On Routing Table Growth

Tian Bu , Lixin Gao, and Don Towsley

The Internet has experienced explosive growth since its commercialization. The Internet is divided into thousands of autonomous systems (ASes), each of which consists of networks of hosts or routers administrated by a single organization. Hosts and routers are identified with 32-bit IP addresses, which brings to a total of 2^{32} (more than 4 billion) possible IP addresses. To ensure the scalability of the Internet routing infrastructure, IP addresses are aggregated into contiguous blocks, called *prefixes*. Routers exchange reachability information for each prefix using the Border Gateway Protocol (BGP). As a consequence, each BGP routing table entry contains reachability information for a single prefix. The size of a BGP routing table is the number of prefixes contained in the routing table. The size of routing tables has risen from 10,000 to 100,000 over the past six years This dramatic growth of the routing table can decrease the packet forwarding speed and demand more router memory space. Some experts have predicted that if the router memory is to keep pace with the growth of the routing tables, each router will require gigabits of memory within the next two years.

In this work, we explore the extent that various factors contribute to the routing table growth and predict the future rate of the growth of the routing table. A prefix consists of a 32-bit IP address and a mask length (e.g., 1.2.3.0/24 represents IP block 1.2.3.0-1.2.3.255). Since the introduction of Classless Inter-domain Routing (CIDR) a prefix can be of any length. This enables more aggressive *route aggregation* in which a single prefix is used to announce the routes to multiple prefixes. For example, prefixes 1.2.3.0/24 and 1.2.2.0/24 can be aggregated as prefix 1.2.2.0/23, and prefixes 1.2.2.0/23 and 1.2.3.0/24 can be aggregated as prefix 1.2.2.0/23. Route aggregation, however, might not always be performed. First, an AS can aggregate its prefix with its provider's only when the AS is *single-homed*, i.e., the AS has only one provider. For a *multi-homed* AS, which has multiple providers, its prefix(es) cannot be aggregated by all of its providers. Second, an AS may choose not to aggregate prefixes originated by it. One reason that an AS originates several prefixes is that an AS *fails to aggregate* aggregatable prefixes originated by it. The second reason that an AS originates several prefixes is load balancing. An AS originates several prefixes so as to perform *load balancing* by reaching different prefixes via different AS paths. The third reason that an AS originates several prefixes is address fragmentation. *Address fragmentation* is caused by a set of prefixes originated by the same AS that cannot be summarized by one prefix.

We explore the extent that factors such as multi-homing, failure to aggregate, load balancing, and address fragmentation contribute to routing table size. We examine the BGP routing table from the Route Views server, present techniques to quantify and perform measurement study on these factors. We find that multi-homing introduces around $20-30\%$ extra prefixes. Next, we explore how load balancing can contribute to routing table size and show that load balancing introduces around $20-25\%$ extra prefixes. However, multi-homing and load balancing are necessary trends and cannot be eliminated.

This leads us to consider how the failure to aggregate can affect the routing table size and find that failure to aggregate increases the routing table size by only $15-20\%$. Finally, we explore the extent that address fragmentation contributes to the routing table size and find that address fragmentation contributes to more than 75% of routing table size. Clearly, address fragmentation contributes to the routing table size the most. This leads us to introduce the concept of the *prefix cluster*, the maximal set of prefixes originated by the same AS that are not aggregated due to either failure to aggregate or address fragmentation. In other words, a prefix cluster is a maximal set of prefixes among which no load balancing is performed, i.e., that are announced identically by any router. We show that the number of prefix clusters is no more than 20% of the number of prefixes.

It is important to predict the future growth pattern of prefixes and prefix clusters. To do so we take advantage of an observation that both the number of prefixes and the number of prefix clusters originated by an AS can be approximated by power-laws. Using these power-law approximations, we estimate the number of prefixes and the number of prefix clusters given the number of ASes. We can predict the number of prefixes and prefix clusters as the number of ASes grows. We observe that the number of prefixes grows much faster than the number of prefix clusters does. To the best of our knowledge, this is the first study on the explosive growth of routing table by systematically comparing factors that contribute to the growth and by observing routing table growth patterns.

The authors were supported in part by NSF grant ANI-9977555, ANI-0085848, and NSF CAREER Award grant ANI-9875513. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Tian Bu and Don Towsley are with Department of Computer Science, University of Massachusetts, Amherst, MA 01003; {tbu,towsley}@cs.umass.edu

Lixin Gao is with Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA 01003; lgao@ecs.umass.edu

Full version of the paper is available at http://www-net.cs.umass.edu/ ˜ tbu