Improving Internet Routing Scalability with AS Landmarks

Yangyang Wang, Jun Bi
Institute for Network Sciences and Cyberspace, Tsinghua University
Department of Computer Science, Tsinghua University
Tsinghua National Laboratory for Information Science and Technology (TNList), Beijing, China
Email: wangyy@cernet.edu.cn, junbi@tsinghua.edu.cn

Abstract—**Internet routing scalability issue is one of challenges of future Internet architecture. The global routing table is inflating with the growing number of networks connected to the Internet and the more-specific prefixes from multi-homing and traffic engineering operation requirements. Many engineering solutions have been proposed. In this paper, we propose a scalable routing architecture in a long-term intention for future Internet. Different from previous work, this architecture maintains AS landmark information for improving route forwarding table (FIB) aggregatability. Our evaluations on real BGP routing tables show that it can lead to an aggregated FIB of size less than 6%~22% of the original FIB size. It can be incrementally deployed in the Internet.**

Index Terms—**Internet routing, routing scalability, BGP, route aggregation, future Internet.**

I. Introduction

Internet routing scalability issue is one of challenges of future Internet architecture, and has attracted attentions from research communities and Internet operators [1]. The routing table size in Default Free Zone (DFZ) has been growing rapidly in recent years, and has exceeded 500K entries in August 2015 [18][19]. This growth can consume more capacities of memory, computation and electric power in high-efficient forwarding devices. The expansion of routing table size is attributed to both the growing Internet size and its flattening topology, and prefix fragments from practical operations, such as multi-homing and traffic engineering [1] [16]. Many solutions to improve Internet routing scalability have been proposed in the past years. These solutions generally can be divided into revolutionary and evolutionary categories in terms of the degree of required changes to today's Internet. Revolutionary solutions need significant architectural changes to the Internet. For example, the typical schemes include core-edge separation [2] (e.g., LISP [3], eFIT [4], Ivip [5]), and core-edge elimination (e.g., ILNP [6], Shim6 [7]). They need protocol changes on hosts, transit network routers, and DNS. Some studies, such as content-centric networking [8], attempt to thoroughly redesign the Internet with fundamentally different networking principles. Architectural redesign also implies high cost of conversion from old architecture to new designs. On the contrary, some studies focus on practical methods that can be directly applied to the current Internet with local upgrade of routers' software or configuration. These solutions include virtual aggregation (VA) [9], FIB route compression [10] [11], etc. They provide feasible alternatives to save space for routing table growth before a thorough resolution has been accepted and deployed. These previous work will be summarized in detail in section 2.

In this paper, we propose a new routing scheme that use autonomous system (AS) numbers to reduce the size of forwarding tables (FIBs) in the routers of transit core networks. There are a few routing architectures that use AS-grained routing in the Internet [12] [13] [14]. However, routing with AS numbers decreases the granularity of routing policies, and all traffic load will be directed to the specific AS nodes. Different from the existent work, our proposed architecture selects a few ASes as landmarks, and data packets may be forwarded based on matching IP prefixes, or the AS numbers of these landmarks. Routers in transit networks can aggregate route entries in FIBs according to AS numbers, and also need to create a mapping table between destination prefixes and these AS number based on BGP routing tables. Thus, it just improves the data forwarding plane to get FIB size reduction, and doesn't change the control plane of routing propagation and routing calculation. This architecture doesn't need construct an extra global mapping system as in LISP [3], and can be deployed individually by ISPs. Our evaluation shows that FIB size can obtain a significant reduction based on selected AS landmark information.

This paper is organized as follows. In the Section 2, we introduce the existing related work. After that, in Section 3, we briefly investigate the hierarchical property of Internet AS topology leveraged for our routing architecture in this paper. In Section 4, we describe the architecture of our proposal. In section 5, an extension on hosts to support landmarks is described. Section 6 presents a primary evaluation. Section 7 is conclusion and future work.

II. Related Work

In this section, we summaries closely related solutions to Internet routing scalability issue, as well as their benefits and costs.

The approaches of core-edge separation and core-edge elimination [2] have ever got more approvals. The typical core-edge separation solutions include LISP [3], eFIT [4], Ivip [5], etc. Core-edge separation solutions separate edge networks address space from transit networks address space, which prevent the prefixes of edge networks from being injected into the transit networks so that it can reduce the routing table size in the transit core. However, it needs to build a globally coordinated mapping system to stitch the separated address spaces.

The core-edge elimination solutions are mostly based on the ID/locator separation of IP addresses. In today's Internet, IP addresses are considered being used as an identifier of a host, as well as a locator where the host is connected to in the Internet. The typical core-edge elimination solutions based on ID/locator separation include Shim6 [7], ILNP [6], etc. In these solutions, a host uses an identifier in the transport sessions, while the IP layer uses one or more IP addresses as locators. Therefore, hosts in an edge network connecting with multiple provider networks can be assigned multiple provider-allocated addresses that can be aggregated in provider networks and further reduce global routing table size. ID/locator separation also needs a mapping system between IDs and locators.

Ballani et al. propose Virtual Aggregation (VA) [9] that can be deployed in individual ISP network independently, and reduce the forwarding table size in ISP networks. Zhang et al. [6] present an evolutionary approach from intra-domain VA to inter-domain VA, which finally conduct to core-edge separation of the Internet. Some studies focus on route compression to reduce FIB size locally [10][11][17]. Route compression requires that prefixes are consecutive, or have covering relationship in address space. Otherwise, routes cannot be aggregated.

Some researches propose AS-based routing architectures. HLP [12] is a hybrid routing architecture. It simply uses AS number in the inter-domain routing among ASes. Shue et al. [13] present a clean-slate Internet design with ID/locator split. It uses string name as host identifier, and AS number as locators. Rolf [14] presents a routing scheme where the AS numbers and IP addresses of source and destination hosts are carried as the immutable part in data packet header. And a mutable IP address field used to routing within an AS is changed when data packets traverse the AS border routers. All of these architectures above do inter-domain routing on Internet AS topology. However, AS granularity is coarse to express routing policies against IP address prefixes, and AS numbers cannot be aggregated.

Our proposal is not routing on AS nodes (AS-level routing). It just aggregates prefixes-based routes in forwarding tables using a few AS numbers (called landmarks in this paper). The prefixes that are not consecutive or have no covering relationship may also be aggregated in our proposed routing scheme. There are no changes in routing updates and decisions, and no coordinated mapping system. Because a majority of routing path will depend on the small number of high tier networks (i.e., Tier-1 and larger ISPs), the size of aggregated forwarding table can become very small, which is only about the sum of the number of AS landmarks and the number of prefixes not in mapping table. It supports incremental deployment, and partial deployment will also bring corresponding benefits to some extent.

III. Hierarchical Property of Internet AS Topology

Our proposed routing architecture in this paper leverages the hierarchical property of the Internet AS topology. In this section, we briefly investigate this hierarchical property.

The Internet AS-level topology has hierarchical property. In general, Tier-1 ASes, such as AT&T (AS7018), Cogent (AS174), etc., locate at the top of the Internet [20], and peer with every other Tier-1 ASes to provide global reachability in the Internet. And, tier-2 ASes purchase traffic transit from tier-1 networrks, and provide network access for tier-3 ASes. The eyeball ASes provide Internet access to end customers. Our proposal selects a few ASes at higher tiers as landmark ASes, by which IP packets will be delivered to their end customer networks. Although existing researches find the Internet AS topology has a flattening trend, our investigation shows that it also keeps notable hierarchical property. This will benefit our proposed routing architecture.

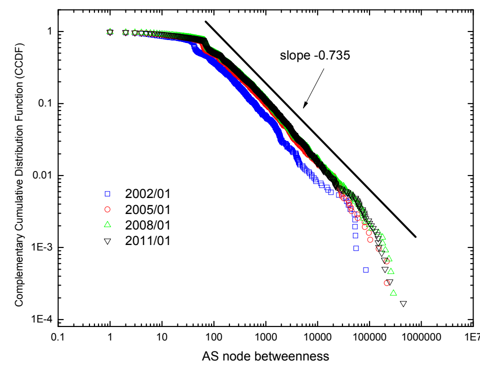

Figure 1. The distribution of AS node betweenness of BGP AS paths.

We investigate the metric AS node betweenness. The betweenness of a node in a graph is defined as the number of all shortest path traverse this node. We calculate the number of AS paths traversing each AS node observed from BGP AS paths collected from all vantage points of RouteViews [21] and RIPE RIS [22] every 3 years from 2002 to 2011. The result in Figure 1 shows power-law distributed curves, which indicate that a plenty of routing paths traverse a few AS nodes. These nodes are important for Internet connectivity and reachability.

IV. Routing Architecture Design

A. Overview of architecture

In this section, we give an overview for our proposed architecture. In this architecture, a few AS nodes are selected as landmarks, and every AS knows the reachability information of some of the landmarks. The reachability information is formed as a set of tuples of AS number and nexthop, just like the entries of destination prefix and nexthop in today's IP routing tables. Figure 2 shows an example of the packet forwarding process between host A and host B. In Figure 2, AS X and AS Y are selected as landmarks, and AS C has an aggregated FIB that contains the routes to X and Y, and has a mapping table between a few destination prefixes and corresponding AS landmarks. In this example, mapping entries are <1.1.1.0/24, X> and <2.2.2.0/24, Y>. When a packet from host A to host B arriving at a border router Rc in AS C, Rc finds that the aggregated forwarding table has no prefixes matching host B's IP address and this packet also doesn't carry any AS information. Thus, Rc will look up the mapping table and find AS X is associated to the prefix matching host B's IP address. And then, Rc adds X into the header of this packet, and forwards it to the nexthop. We assume that this packet traverses AS D. When a border router Rd in AS D receives this packet, it looks up forwarding table for IP address of host B. But the aggregated forwarding table in Rd has only routes to some remote AS landmarks and a few prefixes to nearby networks. After the missing of longest prefix match for host B's IP address, Rd looks up forwarding table for landmark X carried in this packet and gets the correct nexthop. Actually, when Rd detects that the packet carries an AS landmark, it can look up forwarding table at the same time for both host B's IP and AS number X. In this way, this packet finally will be routed to X. When the router Rx of AS X receives this packet, it will remove the carried landmark information same to AS X, and forward it to the destination host B. This is a basic illustration about packet forwarding in this architecture.

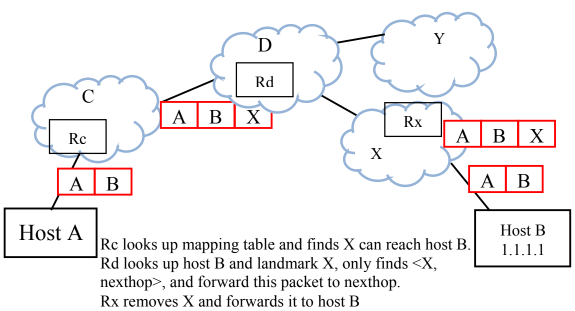

Figure 2. An example of host communications in the architecture.

B. Mapping tables creation

From the example above, we can see that mapping table is important in this architecture. After aggregation, the entire routable IP address space in original FIB is divided into two parts, one part is remained in the aggregate routing table, and another part is removed into mapping tables. It seems like the mapping systems of LISP [3] and other core-edge separation architecture. But it is different from the previous work in the creation and maintenance of mapping tables. In this architecture, every AS can create its own mapping table individually from BGP routing tables. And a mapping table and an aggregated forwarding table are calculated based on the routing table. In this calculating process, landmarks are determined according to configured policies. Figure 3 briefly describes the algorithm for creating mapping tables from routing tables. And which routes are kept in the aggregated forwarding table and which routes are removed to the mapping table can be controlled by local policies.

Figure 4 shows an example of creating mapping table and aggregated forwarding table with four entries of routes observed in a border router of AS X. Suppose that, according to some policies, AS1 and AS2 are selected as landmarks. The border router of AS X creates mapping entries <3.3.3.0/24, AS2> and <2.2.2.0/24, AS2>. And then, it generates a new entry <AS2, r1> and insert it into the aggregated forwarding table to replace the removed route entries of 2.2.2.0/24 and 3.3.3.0/24. Thus, this aggregated forwarding table has a reduced size than original routing table.

INPUTS: a set of landmarks *L*; a set of prefixes *P* and the routes to these prefixes. For prefix *p*, the AS path to *p* is denoted by *as-path(p)*, the nexthop of p is denoted by *nexthop(p)*
OUTPUTS: a mapping table M and a forwarding table F
For every prefix *p* in the set *P*:
  if a landmark $l \in L$ is in *as-path(p)*
    choose the one closest to destination p, denoted as *lr*
    create an entry <*p*, *lr*>, and insert it to M
    create an entry <*lr*, *nexthop(p)*>, and insert it to F
  elseif none of landmarks is in *as-path(p)*
    create an entry <*p*, *nexthop(p)*>, and insert it to F

Figure 3. Algorithm for creating mapping table.

Prefix	Nexthop	AS path
1.1.1.0/24	r1	AS3, AS4, AS5
2.2.2.0/24	r3	AS9, **AS2**, AS6
3.3.3.0/24	r3	AS9, **AS1**, **AS2**, AS7
4.4.4.0/24	r2	AS8, AS5

Figure 4. An example of routing table entries. AS1 and AS2 are selected as landmarks in this example

However, in practice, there are possible scenarios that one AS landmark corresponds to more than one different nexthops from different prefixes. For example, if the route to 3.3.3.0/24 in Figure 3 is <3.3.3.0/24, r1, (AS3, AS1, AS2, AS7)>, then it will create two forwarding entries <AS2, r3> and <AS2, r1>, and mapping entries <2.2.2.0/24, AS2>, <3.3.3.0/24, AS2>. All of the packets with destinations to 2.2.2.0/24 and 3.3.3.0/24 will be added landmark AS2 in their packet header after looking up mapping table. But AS2 corresponds to two different nexthops in aggregated forwarding table. In this case, there are two ways to resolve this problem. 1) The first way is to choose only one nexthop (e.g., choose the nexthop having shortest AS path to destination) for AS2, either <AS2, r1> or <AS2, r3>. It doesn't impact the reachability to destination prefixes 2.2.2.0/24 and 3.3.3.0/24 because AS2 is reachable and AS2 knows the routes to the two destination prefixes. But it could lead to inconsistency with finer-grained routing information expressed in prefixes, and it requires every traversed AS in the path to AS2 has the recheability information to AS2. In this example, AS9, AS1, and AS3 are required to have a route to AS2 in their forwarding tables. However, this resolution can benefit FIB reduction because more route entries of IP prefixes are aggregated by the small number of AS numbers, and moved into mapping tables. 2) The second way is to keep the same finer-grained routing as that at prefix-level. In this example, we can keep the route to 2.2.2.0/24 in aggregated forwarding table, and only create the mapping entry <3.3.3.0/24, AS2> and AS forwarding entry <AS2, r3>. But its forwarding table could be larger than that of the first way, because fewer prefixes are moved into mapping table in order to keep the finer-grained routes in the aggregated forwarding table being equal to the original routing table. We call the first way "**strict AS landmark routing**", and the second way "**constrained AS landmark routing**". These names are to be used in the following paragraphs and the evaluation in section VI.

C. A centralized structure for landmark selection

In subsection B, all landmarks selection and mapping table calculation can be done individually on the local routers in an AS. This could lead to that each AS may have a different landmark decisions due to its own views of Internet AS topology (In fact, almost all of larger ASes and the connections among them can be seen by the ASes in the Internet, the difference is small for selecting larger ASes as landmarks). We can consider it as a distributed structure for landmark selection.

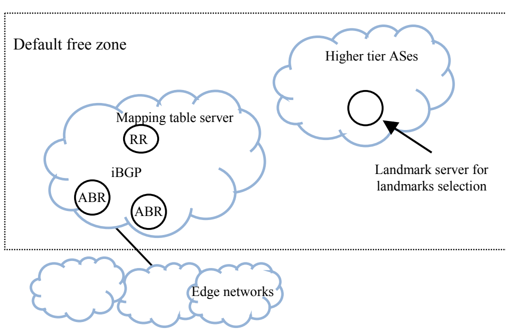

Figure 5. A centralized structure for landmark selection and mapping table creation.

In order to create a globally uniform landmark selection and mapping table creation, we propose a centralized structure. It is shown in Figure 5. The high-tier ASes deploy several public landmark servers used to collect AS topology and determine AS landmarks in terms of AS degree rank or other selection strategies. Every AS in default free zone can run mapping table servers that periodically connect the public landmark servers to get selected landmarks. In general, the set of landmarks is stable when it usually selects the larger ASes as landmarks. Thus, it doesn't need to communicate with the landmark servers frequently. The mapping table servers in an AS can be deployed on BGP route reflectors (RRs) that know the routing table of that AS. And the mapping table servers calculate mapping table and aggregated forwarding table, and then distribute them to AS border routers (ABRs) through iBGP sessions. This structure can conduct to globally consistent landmark selection and save the computation cost for creating mapping tables in each AS.

V. Extension to Hosts

The design described in the above sections does not need protocol changes on hosts. All of work can be handled in access networks and transit networks. In this section, we show that a bit of upgrade on host protocol stack can get more features to Internet routing and host mobility.

A. Host support to AS landmarks

The IP protocol stack on hosts can be modified to support AS landmark option in IP packet headers. Such that, hosts have the ability to directly recognize and send packets carrying with AS information in packet header. This will save the cost of AS addition and removal in packet headers in transit networks. As illustrated in Figure 6. The process of communication between host A and host B is as follows.

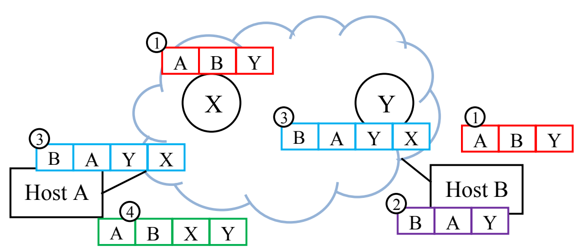

Figure 6. Host porotol stack supports AS landmark option in packet header. Packet format (B, A, Y, X) means source address is host B's IP address, destination is host A's IP address, source landmark is Y, and destination landmark is X; the numbers in circle mean the step sequence.

Step 1: Host A sends a packet to host B. When this packet arrives at AS X, it looks up forwarding table and finds no entry matched. Then, it looks up mapping table and finds out landmark AS Y can reach host B. This landmark Y is referred to as destination landmark, and X adds AS number Y to the packet header. When this packet arrives at Y, Y will forward it to host B without removing the carried AS number in this packet.

Step 2: After host B receives this packet, its protocol stack recognizes the carried landmark Y in the packet. When host B sends a reply back to host A, it will add landmark Y to the reply packet as a source landmark. However, the source landmark does not mean the packet is delivered to host A by traversing this source landmark. It is just used as a destination landmark for host B in the reverse direction in Step 4.

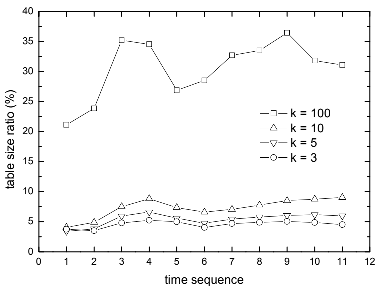

Figure 7. Select ASes of degree larger than *k*, and strict AS landmark routing

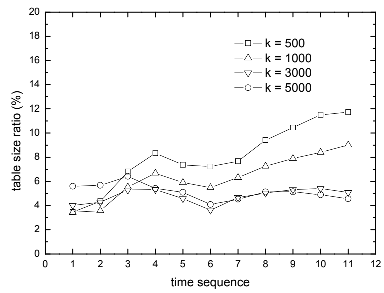

Figure 8. Select ASes of degree top *k*, and strict AS landmark routing

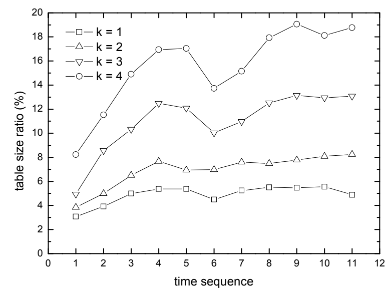

Figure 9. Select ASes of rank larger *k*, and strict AS landmark routing

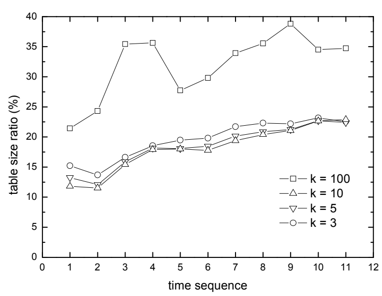

Figure 10. Select ASes of degree larger than *k*, and constrained AS landmark routing

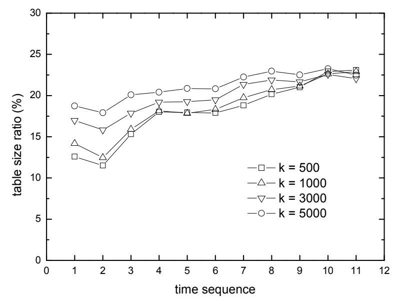

Figure 11. Select ASes of degree top *k*, and constrained AS landmark routing

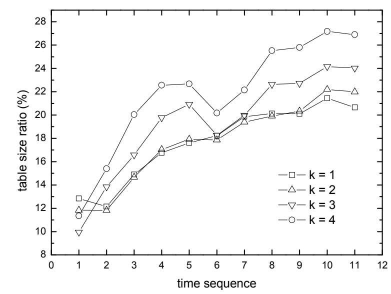

Figure 12. Select ASes of rank larger *k*, and constrained AS landmark routing

Step 3: When the reply packet arrives at AS Y, Y will look up its forwarding table and then mapping table. It finds out AS X can reach host A, and adds AS number X into the packet header as a destination landmark. At last, this reply packet will arrive at host A.

Step 4: Host A receives this reply, and records the source landmark Y and destination landmark X. When host A sends packets to host B again after that, it will add the two landmarks into these packets, and use X as the source landmark, and Y as the destination landmark. When AS X forwards these packets, it can find out matched next-hop by looking up forwarding table for the destination landmark Y, without the need to look up mapping table.

From the above process steps, we can see that the mapping table lookup, and the landmark addition and removal are saved in the subsequent packets after the first round trip of communication. Note that source landmark in a packet does not mean the packet will traverse this landmark on the path from the source host to the destination host. For reasons of routing policies, load balancing, and traffic engineering, etc., the AS path between a pair of source and destination hosts may be symmetry in the Internet.

B. Host-controlled multipath routing

If a host, denoted as host A, knows it can be reached through multiple landmarks from previous communications with other hosts, it can tell these landmarks to a newly communicated host B. And host B can direct its packets to host A through multiple landmarks by add various landmark options in each packets sent to host A. If one of path to host A fails, host B can retour the failed path through other landmarks. However, this feature may invade traffic engineering. We just show that it can support multipath routing to some extent.

VI. Evaluation

In this section, we present a primary evaluation on the effect of our proposed routing method for improving Internet routing scalability.

A. Data sources

We collect the BGP routing table snapshots from the collector "route-views2" of the Oregon Route Views Project [15]. This collector establishes BGP peer session with multiple ASes. We randomly choose the peer of AS3356 and extract 10 snapshots of BGP routing table of the first day of every year from 2002 to 2012. We use these routing table data to evaluate the FIB aggregation effect using AS landmarks.

For each routing table snapshot, we calculate the AS-level topology and AS node degree. We use three methods to select AS landmarks. 1) Sort AS node in descending order of AS node degrees, and select the top *k* nodes; 2) Sort AS nodes in descending order of AS node degrees, and select the nodes of degree larger than *k*; 3) We use the method in [15] to rank AS nodes according topology hierarchy, and select the nodes with rank larger than *k*.

B. Definition of table size ratio

Assume that the number of original FIB entries is t. Given a set of landmarks, an aggregated FIB and a mapping table can be calculated. The aggregated FIB size is denoted with w, and the mapping table size is m. The number of landmarks is l. After creating the mapping table, a part of routes are moved from original FIB to the mapping table, and the rest routes are inserted to the aggregated FIB. The number of routes that appear in both original FIB and aggregated FIB is denoted as u, while the number of routes moved to mapping table is just the mapping table size m. Thus, the original FIB size $t = u + m$, and the aggregated FIB size $w = u + l$. The ratio r of aggregated FIB size to original FIB size reflects the extent of FIB reduction.

$$r = \frac{w}{t} = \frac{u+l}{u+m}$$

C. Results

The results in the figures from Figure 7 to Figure 12 plot the fluctuation of table size ratios with growing time series. The results in Figure 7, 8 and 9 are the table size ratios of ***strict AS landmark routing*** (mentioned in section IV) with three kinds of landmark selection polices. They can generate much smaller table size than that of ***constrained AS landmark routing*** shown in Figure 10, 11 and 12. The aggregated FIB size can be less than 6% of the original FIB size, and grows much slowly. As explained in section IV, the strict AS landmark routing can replace more prefixes entries with a coarse-grained routes using AS landmarks, but it may invade finer-grained routing of traffic engineering or load-balancing expressed with IP address prefixes. However, it makes sense in a very large-scale future Internet where finer-grained control is difficult and in no sense. The FIBs in Figure 10~12 keep the equal finer-grained routing effect to original FIB tables, and get size ratios about 22% of the original FIBs in recent years around 2011. But the ratio grows faster because more and more entries cannot be aggregated.

VII. Conclusion and Future work

In this paper, we propose a routing architecture to improve Internet routing scalability by forwarding table aggregation with a small number of AS landmarks. These AS landmarks are used in forwarding table entries and packet header. And, packet delivery can be performed on IP addresses and AS landmarks. It can be deployed incrementally in BGP routing system. The landmarks can be embedded in a particular form of IP addresses, or carried in IPv4/IPv6 header options. Its needs mapping tables, but that can be created from local BGP routing tables. From a primary evaluation, the result shows a considerable reduction of forwarding table size (about 6% and 22% of original size in different cases). In future work, we plan to find a method to create a mapping table that can produce an optimal forwarding table with minimum size under a given a set of landmarks, and make a further evaluation of its impact on BGP routing paths.

Acknowledgment

This work is Supported by the National High-tech R&D Program ("863" Program) of China (No.2013AA013505) and the National Science Foundation of China (No.61472213).

References

[1] D. Meyer, L. Zhang, and K. Fall, "Report from the IAB Workshop on Routing and Addressing", RFC4984
[2] Dan Jen, Michael Meisel, et al. "Towards A New Internet Routing Architecture: Arguments for Separating Edges from Transit Core". HotNets 2008
[3] D. Farinacci, V. Fuller, D. Meyer, and D. Levis. Locator/ID Separation Protocol (LISP). draft-farinacci-lisp-12, Sep. 2009.
[4] eFIT project, http://irl.cs.ucla.edu/efit/index.html
[5] Ivip, http://www.firstpr.com.au/ip/ivip/ .
[6] R. Atkinson, S. Bhatti, S. Hailes, Evolving the Internet Architecture Through Naming, IEEE Journal on Selected Areas in Communication (JSAC), vol. 28, no. 8, pp1319-1325, Oct 2010.
[7] E. Nordmark and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6" , RFC5533
[8] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, R. L. Braynard (PARC) Networking Named Content, CoNEXT 2009, Rome, December, 2009.
[9] Hitesh Ballani, Paul Francis, Tuan Cao and Jia Wang. "Making Routers Last Longer with ViAggre".USENIX NSDI 2009, Boston, MA, April 2009.
[10] Richard Draves, Christopher King, Venkatachary Srinivasan, Brian Zill: Constructing Optimal IP Routing Tables. INFOCOM 1999:88-97
[11] Xin Zhao, Yaoqing Liu, Lan Wang, Beichuan Zhang: On the Aggregatability of Router Forwarding Tables. INFOCOM 2010:848-856
[12] Lakshminarayanan Subramanian, Matthew Caesar, Cheng Tien Ee, Mark Handley, Zhuoqing Morley Mao, Scott Shenker, Ion Stoica: HLP: a next generation inter-domain routing protocol. SIGCOMM 2005:13-24.
[13] Craig A. Shue, Minaxi Gupta: An Internet without the Internet protocol. Computer Networks (CN) 54(18):3232-3245 (2010).
[14] R. Winter, "Two Steps Towards Practical Compact Routing", 21st International Teletraffic Congress (ITC21), Paris, September 2009.
[15] Lakshminarayanan Subramanian, Sharad Agarwal, Jennifer Rexford, Randy H. Katz: Characterizing the Internet Hierarchy from Multiple Vantage Points. INFOCOM 2002.
[16] Bu, Tian, Lixin Gao, and Don Towsley. "On characterizing BGP routing table growth." Computer Networks 45.1 (2004): 45-54.
[17] Sobrinho, João Luís, et al. "Distributed Route Aggregation on the Global Network." Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies. ACM, 2014.
[18] http://www.bgpmon.net/what-caused-todays-internet-hiccup/
[19] http://mailman.nanog.org/pipermail/nanog/2014-June/067560.html
[20] http://en.wikipedia.org/wiki/Tier_1_network
[21] University of Oregon Route Views Project, http://www.routeviews.org/
[22] RIPE RIS Project, https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data