is congruent with the economy of the Internet. Moreover, betweenness ignores the IP address spaces of ASes. That is, a path incident on a large IP address space AS is equivalent to a path incident on a small IP address space in terms of the impact. IP spatial stress centrality, on the other hand accounts for the IP address spaces of source and destination ASes of inter-AS paths (Equation 3).

5.2.2 Incomplete Rank Correlations

Above, we studied the correlation between IP spatial stress centrality and degree, customer-cone size and betweenness over the whole dataset. However, one may be interested in only top-k critical ASes and the amount of concordance/discordance among top-k lists may become more important then the correlation over the whole dataset. Moreover, the discrepancies among different measures in complete rank correlations do not necessarily imply that they hold in incomplete rank correlations as well. In the following we analyze the correlation of top-k lists between IP spatial stress centrality and degree, customer-cone size and betweenness.

One particular problem in rank correlation for top-k lists is incomplete rankings. That is, two top-k lists might have objects that show up in one list but not appear in the other list. We extend Emond-Mason τ_x for incomplete rankings by appending the objects appearing in one list to the other while preserving the order in the former list. The append operation is applied to both lists to achieve a common domain between both top-k lists. The intuition behind the append operation is that the objects appearing in one top-k list but not appearing in the second are ranked lower than all objects in the second top-k list. Otherwise, they would have appeared in the top-k of the second list. Hence, we properly penalize for the differences between the two top-k lists. On the other hand, we do not penalize for the within-order of list differences by preserving the order in the former list. Because, the general assumption is that one does not have access to the order of the objects beyond k in top-k lists. Therefore, our incomplete rank correlation gives an upper correlation bound.

Figures 4a and 4b respectively show the incomplete rank correlations and the percentage of common ASes between IP spatial stress centrality and degree, customer-cone size and betweenness sampled at every 15 ASes.

In Figure 4a, the rank correlation between IP spatial stress and degree increases to 0.21 at top-60 list followed by a sharp decrease until 0.09 at top-270. In the same interval the ratio of common ASes (Figure 4b) decreases from 0.53 to 0.28 and fluctuates around 0.28. The behavior shows that the common ASes between IP spatial stress and degree quickly decays and causes

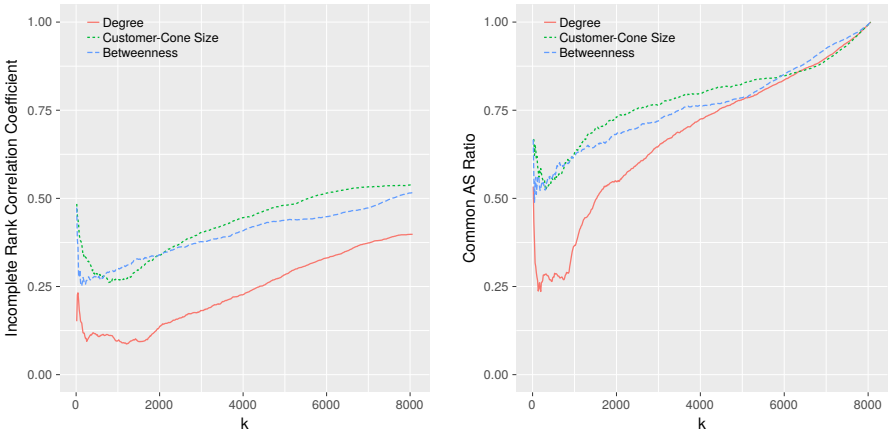

(a) Top-k incomplete rank correlation.

(b) Top-k common AS ratio.

Fig. 4: Comparison of top-k lists between IP spatial stress centrality and degree, customer-cone size and betweenness.

low correlation between the two centrality measures. The ratio of common ASes fluctuates roughly around 0.28 until top-885 list, then experiences a sharp increase until top-1800 list and then, the increase continues at a slower rate. However, the incomplete rank correlation unexpectedly gets lower between top-885 and top-1800 lists. Analyzing the dataset further shows that although the ratios of common ASes between top-885 and top-1800 increase, the ASes are ranked very differently in top-k lists of IP spatial stress and degree centralities. Therefore, their incomplete rank correlations do not increase along with the ratios of common ASes. The correlation and ratio of common ASes increase together after top-1800 list yet, the correlation only reaches to 0.39 for the whole dataset.

In Figure 4a, the rank correlations between IP spatial stress and customer-cone size decrease from 0.48 at top-15 to 0.26 at top-840. It stays roughly around 0.27 until top-1320 and increases thereafter. Figure 4b, on the other hand demonstrates a decrease in the ratio of common ASes from 0.57 at top-15 to 0.52 at top-315. The ratio monotonically increases after top-315. We have observed a similar behavior such that the ratio of common ASes increases between top-315 and top-1320, however the rank correlation continues to decrease because the ASes in those lists are ranked very differently with respect to IP spatial stress and customer-cone size.

Lastly, Figure 4a shows that the rank correlations between IP spatial stress and betweenness decrease from 0.47 at top-15 to 0.25 at top-135. After top-135, the correlations demonstrate an increasing trend up to 0.51 for the whole dataset. We observe a similar pattern in Figure 4b for the ratio of common ASes between IP spatial stress and betweenness centralities. That is, the correlation decreases/increases with respect to the ratio of common ASes in both top-k lists.

In summary, IP spatial stress centrality is significantly different from other measures, i.e., degree, customer-cone size and betweenness, regarding the whole dataset as well as the top-k critical AS lists. Unlike the other measures, IP spatial stress centrality incorporates business relations and IP address spaces to achieve a better measure for identifying critical ASes in the Internet.

6 Threats and Defense Mechanisms

In this section, we discuss the defense mechanisms against the cyber attacks targeting the routing infrastructure of the Internet through critical ASes rather than the attacks targeting the individual hosts in the Internet.

BGP speaking routers use TCP sessions to communicate the routing updates. Perpetrators can conduct well-known attacks including BGP message eavesdropping, modification, insertion, deletion and replay. In addition, BGP neither enforce any widely supported, strong authentication mechanisms nor does it strictly impose AS number, IP address prefix, route origination or AS path validation. Therefore, it is necessary to ensure the confidentiality, integrity and assurance of BGP messages in a scalable fashion [28]. An early work [29] proposed a general security mechanism by taking advantage of Public Key Infrastructures (PKIs) which allow routers to identify each other. Although the proposed mechanism induced high overhead on routers and suffered from scalability issues [30], it demonstrated that PKIs can play an important role in routing security. Secure BGP (S-BGP) [31] is a comprehensive routing security framework focusing on the BGP protocol. S-BGP uses two PKIs where the first one is employed for IP prefix attestation and the second one is for AS number attestation. This scheme requires the route attestation information via an attribute in BGP UPDATE messages. AS numbers and IP prefix ownerships in any routing update messages are authenticated through the PKIs. Also each AS on a path is required to include attestation information in BGP advertisements. Although S-BGP is comprehensive, it induces high overhead on BGP speaking routers. Secure Origin BGP (SoBGP) [32] reduces the overhead of S-BGP. Similar to S-BGP, SoBGP takes advantage of PKIs to authenticate and authorize ASes. It defines a new BGP message type, SECURITY, which delivers the necessary certificates to validate routes. Using the SECURITY messages, routers create a network topology map and validate the received BGP updates. Inter-domain Route Validation (IRV) [33] is a protocol for decentralized route security. The approach requires each AS to deploy an IRV server into its network. BGP speaking routers can query the IRV servers to validate the advertised routes. Secure Blockchain Trust Management (SBTM) [34] is a trust management systems to secure the inter-domain routing by taking advantage of blockchain-based PKI. An IP prefix is typically originated by a single AS in the Internet, because prefixes originated by more than one AS, Multiple Origin AS (MOAS), may imply prefix hijacking. Nevertheless, some large ISPs legitimately use MOAS for traffic engineering practices. Therefore, prefix origin authentication methods are employed to solve prefix ownership conflicts [35]. Prefix Hijacking Alert System (PHAS) [36] maintains a database

of routing information to identify IP prefix hijacking events. The routing information is collected from BGP monitoring projects such as Routeviews and RIPE and prefix ownership conflicts are reported to the ASes via email.

In addition to the attacks exploiting the BGP protocol, perpetrators can conduct DoS attacks and the variants [37, 38, 39] to flood the routers and links of the critical ASes. The defense mechanisms against the DoS attacks can be divided into three parts: attack detection, attack reaction and attack prevention. Attack detection mechanisms monitor or sample traffic to detect the DoS/DDoS attacks. MUltiLevel Tree for Online Packet Statistics (MULTOPS) detects bandwidth attacks by monitoring the packet rate between hosts in the Internet [40]. MULTOPS is based on the assumption that the packet rates between two machines remain proportional during regular operation hours. Hence, a dramatic increase in packet rates indicates the existence of a targeted attack. Another approach is based on the assumption that the set of the source IP addresses do not drastically change during the regular operation hours [41]. Hence, a drastic increase in terms of new IP addresses indicates the existence of a distributed, targeted attack. Attack reaction techniques involve resource management to mitigate the impact of DoS/DDoS attacks. Aggregate based Congestion Control (ACC) mechanism suggests monitoring and controlling high bandwidth aggregates at routers [42]. An aggregate corresponds to a collection of packets sharing a common property such as source address, destination address, protocol type or application type. The ACC mechanism identifies the aggregates causing congestions and rate limit the aggregates at the local or upstream routers. This method is effective not only for DDoS attacks but also for flash crowds. A short term but immediately effective solution requires deploying redundant network resources to absorb the rogue traffic during an attack. High profile service providers such as Microsoft and Yahoo dynamically increase service and network resources during attacks [43]. Although this type of reaction minimizes the impact of attacks in a timely fashion, it can still fail under persistent attacks. Attack prevention techniques aim to control targeted attacks before they reach to the victims. Ingress filtering [44] is an effective approach to drop the rogue traffic at the first AS. It requires each ISP checking the source IP addresses of the outgoing packets and filtering them if the source IP addresses do not belong to their IP address spaces, i.e., spoofed. ScoreForCore [45] is a statistical packet filtering mechanism to defend a victim site against DDoS attacks. In the proposed method, each packet's score is computed based on its attributes, including IP address, port number, packet length, TTL value and TCP flags. Then, the packets having a score below than a threshold value are discarded.

Attacks targeting critical ASes result in larger scale traffic disruptions in the Internet, which in turn may affect other critical services running on top of the Internet. On the other hand, there is no off-the-shelf defense mechanism that handles all types of threats. Therefore, it is necessary for critical ASes to deploy several defense mechanisms covering different types of attacks.

7 Conclusions

In this study we introduced IP address spatial path stress centrality as a measure to identify and group the critical ASes in the Internet. Evaluating the criticality of ASes not only guides adversaries to disrupt the Internet traffic with minimum resources but also provides network practitioners with insight on the pivotal ASes in the Internet. We define the criticality of an AS as the amount of potential traffic that it carries between pairs of ASes. Hence, the criticality of an AS is proportional to the number of the policy-preferred inter-AS paths passing through it as well as the traffic intensities of the paths.

Our empirical results show that the transit ASes in the observed Internet topology can be grouped into 28 levels of criticality such that 15 ASes form the set of the most critical ASes. These top-15 ASes frequently appear on the high intensity AS-to-AS paths. The most critical AS in the observed Internet topology is found to be $AS1299$ (Telianet) followed by $AS174$ (Cogent Communications), and $AS3356$ (Level 3 Communications). Nine of the top-15 ASes are tier-1 ASes whereas six of them are just transit ASes. Further investigation has shown that these six ASes are more critical than some of the tier-1 ASes, because they either have high numbers of peering ASes or they have many descendants with larger IP address spaces.

We compared the IP address spatial path stress centrality to the three well known and widely used centrality measures, namely customer-cone size, node degree and betweenness. Experimental results demonstrate significant difference in both complete and incomplete rank correlations among them. Because, the proposed IP spatial stress centrality incorporates business relations and IP address spaces to achieve a better measure for identifying critical ASes in the Internet.

References

1. Hawkinson J, Bates T (1996) Guidelines for creation, selection, and registration of an Autonomous System (AS). RFC 1930
2. Tozal ME (2016) The Internet: A system of interconnected autonomous systems. In: IEEE Systems Conference, Orlando, FL, USA
3. Luckie M, Huffaker B, Claffy K, Dhamdhere A, Giotsas V (2013) As relationships, customer cones, and validation. In: Internet Measurement Conference (IMC), Barcelona, ESP
4. Gao L (2001) On inferring autonomous system relationships in the internet. IEEE/ACM Transactions on Networking 9(6):733–745
5. Giotsas V, Luckie M, Huffaker B, Claffy K (2014) Inferring Complex AS Relationships. In: ACM IMC
6. CAIDA (2016) http://data.caida.org/datasets/as-relationships/serial-2/20160501.as-rel2.txt.bz2, 06/01/2016
7. Nur AY, Tozal ME (2016) Defending cyber-physical systems against dos attacks. In: IEEE International Conference on Smart Computing, St. Louis, MO, USA

8. Kang MS, Lee SB, Gligor VD (2013) The crossfire attack. In: IEEE Symposium on Security and Privacy, San Francisco, CA, USA
9. Bellovin S, Gansner ER (2004) Using link cuts to attack internet routing. Tech. rep., ATT Research
10. Schuchard M, Mohaisen A, Kune DF, Hopper N, Kim Y, Vasserman EY (2010) Losing control of the internet: Using the data plane to attack the control plane. In: ACM Conference on Computer and Communications Security, Chicago, IL, USA
11. Butler K, Farley TR, McDaniel P, Rexford J (2010) A survey of bgp security issues and solutions. Proceedings of the IEEE 98(1):100–122
12. Tozal ME (2017) Autonomous system ranking by topological characteristics: A comparative study. In: IEEE Systems Conference, Montreal, CAN
13. Tozal ME (2018) Policy-preferred paths in AS-level Internet topology graphs. Theory and Applications of Graphs 5(1):1–32
14. Tozal ME (2016) Enumerating single destination, policy-preferred paths in AS-level Internet topology maps. In: IEEE Sarnoff Symposium, NJ, USA
15. Rueda DF, Calle E, Marzo JL (2017) Robustness comparison of 15 real telecommunication networks: Structural and centrality measurements. Journal of Network and Systems Management 25(2):269–289
16. Wang Y, Zhang K (2016) Quantifying the flattening of internet topology. In: International Conference on Future Internet Technologies
17. Latora V, Nicosia V, Russo G (2017) Complex Networks: Principles, Methods and Applications. Cambridge University Press
18. Zimmerli L, Tellenbach B, Wagner A, Plattner B (2009) Rating autonomous systems. In: Internet Monitoring and Protection (ICIMP)
19. Clérot F, Nguyen Q (2005) A social network approach for the ranking of the autonomous systems of the internet. In: Link Analysis Workshop
20. Wagner C, François J, State R, Dulaunoy A, Engel T, Massen G (2013) Asmatra: Ranking ass providing transit service to malware hosters. In: Integrated Network Management, IFIP/IEEE
21. Dimitropoulos X, Krioukov D, Riley G, claffy K (2006) Revealing the autonomous system taxonomy: The machine learning approach. In: Passive and Active Network Measurement Workshop (PAM), Adelaide, Australia
22. Rekhter Y, Li T, Hares S (2006) A Border Gateway Protocol 4 (BGP-4). RFC 4271 (Draft Standard)
23. Wang H, Song M (2011) Ckmeans.1d.dp: optimal k-means clustering in one dimension by dynamic programming. The R Journal 3(2):29–33
24. Freeman LC (1977) A Set of Measures of Centrality Based on Betweenness. Sociometry 40(1):35–41
25. Shavitt Y, Weinsberg U (2012) Topological trends of internet content providers. In: SIMPLEX
26. Emond EJ, Mason DW (2002) A new rank correlation coefficient with application to the consensus ranking problem. Journal of Multi-Criteria Decision Analysis 11(1):17–28
27. Masoud MZ, Hei X, Cheng W (2013) A graph-theoretic study of the flattening internet as topology. In: IEEE International Conference on Net-

works, Singapore, Singapore
28. Hiran R, Carlsson N, Shahmehri N (2016) Does scale, size, and locality matter? evaluation of collaborative bgp security mechanisms. In: IFIP Networking Conference, IEEE, pp 261–269
29. Perlman RJ (1988) Network layer protocols with byzantine robustness. PhD thesis, Massachusetts Institute of Technology
30. Nicholes MO, Mukherjee B (2009) A survey of security techniques for the border gateway protocol (bgp). IEEE communications surveys & tutorials 11(1):52–65
31. Kent S, Lynn C, Seo K (2000) Secure border gateway protocol (s-bgp). IEEE Journal on Selected areas in Communications 18(4):582–592
32. Ng J, et al (2004) Extensions to bgp to support secure origin bgp (sobgp). Tech. rep., Internet Draft, Apr
33. Goodell G, Aiello W, Griffin T, Ioannidis J, McDaniel PD, Rubin AD (2003) Working around bgp: An incremental approach to improving security and accuracy in interdomain routing. In: NDSS, vol 23, p 156
34. Gómez-Arevalillo ADLR, Papadimitratos P (2017) Blockchain-based public key infrastructure for inter-domain secure routing. In: International Workshop on Open Problems in Network Security (iNetSec), pp 20–38
35. Qiu SY, Monrose F, Terzis A, McDaniel PD (2006) Efficient techniques for detecting false origin advertisements in inter-domain routing. In: Secure Network Protocols, 2006. 2nd IEEE Workshop on, IEEE, pp 12–19
36. Lad M, Massey D, Pei D, Wu Y, Zhang B, Zhang L (2006) Phas: A prefix hijack alert system. In: USENIX Security symposium, vol 1, p 3
37. Nur AY, Tozal ME (2018) Record route ip traceback: Combating dos attacks and the variants. Computers & Security 72:13–25
38. Kalkan K, Gür G, Alagöz F (2017) Filtering-based defense mechanisms against ddos attacks: A survey. IEEE Systems Journal 11(4):2761–2773
39. Wisthoff M (2018) Ddos countermeasures. In: Information Technology-New Generations, Springer, pp 915–919
40. Gil TM, Poletto M (2001) Multops: A data-structure for bandwidth attack detection. In: USENIX Security Symposium, pp 23–38
41. Peng T, Leckie C, Ramamohanarao K (2004) Proactively detecting distributed denial of service attacks using source ip address monitoring. In: International Conference on Research in Networking, Springer, pp 771–782
42. Mahajan R, Bellovin SM, Floyd S, Ioannidis J, Paxson V, Shenker S (2002) Controlling high bandwidth aggregates in the network. ACM SIGCOMM Computer Communication Review 32(3):62–73
43. Peng T, Leckie C, Ramamohanarao K (2007) Survey of network-based defense mechanisms countering the dos and ddos problems. ACM Computing Surveys (CSUR) 39(1):3
44. Baker F, Savola P (2004) Ingress Filtering for Multihomed Networks. RFC 3704
45. Kalkan K, Alagöz F (2016) A distributed filtering mechanism against ddos attacks: Scoreforcore. Computer Networks 108:199–209