Measuring the Effects of Route Prepending for Stub Autonomous Systems

Samantha S. M. Lo and Rocky K. C. Chang

Department of Computing
The Hong Kong Polytechnic University
Hung Hom, Kowloon, Hong Kong, SAR, China
Email: {cssmlo|csrchang}@comp.polyu.edu.hk

This paper considers the widely practiced *AS-path prepending* (or just prepending) approach which is usually applied to control the inbound traffic of a multihomed AS. To explain how the prepending approach works, consider a multihomed AS with links 1 and 2, which are connected to two different ISPs. For instance, link 1 is much more congested with inbound traffic than link 2. Therefore, the network administrator may add several of its AS number in the BGP routes advertised to link 1, while the BGP routes advertised to link 2 do not have this artificial addition. Since one of the major criteria of selecting the best route is based on a shortest AS-path length, the prepended route is more likely to be rejected by the upstream ISPs, thus reducing the traffic coming into link 1. Moreover, the affected traffic will instead be received from link 2.

Compared with other methods in controlling the inbound traffic, the AS-path prepending method offers many advantages, such as providing network resilience and not increasing the routing table size. However, there is a lack of in-depth studies on the prevalence and the effectiveness of this approach. Although the practical experience has indicated that the prepending is very effective in changing routes, there is a void in understanding why and under what conditions this approach is effective. We attempt to address some of these issues in this paper.

I. Active measurement of the prepending method

We have set up an active measurement facility in a stub, dual-homed AS, which we simply call *Home AS* in order to conceal its identity. As depicted in Figure 1, a BGP router announces routes for a *beacon prefix* to $AS1$ (a tier-1 ISP) and $AS2$ (a regional ISP). The beacon prefix is a set of addresses with prefix $/21$ in the Home AS that are not in use. In other words, the Home AS normally does not expect to receive traffic destined to the beacon prefix. As a result, our active measurement will not affect the normal operation of the Home AS[1] and the Internet.

In all experiments, we prepend only link 1, and use a set of 16 route servers and 42 looking glasses as the set of (virtual) traffic sources. The overall objective is to study the impact of the prepending on link 1 on the routes for the sources to reach the Home AS. The route servers are two to four ASes away from the Home AS; the looking glasses are much farther away: most of them are five to seven ASes away. Therefore, we could also evaluate the impact of AS-path prepending with respect to the AS-path length.

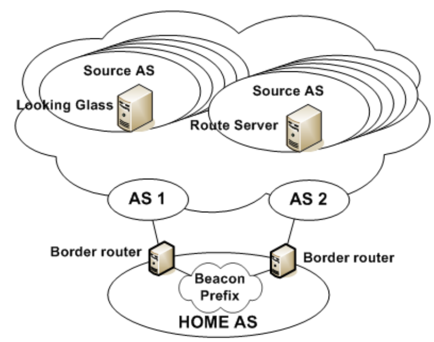

Fig. 1. The active experiment setup at the Home AS.

Since we apply prepending on link 1 only, this prepending does not affect the traffic that is already using link 2. Before effecting prepending on link 1, we therefore first identify the set of routes which use link 1. After announcing a new prepending value, we then look for route changes as a result of the prepending. For the route servers, we can observe the changes from their BGP routing tables; for the looking glasses, we perform reverse `traceroute` and derive the AS-path from the `traceroute` data.

Moreover, we have performed *forward prepending* by increasing the prepending length from zero (i.e., no prepending) to five, and *backward prepending* by decreasing the prepending length from five to zero. The maximum prepending length is five, because we do not observe further route changes beyond a prepending length of four.

II. The measurement results

Figure 2(a) shows the link usages by the incoming traffic from the set of router servers. All the 16 route servers use link 1 when there is no prepending (prepending length of 0). With a prepending length of one or two, none of them switch

[1]The stub AS is AS17764, and the prefix is `202.125.232.0/21`. We performed the measurement experiments in January to February 2005.

to link 2. However, a further increment beyond two brings about an abrupt change in the incoming link for nine of them. Altogether, eleven route servers (68%) respond to the AS-path prepending method, and five of them do not. After decreasing the prepending length back to zero, we observe that the distributions of the link usages are not the same for the cases of $1 \rightarrow 2$ and $3 \rightarrow 2$. In other words, the link usage distribution for a given prepending length of two depends on whether the prepending is obtained by forward prepending or backward prepending. In general, the unbalancing phenomenon occurs when the link usage distributions are not the same for the cases of $m-1 \rightarrow m$ and $m+1 \rightarrow m$, where $m > 1$ is the prepending length.

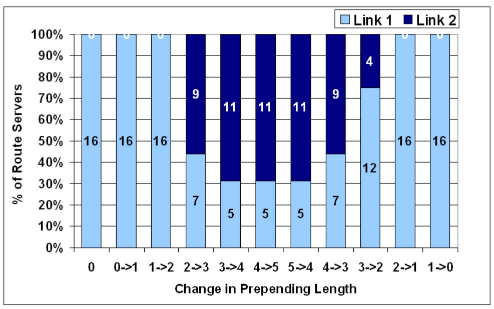

(a) Route servers

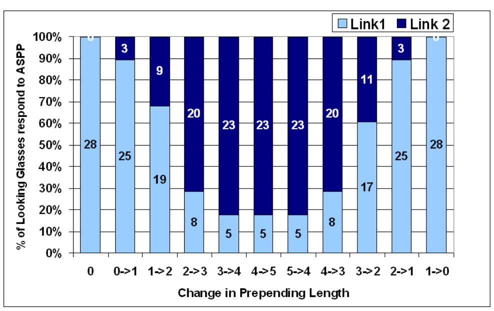

(b) Looking glasses

Fig. 2. The incoming link usages by the route servers and looking glasses under different prepending lengths.

Figure 2(b) shows the link usages by the incoming traffic from the set of looking glasses. When there was no prepending, there are 28 of them using link 1 and 14 of them using link 2. As compared with route server results, the link changes take place more gradually. Each increment in the prepending length results in link changes for some looking glasses until the prepending length reached four. Altogether, 23 looking glasses (82%) respond to the AS-path prepending method. After decreasing the prepending length back to zero, we also observe a slight unbalance between $1 \rightarrow 2$ and $3 \rightarrow 2$.

In our measurement, all the routes are restored to use link 1 after decreasing the prepending length back to zero. Therefore, we do not need to reset the connection to link 2 in order to revert the routes to link 1. However, this may not be true for other other ASes.

The AS-path length for reaching the Home AS via link 1	via link 2				
	5	6	7	8	9
4	1	**32**	5	1	1
5	0	0	**3**	0	0
6	0	1	0	**2**	1

TABLE I

THE NUMBER OF DIRECT-RESPONSIVE ASES FOR EACH COMBINATION OF AS-PATH LENGTH FOR REACHING THE HOME AS VIA LINK 1 AND AND THAT VIA LINK 2.

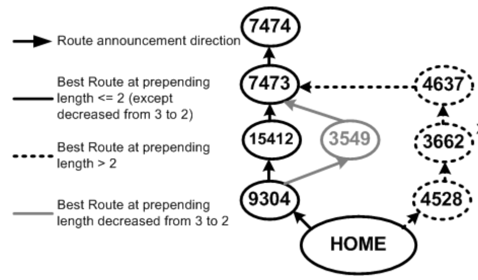

Fig. 3. Example of route change. AS7473 is the direct-responsive AS and AS7474 is the indirect-responsive AS.

To sum up, the active measurement results in the Home AS have shown that the prepending method is quite effective in changing the best routes for the set of sources. This particular set of data also shows that the method can influence more routes when the sources are located farther away from the Home AS with the same range of prepending length. Moreover, the two sets of sources share the commonality that a prepending length of more than four does not cause further route changes.

A. Which ASes are responsive to prepending?

Route changes are mainly caused by *responsive ASes* which switch their routes from link 1 to link 2 after sufficient prepending. We further classify responsive ASes according to whether they have changed their next-hops for reaching the Home AS. Consider the example in Figure 3. AS7473 changes the next-hop after receiving a prepended route from AS15412. However, AS7474 does not change its next-hop; instead, its route change is the result of the next-hop change by AS7473. We refer these two classes of responsive ASes to as *direct-responsive ASes* and *indirect-responsive ASes*, respectively. There are 47 direct-responsive ASes and 26 indirect-responsive ASes out of 116 ASes in our measurement.

Table I shows the number of direct-responsive ASes and their AS-path lengths to the Home AS. We have found that the AS-path length of the majority of the direct-responsive ASes are four for reaching link 1 and six for reaching link 2. That is, the difference in the AS-path lengths for reaching the two links is two. This explains why the greatest route change takes place at increasing the prepending length from two to three.