Measuring the Internet Routing Scalability from the Perspective of Address Allocation

Dan Wu*†, Zhiliang Wang*‡, Xia Yin*†, Xingang Shi*‡, Jianping Wu*†
* Tsinghua National Laboratory for Information Science and Technology (TNList)
† Department of Computer Science & Technology, Tsinghua University
‡ Institute for Network Sciences and Cyberspace, Tsinghua University, Beijing, P.R. China, 100084
Email: {wudanzy, wzl, yxia, shixg, jianping}@csnet1.cs.tsinghua.edu.cn

***Abstract*—The internet faces severe routing scalability problem - the rapid increasing of prefix and updates. There are two classes of factors that contribute to scalability problem: inappropriate allocation of IP addresses and inappropriate usage of them. Existing works mainly focus on inappropriate usage of allocations, so this work focuses on inappropriate allocation of IP addresses. We carry out the measurement from two perspectives. From the perspective of each AS, we classify preprocessed prefixes according to their contribution to aggregating addresses inside the AS. We find that factors due to inappropriate allocation of IP address contribute more than that of inappropriate usage of them in values. From the perspective of Internet, we propose a new metric Brother Prefix Distance, which implies the likelihood of aggregating addresses of different ASes. We find that Brother Prefix Distance increase over time, which implies that impact of address allocation on routing scalability increases over time.**

***Index Terms*—scalability; network measurement**

I. Introduction

Nowadays, one of the major concerns about Internet is its scalability problems, including the rapid increasing of BGP routing table size and routing updates. Those factors that contributes to rapid increasing of BGP routing table can be classified into two groups: factors related to allocation of IP addresses and factors related to usage of IP addresses. The latter group attracts considerable attention and has been fully studied, but the former group, which contains factors related to allocation of IP addresses, has not been studied in detail. We argue that the allocation of IP addresses, as the foundation of their usage, impacts the scalability of Internet deeply.

In Internet practice, a feasible and effective way of improving scalability is to leverage the aggregation of IP addresses [1]. Aggregation can be used in routing table or forwarding table, which combines multiple entries into a new one and uses the new entry instead of old ones. Therefore, a good allocation scheme should facilitate the aggregation of IP addresses. On the one hand, it should facilitate the aggregation of IP addresses in the same AS. On the other hand, it should also facilitate the aggregation of IP addresses in different ASes.

In this paper, we measure the scalability of Internet and focus on factors due to address allocation. Our measurement is carried out from the following two perspective. First, from the viewpoint of AS, we measure the likelihood of aggregation of addresses allocated to each AS. Second, from the viewpoint of whole Internet, we propose a new metric named brother prefix distance, which implies the likelihood of aggregating a prefix in remote routers.

II. From the viewpoint of ASes

From the viewpoint of ASes, prefixes are classified according to its contribution towards aggregation inside an AS. We simulate the process of allocating prefixes to an AS. First, all prefixes of the same origin AS is sorted by its appearance time, which is the date the prefix first appears in a BGP table. Then, we allocate each prefix to the origin AS in the order of appearance. Each prefix is aggregated with already allocated ones if possible. All prefixes can be classified into four classes:

Continuous Fragments are more than two prefixes that are allocated at the same time and continuous but cannot be aggregated, which are usually due to the fact that the range of IP addresses in an allocation can not be aggregated into a single prefix.

Discontinuous fragments are prefixes that are not continuous with other fragments, which are due to the lack of address reservation for an AS. When the AS requires another block of IP addresses, addresses continuous with existing blocks are already allocated to other ASes.

Aggregated non-fragments are prefixes that are continuous with previous fragments and can be aggregated. Contrary to discontinuous fragment, an aggregated non-fragment is an instance of successful address reservation for an AS. Aggregated non-fragments are advertised into routing system because of inappropriate usage of allocated addresses.

Absorbed non-fragments are prefixes that are covered by previous fragments, so they can be absorbed. They are usually caused by traffic engineering issues and are harmful to the routing scalability.

Figures 1a and 1b show the classification of IPv4 and IPv6. Of the four classes of prefixes, the number of discontinuous fragments is the most, next is absorbed non-fragments. In all prefixes, more than a half of them are fragments (continuous and discontinuous), which implies that factors due to address allocation are still main contributors to routing scalability. Moreover, IPv6 is still in its early years, the proportion of fragments is much higher than that of IPv4. In IPv4, about 20% of fragments are continuous fragments and about 80% of non-fragments are absorbed fragments which is caused by traffic engineering. For the changes over time, the proportion

978-1-4673-8590-9/15/$31.00 ©2015 IEEE

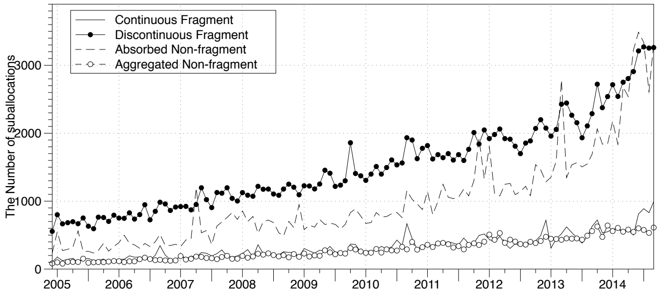

(a) The classification of IPv4 prefixes

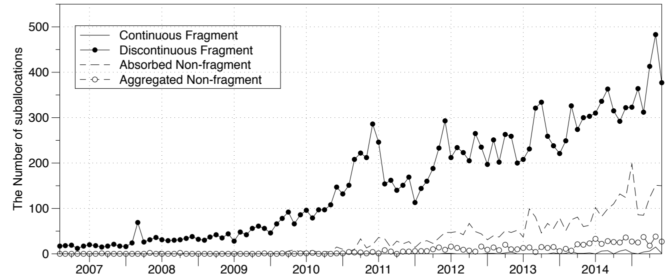

(b) The classification of IPv6 prefixes

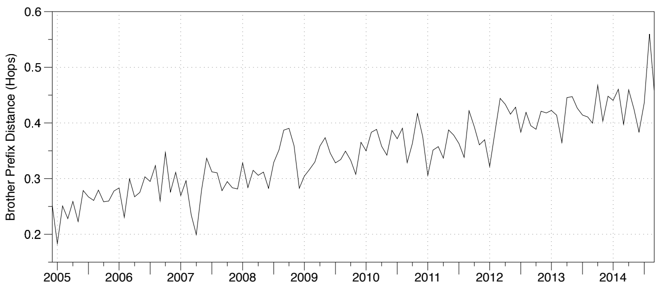

(c) IPv4 Brother Prefix Distance

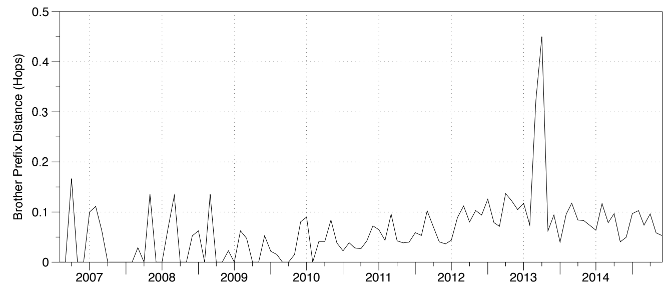

(d) IPv6 Brother Prefix Distance

Fig. 1. Analysis results of Prefix Classification and Brother Prefix Distance

of fragments decreases, which implies that prefixes caused by address usage increase more rapidly.

III. From the viewpoint of Internet

If prefixes can be presented as a binary tree, only brother nodes can be aggregated, which are called brother prefixes. So the topological distance between the ASes originating this prefix and its brother prefix implies the likelihood of aggregating this prefix. Therefore, we come up with a metric called Brother Prefix Distance (BPD), which measures the topology distance between a prefix and its brother prefix. For simplicity, we use the term “distance of two prefixes” to denote the topological distance between two ASes who advertise the two prefixes. We sort all prefixes (except absorbed non-fragments) according to its allocation time and allocate each prefix to its AS in order of allocation. Each time a prefix is allocated, the BPD of this prefix is calculated, there are four cases:

None means the brother prefix is not allocated. Since BPD measures the likelihood of aggregation, and it is likely to aggregate this prefix if extra routable space is allowed [1]. Therefore, the prefix has the potential to be aggregated, and its brother prefix distance is regarded as zero.

One means an existing prefix is exactly its brother prefix. This case is simple, the distance between two prefixes is the brother prefix distance.

Many-Incomplete means that the brother prefix overlaps with many existing prefixes, which together don’t cover the brother prefix. This case is complicated, we calculate the average distance between the prefix and those overlapped prefixes.

Many-Complete means that its brother prefix overlaps with many existing prefixes, which together cover the brother prefix. This case is also complicated and similar to Many-Incomplete, and the brother prefix distance is the average distance between the prefix and those overlapped prefixes.

Figures 1c and 1d show that the average value of BPD is increasing, either in IPv4 or IPv6. This implies that newly allocated prefixes are more difficult to be aggregated. Since IPv6 is still in its early years, its average BPD value is much lower than that of IPv4 and the curve is smoother.

IV. Conclusion

In this paper, we measure the Internet routing scalability from the perspective of address allocation. We find that more prefixes are due to allocation of IP addresses rather than usage of them. Then, we propose Brother Prefix Distance and find it increasing over time, which implies that impact of address allocation on routing scalability increases over time.

Acknowledgment

This work is partially supported by the National High Technology Research and Development Program of China (863 Program) No. 2015AA015603, 2015AA016105, the National Natural Science Foundation of China (Grant No. 61202357, 61373161, 61402253).

References

[1] X. Zhao, Y. Liu, L. Wang, and B. Zhang, “On the aggregatability of router forwarding tables,” in *INFOCOM, 2010 Proceedings IEEE*, 2010, pp. 1–9.