

Understanding Transit-Edge Routing Separation: Analysis and Characterization

Kunpeng Liu^a, Bijan Jabbari^a, Stefano Secci^b

^aCommunications and Networking Lab., George Mason University, VA. ^bLIP6, Pierre & Marie Curie University, France.

E-mail: {kliu3,bjabbari}@gmu.edu, stefano.secci@lip6.fr

Abstract—With the significant increase in the Internet traffic, the uncoordinated routing practices across border gateways are limiting the Internet growth. It is becoming urgent to rethink the principles underlying the Internet infrastructure as well as the design of its major protocols, especially those related to Internet routing and traffic engineering. In its support, an appropriate characterization of the current Internet properties seems necessary as it may provide valuable information for the design of future Internet protocols. In this paper, we analyze Internet routing maps over a two-year period within a Transit-Edge (T-E) routing separation perspective, a promising direction to improve Internet resiliency and security by allowing explicit forwarding through routing locators on the way toward the destination network. We focus on the characterization of the behavior of edge and transit Autonomous Systems (ASes) in terms of interconnection, routing and traffic engineering practices, highlighting similarities and differences. We show in particular that edge networks significantly perform incoming traffic engineering and that this trend is increasing in time¹.

I. INTRODUCTION

The Internet has been evolving from an academic network managed and operated by researchers, to a worldwide and ubiquitous network interconnecting devices of multiple natures. At its inception, many technology choices had to be taken, such as on the forwarding nature of the Internet Protocol, its addressing and the inter-domain routing principle. The history tells us that the Internet Protocol (IP) relies on packet switching with statistical multiplexing, that its addressing is based on a 32-bit space and is now migrating to a 128-bit space, and that the Border Gateway Protocol (BGP) [1] is the inter-domain routing protocol used by Autonomous Systems (AS) to exchange routing information. BGP relies on a flat routing mode using path vectors for each IP network prefix, announced independently in an uncoordinated fashion.

The lack of coordination amongst AS networks appears strategically reasonable as each AS needs to follow first its own interests and objectives. However, the flat routing mode of Internet routing is unable to scale with such a behavior for a very large number of networks. Meanwhile, the number of ASes as well as the announced network prefixes are increasing extremely fast (currently, about 39,000 ASes and 400,000 network prefixes). Such a large and increasing number of prefixes, even if dictated by reasonable traffic engineering and multi-homing practices, are posing many issues from a network

management viewpoint. Coupled with other aspects such as BGP routing convergence, instability and weak resiliency, they are undermining the healthy development of the Internet.

A direction recently evaluated to tackle the Internet routing scalability and resiliency issue is to adopt transit-edge (T-E) routing separation schemes [2]. Allowing a two-level hierarchy routing between edge and transit networks, it is possible to reduce the transit routing table sizes since a very large majority of the Internet networks are at the edges and do not transit traffic. Moreover, novel traffic engineering capabilities can be introduced. In this paper, we measure the Internet topology from a T-E routing separation perspective. By analyzing the recent BGP tables over a two-year period, we aim at characterizing the properties of edge and transit networks from interconnection, routing and traffic engineering perspectives. We first analyze the interconnection degrees, the AS path prepending and IP prefix de-aggregation behaviors, for edge and transit ASes. Last but not least, we measure and characterize the routing stability phenomenon. Our analysis shows that edge ASes do perform actively Internet traffic engineering almost as much as transit ASes do, and that this trend is increasing in time, which suggests that edge ASes would benefit from novel T-E separation protocols.

The paper is organized as follows. Section II describes the technical background of transit and edge networks. Section III and IV analyze the T-E separation characteristics from interconnection and routing perspectives, respectively. Section V summarizes the paper with final conclusions.

II. TRANSIT AND EDGE NETWORKS

The Internet interconnection graph can be partially inferred via BGP routing tables. Routeviews' public routing tables [3] aggregate the daily view of multiple backbone routers, which represents a very detailed mirror on the Internet ecosystem evolution. After a rapid analysis, we find that at present around 84% of the total ASes act as pure destination networks, only appearing at the last position of the AS paths. They are commonly considered as "stub ASes". In practice, some large stub ASes (content providers and delivery networks) functionally fragment their networks into multiple ASes for management reasons, and they may also appear in the penultimate or in the third from last position in AS paths. Nearly 13% additional ASes appear up to the third from last position of BGP AS paths, among which are certainly also some regional Internet Service Providers (ISPs). The sub-network composed

¹This work was funded by the ONR US project "Secure Protocols and Services for Resilient Internetworking".

of these 97% ASes can be treated as the edge of the Internet, which given its interconnection behavior has different traffic engineering requirements and routing purposes than transit networks. In fact, the remaining 3% ASes do transit the global Internet traffic as their principal purpose, and they can be treated as the transit part of the Internet. As of our observation, these transit and edge network ratios have been rather stable even though the Internet has significantly grown.

The T-E routing separation paradigm suggests to insert routing locators at the frontier between transit and edge networks. Different protocols can be conceived to manage identifier-to-locator mappings and to encapsulate or aggregating (tunneling) packets in the transit sub-path. One working at layer 3 only is the Locator-Identifier separation protocol (LISP) [4], currently under standardization (it somehow supersedes host-based approaches such as SHIM6 [5] or HIP [6] that appear as less scalable mechanisms). Besides allowing a very important reduction of the Internet routing table, as discussed in [7], T-E separation can lead to important improvements in terms of routing resiliency. Indeed, the introduction of many routing locators for the same destination drastically increases the Internet path diversity. If adequately managed for traffic engineering, the enlarged path diversity can lead to significant improvements of the Internet resiliency, as explained in [2] where a framework for coordinated edge-to-edge load-balancing and Internet-wide multipath routing is presented.

Therefore, new tools for Internet traffic engineering - currently limited to BGP tweaking practices such as prefix de-aggregation and transient announcements that are increasing the routing table size and are decreasing the Internet service reliability - could arise from T-E separation. At present, the potential achievable performance improvements are attracting attention from content providers and content delivery edge networks, especially with the emergence of Cloud Computing applications that require high connection resiliency and persistent reachability [2]. In the following, we characterize current interconnection, routing and traffic engineering practices of edge and transit ASes via measurement of BGP routing tables.

III. INTERCONNECTION TOPOLOGY ANALYSIS

BGP Routeviews' routing tables are captured from ASes that peer with many large transit carriers, so they represent a transit view on the Internet routes. Meanwhile, the AS interconnection information from the directional perspective of edge ASes is difficult to get. Therefore, it appears appropriate to represent routing maps using an undirected graph. Studying the undirected graph, we can characterize the degree distributions of edge and transit ASes, and analyze the interconnection properties of T-E separation.

A. Degree analysis

The AS degree, defined as the number of AS neighbors, somehow reflects the importance of an AS. In Fig. 1 we plot the complementary cumulative distribution function (CCDF) of the AS degree for edge and transit ASes.

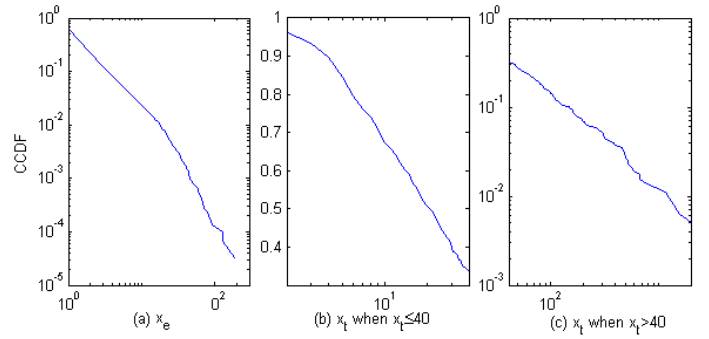


Fig. 1. The degree CCDF of edge and transit ASes

Let x_e and x_t denote the degree of edge and transit ASes, respectively. The CCDFs in Fig. 1 are obtained by analyzing the routing tables of Jan. 2009, but the same profile is approximately maintained for successive routing tables. Note that Fig. 1(a) and Fig. 1(c) use a log-log scale, while Fig. 1(b) uses a log-linear scale. We can see that the x_e CCDF linearly decreases in a log-log scale, and so does the x_t CCDF when the degree is bigger than a relative large threshold, e.g., 40. When x_t is smaller than the threshold, the CCDF decreases almost linearly in a log-linear scale. It is worth recalling that the CCDF of a nonnegative random variable that follows truncated discrete power law distribution can be calculated as $F_c(x) \sim ax^{-\alpha}$, while the CCDF of a random variable that has truncated probability density function (pdf) as $f(x) = b/x$ can be calculated as $F_c(x) \sim -b \ln(x)$. In the following, we define the distribution with PDF $f(x) = b/x$ as inverse distribution; note that the CCDF of power law distribution becomes to linear function in a log-log scale, while that of inverse distribution shows linear characteristic in a log-linear scale. When combining the above results, we find that:

- The degree of edge ASes can be well fit with a power law distribution.
- When the degree of a transit AS is relatively small, it follows a truncated inverse distribution.
- When the degree of a transit AS is larger than a certain threshold, it follows a power law distribution.

To simplify the following analysis, we treat x_e and x_t as continuous random variables. Let the CCDFs for the degree of edge and transit ASes be F_{ce} and F_{ct} , respectively. We investigate the following relations:

$$F_{ce}(x_e) \sim a_e x_e^{-\alpha_e} \quad (1)$$

$$F_{ct}(x_t | 2 \leq x_t \leq d) \sim -b \ln(x_t) \quad (2)$$

$$F_{ct}(x_t | x_t > d) \sim a_t x_t^{-\alpha_t} \quad (3)$$

Please note that in (1) and (3) the CCDFs have right hand side cutoffs C_e and C_t , respectively. From (2), we find $f_t(x_t | 2 \leq x_t \leq d) \sim b/x$. As $\int_2^d f_t(x_t | 2 \leq x_t \leq d) dx = 1$, we get:

$$b \sim \ln^{-1}\left(\frac{d}{2}\right) \quad (4)$$

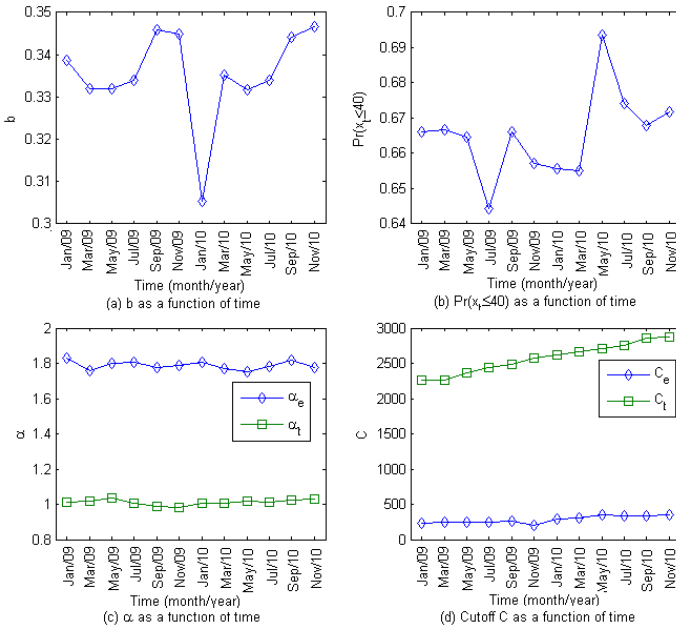


Fig. 2. Model parameters as a function of time

Hence, as long as d is a constant, b as well as the statistics of x_t given $2 \leq x_t \leq d$ will also be deterministic. Through a similar derivation, the relationship between a and α can also be found.

In order to inspect the parameter trends, we choose $d = 40$, and apply the least square error (LSE) as the model estimator to the two-year period routing tables. We first examine the trend of b to validate our previous analysis. From (4), we know that b should be around 0.33 given $d = 40$. The theoretical analysis perfectly fits our measurements reported in Fig. 2(a).

Next, we are interested in the trends of $Pr(x_t \leq 40)$, α_e , α_t , as well as the cutoffs C_e and C_t . In Fig. 2(b), we find that $Pr(x_t \leq 40)$ is quite stable; this indicates the probability that the transit AS degree follows the power law distribution or inverse distribution is quite stable. Fig. 2(c) shows that α_e is larger than 1.5 and smaller than 2, while α_t is very close to 1. Fig. 2(d) shows that the cutoff of x_t is much larger than that of x_e , and C_t as well as C_e show a clear increasing trend during the observation period. Before further analyzing the results, let us discuss the properties of truncated power law distribution with PDF $f(x) \sim r x^{-\alpha-1}$ and two cutoffs c_1 and c_2 (c_1 is the left hand side cutoff, and c_2 is the right hand side cutoff). When $c_2 \gg c_1$ and $c_1 = 1$ or 2, it is easy to show that:

$$E(x) \sim r \frac{c_2^{1-\alpha} - c_1^{1-\alpha}}{1-\alpha} \quad (5)$$

$$E(x^2) \sim r \frac{c_2^{2-\alpha} - c_1^{2-\alpha}}{2-\alpha} \quad (6)$$

When $\alpha \approx 1$, based on (5), we can get the equation

$$\lim_{\alpha \rightarrow 1} E(x) \sim r \ln(c_2) \quad (7)$$




Fig. 3. The roles immutability of ASes as a function of time difference

Combining the observations, we can assert that:

- The x_t average is increasing in the two years, as α_t is very closer to 1 and the cutoff C_t is always raising. This shows the interconnection of transit ASes evolves permanently, by which a lot of new shortest paths can be created to improve the performance of the Internet.
- Following the raise of cutoff C_e , the x_e average is also increasing in the two years. This reflects the fact that more and more edge networks perform upstream multi-homing to improve the network interconnection situation.
- Based on (5)-(7) and simple calculations, we can find that the standard deviations of x_e and x_t are also increasing in last years. This indicates that the distributions for the degree of edge and transit ASes are stretching constantly.

B. T-E Separation Properties

According to the position of each AS in the routing entries, the Internet can be artificially separated into edge and transit networks; obviously, an AS should hold either an edge or a transit role. However, the role of a particular AS may change abruptly, due to interconnection evolution or routing fluctuations; this phenomenon is shown in Fig. 3 (filtering out path prepending). The horizontal axis represents the time difference, and the vertical axis represents the percentage of a kind of ASes that still hold their original ranking after the time interval (defined as AS role immutability). From Fig. 3, we can see the immutability of edge networks drops almost linearly from 98% to 90% when the time difference increase from 2 months to 22 months, while at the same time the immutability of transit networks drops in a more dramatic way from 81% to 59%. Given these observations, we can state that:

- The roles of ASes are quite immutable in a short relative period, like 1 or 2 months.
- Not only the immutability of edge ASes is higher than that of transit ASes, but the role change rate of edge ASes is also much smaller than that of transit ASes.
- T-E separation should not rely on an automated detection of current roles, but should be set statically by transit ASes with little or no coordination with edge ASes.

Such role changes indicate that edge ASes rarely “evolve” as transit ones, but rather the inverse occurs, i.e., ASes in the transit core are pushed towards the edges as the time passes.




Fig. 4. AS node path prepping probability as a function of time




Fig. 5. AS link path prepping probability as a function of time

IV. ROUTING AND TRAFFIC ENGINEERING ANALYSIS

In this section, we characterize edge and transit networks from a routing and traffic engineering point of view. Among all the available traffic engineering techniques in BGP routing, we can mention local preferences for outbound traffic engineering, AS path prepping for inbound traffic engineering, and IP de-aggregation for multi-homing traffic engineering. While the first cannot be inferred with adequate precision from routing table analysis, path prepping and IP de-aggregation can, as reported in the following. Such practices coupled with the BGP convergence issue indirectly affects the BGP routing instability, which is an aspect also analyzed in this section.

A. AS path prepping analysis

With AS path prepping, artificially repeating its own AS number to increase the length of certain AS paths passing through it, an AS can meet inbound traffic engineering goals, i.e., distracting incoming traffic toward more available or preferred entry points. We are interested in the occurrence of path prepping, including the empirical probability for an AS to apply path prepping, as well as for an AS link to be affected by path prepping. We categorize the AS links into three types: links between edge networks, links between edge and transit networks and links between transit networks. Fig. 4 shows the experimental probabilities that edge and transit ASes use path prepping, while Fig. 5 shows the probabilities that the three types of AS links are affected by path prepping. In Fig. 4, we find that not only are the probabilities to employ AS path prepping very close to each other, they but also share the same time profile. In Fig. 5, we see that the AS links

between transit networks are affected by path prepping with the highest probability while the links between edge networks are with the lowest probability. All in all, we can assert that:

- The path prepping occurrence for edge and transit ASes is relatively low, as it is below 0.1 for both.
- The occurrence probabilities are very similar with each other.
- The transit networks perform inbound traffic engineering more frequently than edge ASes.

Edge ASes apply path prepping essentially for inbound load balancing, while transit ASes perform path prepping as a second-level routing rule for provider transit vs. client transit and transit links vs. peering links load-balancing (the first-level rule for such operations typically is the local-preference).

B. Prefix de-aggregation impairment analysis

For security, resiliency as well as load balancing purposes, ASes can artificially fragment large IP prefixes into several smaller prefixes and announce them separately [9], [10]. This behavior is usually known as IP prefix de-aggregation. We analyze the impairment of IP prefix de-aggregation at time t to BGP routing tables in the following way: first, we gather all the IP prefixes announced by a given AS x at time t , noting the total number of prefixes as d_{xt} ; next, we recursively apply a seamless and precise IP aggregating rule to obtain the size of the IP prefixes before IP de-aggregation, which is noted as a_{xt} ; then the IP de-aggregation rate r_{xt} of the AS x can be expressed as:

$$r_{xt} = \frac{d_{xt} - a_{xt}}{a_{xt}} \quad (8)$$

For example, an AS announces 1.2.3.128/25, 1.2.3.0/25 and 128.1.1.0/24, separately; as 1.2.3.128/25 and 1.2.3.0/25 can be aggregated with 1.2.3.0/24, the de-aggregation rate of the AS is $(3-2)/2 = 0.5$. Therefore, any AS that does not employ IP de-aggregation should have a zero IP de-aggregation rate.

Fixing the total number of ASes to N , an AS that can communicate with every announced prefix at time t should have a BGP routing table size close to $\sum_{i=1}^N (a_{it}r_{it} + a_{it}) = \sum_{i=1}^N a_{it}r_{it} + \sum_{i=1}^N a_{it}$. Nevertheless, in an ideal scenario, if there is no IP prefix de-aggregation, its BGP routing table size should only be $\sum_{i=1}^N a_{it}$. Due to IP prefix de-aggregation, the routing table size gets indeed significantly enlarged.

If we consider the overall impact of IP de-aggregation to the sizes of routing tables and let R_t be the impact ratio of the routing tables, then:

$$R_t = \frac{\sum_{i=1}^N a_{it}r_{it}}{\sum_{i=1}^N a_{it}} \quad (9)$$

where, $i \in [1, N]$, a_{it} are unknown constants and r_{it} can be treated as independent identically distributed (IID) random variables due to the partial arbitrary nature of IP prefix de-aggregation. From (9), we know that:

$$E(R_t) = E(r_{it}) \quad (10)$$

Therefore, if we could find an alternative routing mode with some form of hierarchical routing more natively supporting

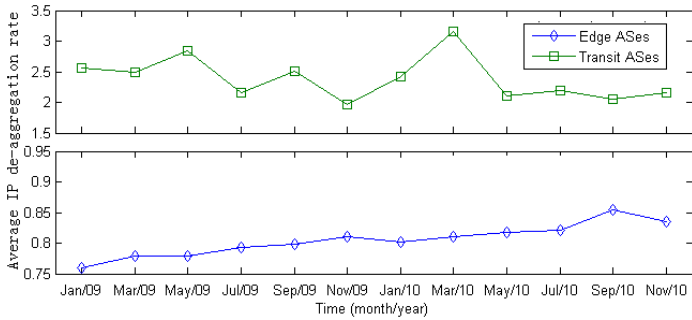


Fig. 6. Average AS prefix de-aggregation rates as a function of time

IP prefix de-aggregation – such as a T-E routing separation protocol – while allowing at least the same level of traffic engineering capabilities, the BGP routing table size could shrink dramatically.

The average prefix de-aggregation rates for edge and transit ASes are shown in Fig. 6. We find that for edge ASes it increases quite clearly in time, while for transit ASes it oscillates and slightly decreases in time. The overall IP prefix de-aggregation rate, mainly depending on edge ASes, has grown from 0.81 to 0.87 in the two-year period, which further stresses the Internet scalability (higher impact on routing tables). From the studies, we can assert that:

- Transit ASes are more used to prefix de-aggregation than edge ASes, which is roughly 3-times more often than edge ASes, and its de-aggregation usage can vary significantly in time and not necessarily increases, while edge ASes usage de-aggregation raises constantly.
- The IP de-aggregation rates of edge and transit ASes directly impair the scalability and efficiency of the Internet, and the average value of impact ratio R is affected by the de-aggregation rate r_{it} (10).
- Following the growth of the overall prefix de-aggregation rate, the impairment of prefix de-aggregation also increases in these two years.

All in all, it is worth stressing that one would expect that edge networks do not perform actively traffic engineering because of the much lower scale of bitrate aggregates than for transit networks. However, we have verified that not only they actively do incoming and multihoming traffic engineering (via BGP path prepending and prefix de-aggregation), but that they do that at a close level to the level at which transit networks do. Moreover, it appears that the this trend is increasing in time (prefix de-aggregation).

C. Routing Instability Analysis

Internet routing instability represents the fluctuation of routing information towards network reachability. Many reasons are behind this phenomenon, including the change of infrastructure, the impact of traffic engineering, the employment of multi-homing, etc. However, high levels of routing instability can lead to serious impairments, e.g., packet loss, increase of network latency and time to convergence, and even the




Fig. 7. AS interconnection diagrams




Fig. 8. Routing instability as a function of the time difference

loss of interconnection availability in wide-area or national networks [11].

In inter-domain routing, routing instabilities can be roughly characterized from the fluctuation of the BGP routing table. In the following, we define the appearance time of an AS-level link i in a routing table as the occurrence count of the link, also define the average of the overall change rate as the routing instability rate, noted as RI . We consider RI as an adequate metric to quantify the routing instability. If we represent an undirected graph at time t with $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$, where \mathcal{V}_t is the set of the nodes and \mathcal{E}_t is the set of links, the RI after time τ can be calculated as follows:

$$RI = \frac{1}{|\mathcal{E}_t|} \sum_{i \in \mathcal{E}_t} \frac{|n_i^t - n_i^{t+\tau}|}{\max(n_i^t, n_i^{t+\tau})} \quad (11)$$

where, $|\mathcal{E}_t|$ is the size of the link set, n_i^t is the occurrence count of link i in the routing table at time t , and $n_i^{t+\tau}$ is the occurrence count of link i in the routing table at time $t + \tau$. If link i cannot be found in the routing table at time $t + \tau$, we set $n_i^{t+\tau} = 0$.

A demonstration of how to use (11) is shown here. Suppose we want to calculate the RI between Fig. 7(a) and Fig. 7(b), then $RI = 1/3 * (|5 - 1|/5 + (3 - 0)/3 + |1 - 5|/5) \simeq 0.87$. As there is considerable difference between Fig. 7(a) and Fig. 7(b), we get a very big RI , which represents the routing instability between the two graphs is in a significantly high degree.

We now artificially partition the AS graph into three layers: edge networks, transit networks, and intermediate networks connecting edge and transit ASes. Then we use (11) to measure the routing instability status of these three networks, which are shown in Fig. 8 and Fig. 9. In Fig. 8, the horizontal




Fig. 9. Routing instability rates as a function of the time

axis is the time difference τ and the vertical axis is the routing instability given the time difference τ . In Fig. 9, the horizontal axis is the time t , and the vertical axis is the routing instability between the routing table at time $t - \tau$ and the routing table at time t on a fixed time different $\tau = 2$ months. We find that the routing instabilities of the three networks all raise gradually in a similar way when the time difference increases. When the time difference is fixed at two months, the routing instabilities of the three networks also vary with a similar pattern.

From the two figures, we can assert that:

- The routing instability of the three network layers is similar, and for all it raises as long as the time difference increases.
- Among the three network layers, routing at intermediate networks is the most stable, while at the transit networks it is the least stable.
- When the time difference is fixed at two months, the routing instabilities of the three network layers also share the similar pattern as time changes.
- The routing instability phenomenon is relatively serious presently, as the minimum value in the two figures is still around 0.2.

Two main factors can be behind such a routing instability: the inner convergence and oscillation problems of BGP, and the incentive of edge and transit networks in performing inbound and outbound traffic engineering operations.

V. CONCLUSION

Transit-edge routing separation functionally proposes to create a two-level hierarchical routing between networks that have different routing behavior. In this paper, we measure real inter-domain routing information to characterize behavior and properties of edge and transit AS networks with a transit-edge routing separation perspective.

From an interconnection viewpoint, we found that the interconnection degree of an edge AS can be well fit with a truncated power law distribution, while that of a transit AS can be fit by the combination of power law and inverse distribution, and we identified the different regimes of edge AS and transit AS degree distributions. From a routing and traffic engineering viewpoint, we discovered that edge and transit ASes have similar probabilities of applying AS path prepending. We categorized the AS links into three types,

and unraveled that they are affected by path prepending with different probabilities. We recognized that the impact ratios of BGP routing tables are directly determined by the IP prefix de-aggregation rate of edge and transit ASes. Moreover, we described a mechanism to measure the routing instability phenomenon, recognizing that the transit networks have the largest routing instability while the intermediate networks have the least routing instability.

From a traffic engineering requirement perspective, one would expect that edge networks do not actively perform traffic engineering because of the much lower scale of bitrate aggregates than for transit networks. However, we have verified that not only edge ASes actively do traffic engineering, but that they do it at a close level to the one at which transit networks do. Moreover, it appears that the multihoming traffic engineering trend, based on prefix de-aggregation, is a practice increasingly adopted by edge ASes.

REFERENCES

- [1] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," *RFC 4271*, 2006.
- [2] S. Secci, K. Liu, K. Rao, and B. Jabbari, "Resilient Traffic Engineering in a Transit-Edge Separated Internet Routing," in *Proc. of IEEE ICC*, 2011.
- [3] D. Meyer, "University of oregon route views archive project," at <http://archive.routeviews.org>.
- [4] D. Farinacci, V. Fuller, D. Oran, D. Meyer, and S. Brim, "Locator/ID separation protocol (LISP)," *draftfarinacci-LISP-03*, 2007.
- [5] E. Nordmark and M. Bagnulo, "Shim6: Level 3 multihoming shim protocol for IPv6," *RFC 5533*, 2009.
- [6] R. Moskowitz and P. Nikander, "Host identity protocol (HIP) architecture," *RFC 4423*, 2006.
- [7] Y. Wang, J. Bi, and J. Wu, "Empirical analysis of core-edge separation by decomposing Internet topology graph," in *Proc. of IEEE GLOBECOM*, 2010.
- [8] D. Cohen, "All the world's a net," *New Scientist*, vol. 174, no. 2338, pp. 24–29, 2002.
- [9] X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, and L. Zhang, "IPv4 address allocation and the BGP routing table evolution," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 1, pp. 71–80, 2005.
- [10] G. Molla, G. Castro, J. Baliosian, M. Bruin, and M. Yannuzzi, "Understanding IPv4 prefix de-aggregation: challenges for routing scalability," 2010.
- [11] C. Labovitz, G. Malan, and F. Jahanian, "Internet routing instability," *Networking, IEEE/ACM Transactions on*, vol. 6, no. 5, pp. 515–528, 1998.