

The Extent of AS Path Inflation by Routing Policies

Lixin Gao and Feng Wang

Department of Electrical and Computer Engineering, University of Massachusetts

Amherst, MA 01002

{lgao, fewang}@ecs.umass.edu

Abstract— A route in the Internet may take a longer AS path than the shortest AS path due to routing policies. In this paper, we systematically analyze AS paths and quantify the extent to which routing policies inflate AS paths. The results show that AS path inflation in the Internet is more prevalent than expected. We first present the extent of AS path inflation observed from the Route View routing tables. From an ISP, at least 55% of AS paths are inflated by at least one AS hop and AS paths can be inflated by as long as 6 AS hops. We then employ two typical routing policies to show the extent of AS path inflation for all AS pairs, we find that at least 45% of AS paths are inflated by at least one AS hop and AS paths can be inflated by as long as 9 AS hops. Quantifying AS path inflation in the Internet has important implications on the extent of routing policies and traffic engineering performed on the Internet, and BGP convergence speed.

I. INTRODUCTION

The Internet connects thousands of Autonomous Systems (ASs) operated by different Internet Service Providers (ISPs), companies, and universities. Routing within an AS is controlled by intradomain protocols such as static routing, OSPF, IS-IS, and RIP. Border Gateway Protocol (BGP) [1], [2] is an interdomain routing protocol that allows ASs to apply local policies for selecting routes and propagating routing information. These routing policies are typically constrained by contractual commercial agreements between administrative domains. It is well known that an AS may take a longer AS path than the shortest AS path possibly as a result of these routing policies. However, the extent to which routing policies inflate AS paths in the Internet has not been systematically analyzed or quantified.

Quantifying AS path inflation in the Internet has important implications on the extent of routing policies and traffic engineering performed on the Internet, and BGP convergence speed. First, since ISPs typically do not make their routing policies public, it is not clear how prevalent these routing policies are and to what extent AS paths are inflated due to routing policies. Second, BGP protocol studies [5], [6] have shown that BGP convergence speed is directly correlated with AS path length. The extent of AS path inflation indicates the extent to which routing policies can increase BGP convergence time.

In this paper, we systematically study AS paths and quantify the extent that AS paths are inflated by routing policies. Our results show that AS path inflation in the Internet is more prevalent than expected. We derive chosen AS paths and shortest AS paths from the Route View BGP routing tables [8]. In particular, we collect statistics of AS path length from ISPs of various sizes: a tier-1 ISP, a tier-2 ISP, and a tier-3 ISP. From the tier-1

ISP, about 20% of chosen AS paths are longer than the shortest AS paths and AS paths can be inflated by as long as 4 AS hops. From the tier-2 ISP, at least 55% of chosen AS paths are longer than the shortest AS paths and AS paths can be inflated by as long as 6 AS hops. From the tier-3 ISP, more than 22% of chosen AS paths are longer than the shortest AS paths and AS paths can be inflated by as long as 5 AS hops. In order to understand the overall extent of AS path inflation, we collect statistics on all chosen AS paths that are visible from the Route View server. In particular, we find that about 20% of AS pairs take longer paths than the shortest AS paths and AS paths can be inflated by as long as 10 AS hops.

We present two typical routing policies to show the extent to which an AS path can be inflated by the routing policies. The first routing policy assumes that each AS obeys commercial agreements with its neighboring ASs. Commercial agreements between pairs of ASs can be classified into customer-provider, peering, and sibling [3], [4]. A provider can transit traffic for its customers or siblings. However, a customer does not transit traffic between two of its providers and peers. We compute the AS paths that conform to this routing policy. Our results show that this routing policy inflates AS paths for only 4% AS pairs. The second routing policy assumes that an AS prefers customer routes over provider or peer routes. We derive AS paths that conform to the second routing policy, and find that more than 45% of AS pairs use a longer AS path than the shortest AS path.

The only study on the Internet path length [7] assumes that each AS chooses the shortest AS path. With this assumption, the authors conclude that 20% of Internet paths are inflated by more than five router-level hops. Our study complements their work on that we focus on AS path inflation instead of router-level hop inflation. In addition, we quantify AS path inflation by analyzing real BGP routing tables and explore routing policies that conform to commercial relationships.

The remainder of the paper is structured as follows. In Section II, we quantify the extent of AS path inflation in the Internet by analyzing BGP routing tables. Section III presents AS path inflation for all AS pairs by examining two typical routing policies guided by commercial relationships. We derive the extent to which AS paths are inflated by these routing policies for all AS pairs. We conclude the paper in Section IV with a summary.

II. AS PATH INFLATION OBSERVED FROM ROUTE VIEW

In order to quantify the extent of AS path inflation, we derive the chosen AS path length and the shortest AS path length between a pair of ASs. To this end, we use routing tables from the Route View router in Oregon [8], which has 41 peering sessions. We construct an AS graph $G = (V, E)$, where the node set V consists of ASs and the edge set E consists of AS pairs that exchange traffic with each other. The shortest path between

The authors are supported in part by NSF grant ANI-0208116, 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.

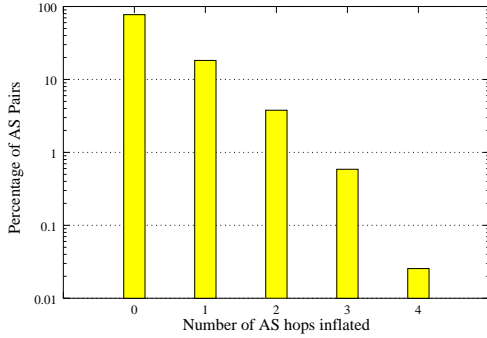


Fig. 1. Percentage of ASs whose AS paths from tier-1 ISP are inflated by a fixed number

a pair of ASs can be derived by using Dijkstra’s algorithm [10]. Although some AS level connections might not be included in the AS graph constructed from Route Views, this can only result in an overestimation of shortest AS path length.

The chosen AS path length between a pair of ASs is computed by the AS path appearing in the BGP routing table. We eliminate the AS prepending effect on the routing table. In other words, if an AS appears in an AS path several times, then we count the AS only once in the AS path length. Since an AS pair might use different AS paths for different destination prefixes, we choose to use the longest for the chosen AS path length to see the extent of AS path inflation.

We analyze routing tables during year 2000 and 2001, and present the results for 12/11/2000 only throughout this paper since the results for other dates are similar. To demonstrate the AS path inflation from different ASs, we first illustrate the extent of AS path inflation from three ISPs of various sizes: a tier-1 ISP, a tier-2 ISP, and a tier-3 ISP. An ISP is called a tier-1 ISP if it accesses the global Internet and does not buy network capacity from other ISPs. Providers that buy part or all of their interconnectivity from tier-1 ISPs are tier-2 ISPs. A local ISP is defined as a tier-3 ISP. We then show the extent of AS path inflation for all AS pairs whose chosen AS paths are visible from the Route View routing table.

To quantify the difference between the chosen AS path length and the shortest AS path length from the tier-1 ISP, we plot the number of ASs whose paths from the tier-1 ISP are inflated by a fixed number in Figure 1. The plot shows that about 20% of ASs whose paths from the tier-1 ISP are inflated by at least one AS hop, and AS paths can be inflated by as long as 4 AS hops. Figure 2 shows the distribution of inflation with respect to the shortest AS path. We see that most AS path inflations occur within AS path length 1 to 4. Most AS paths are inflated by one hop. In addition, we see that the tier-1 ISP is at most 6 AS hops to any AS in the shortest path while it can be as long as 8 hops in the chosen path.

In Figure 3, more than 55% of AS pairs have a longer chosen AS path than the shortest AS path, and AS path from the tier-2 ISP can be inflated by as long as 6 AS hops. Figure 4 shows that from the tier-2 ISP, most inflations occur within AS path length 1 to 5, and the chosen AS paths can be as long as 10 hops from the tier-2 ISP while the shortest AS paths are at most 6 hops. Note that the AS path inflation is more significant for the tier-2 ISP than for the tier-1 ISP. This means that a tier-2 ISP might be affected by routing policies more severely than a tier-1 ISP is.

In Figure 5, more than 22% of AS pairs have a longer chosen AS path than the shortest AS path and AS paths can be inflated

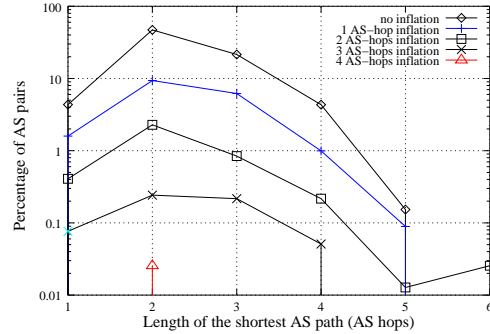


Fig. 2. Distribution of AS path inflation from tier-1 ISP




Fig. 3. Percentage of ASs whose AS paths from tier-2 ISP are inflated by a fixed number

by as long as 5 AS hops. Note that although AS path inflation is severer from the tier-3 ISP than from the tier-1 ISP, it appears that AS path inflation is comparable with that from the tier-2 ISP and from the tier-3 ISP. Figure 6 shows that most AS path inflations occur within AS path length 2 to 6. We see that chosen AS paths can be as long as 10 AS hops while the shortest AS path is at most 7 AS hops.

AS path inflation which we observe above is derived from one vantage point, such as a tier-1 ISP. Now we investigate AS path inflation from multiple vantage points. We select those AS pairs whose chosen paths are visible from the Route View routing table. Figure 7 shows about 20% of AS pairs have a longer chosen AS path than the shortest AS path and AS path can be inflated by as long as 10 AS hops. In Figure 8, we see that chosen AS paths can be as long as 13 hops while the shortest AS paths are at most 7 hops. Since the majority peers of the Route View server are tier-1 ISPs, the result can be biased towards tier-1 providers. The extent of AS path inflation for all AS pairs can be larger than we can see from the discrepancy between the non-tier-1 ISPs and the tier-1 ISPs. In the next section, we confirm this by measuring AS path inflation resulted from two typical routing policies for all AS pairs.

III. AS PATH INFLATION BY TWO ROUTING POLICIES

We can see that the extent of AS path inflation varies from ISPs to ISPs. Because not all ISPs are willing to reveal their routing policies, it is hard to get an overall picture of the AS path inflation. We derive AS path inflation by assuming two typical routing policies in this section. We first present the two routing policies that conform to commercial contractual agree-




Fig. 4. Distribution of inflation from tier-2 ISP

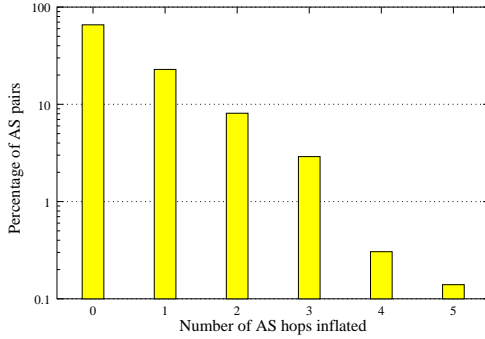


Fig. 5. Percentage of ASs whose AS paths from tier-3 ISP are inflated by a fixed number




Fig. 6. Distribution of inflation from tier-3 ISP

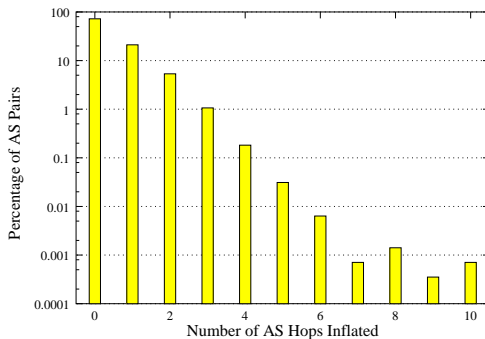


Fig. 7. Percentage of AS pairs whose chosen paths are visible from the Route View server are inflated by a fixed number

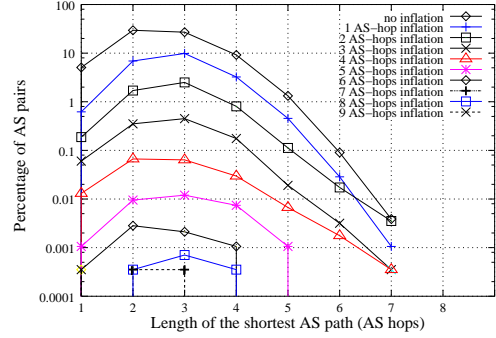


Fig. 8. Distribution of inflation for all AS pairs whose chosen paths are visible from the Route View server

ments. We then compare the derived AS path length with the shortest AS path length for all AS pairs.

A. Two Routing Policies

Routing policies typically conform to the relationships between ASs. A customer pays its provider for connectivity to the rest of the Internet. A pair of peers agree to exchange traffic between their respective customers free of charge. An AS pair has a sibling-sibling relationship if it has a mutual-transit agreement to provide connectivity to the rest of the Internet for each other. To derive AS relationships, we use the algorithm presented in [11].

The AS relationships can be translated into two typical routing policies as follows [12], [4], [14].

No-valley Routing Policy: An AS does not provide transit services between any two of its providers or peers. That is, in an AS path (u_1, u_2, \dots, u_n) , if (u_i, u_{i+1}) has a provider-customer or peer-peer relationship, then (u_j, u_{j+1}) must have either a provider-customer or a sibling-sibling relationship for any $i < j < n$. For example, in Figure 9, AS paths $(1, 2, 3)$ and $(1, 2, 6, 3)$ are no-valley paths while as_path $(1, 4, 3)$ and $(1, 4, 5, 3)$ are not no-valley paths.




Fig. 9. paths $(1, 2, 3)$ and $(1, 2, 6, 3)$ are no-valley path while AS paths $(1, 4, 3)$ and $(1, 4, 5, 3)$ are not no-valley path.

No-Valley-and-Prefer-Customer Routing Policy: In addition to the no-valley routing policy, an AS typically chooses a customer route over a route via a provider or peer since an ISP does not have to pay its customers to carry traffic, and tends to avoid the traffic congestion at peering exchange points. For example, in Figure 9, AS paths $(2, 3, 4)$ and $(2, 6, 3, 4)$ are no-valley paths, and AS2 will prefer the AS path $(2, 3, 4)$ via customer AS3 instead of peer AS6 to reach AS4.

Note that we choose the shortest AS path policy in the two routing policies since this can only underestimate AS path inflation by the routing policies. In reality, it is possible for ASs to have more complicated policies (e.g., for traffic engineering) and lead to additional AS path inflation. In the next section, we show the extent to which the no-valley routing policy inflates AS paths. In Section III-C, we show the extent to which the no-valley-and-prefer-customer routing policy inflates AS paths.

B. Path Inflation by No-Valley Routing Policy

In this section, we compare the AS path length with the shortest AS path length given the no-valley routing policy. We have a modified Dijkstra's algorithm for computing the shortest AS path among all no-valley paths. Figure 10 shows the algorithm in details. The algorithm takes $O(N^3)$ time to compute AS path length for all AS pairs, where N is the number of ASs in the AS graph. For detailed description of the algorithm, see [13].

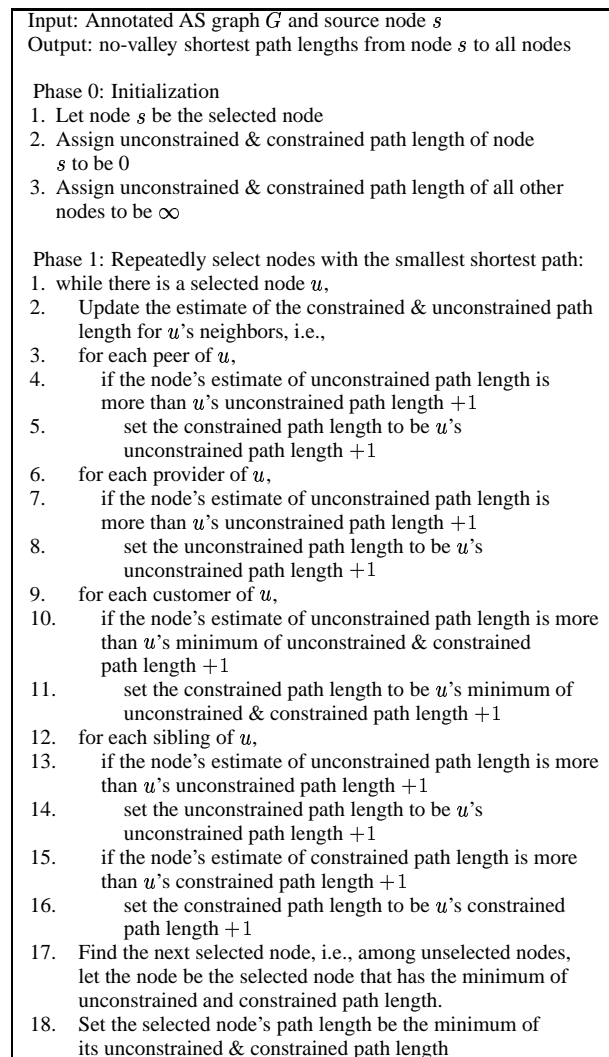


Fig. 10. Algorithm for computing no-valley path length

In Figure 11, around 4% of AS pairs have longer AS paths than the shortest AS paths. In Figure 12, We see that there is a small discrepancy between the shortest AS paths and derived

AS paths. This indicates that ASs typically employ more complicated routing policy than the no-valley routing policy.

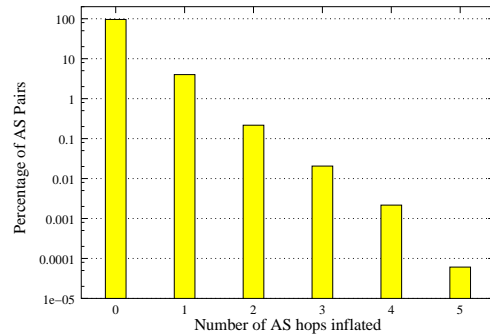


Fig. 11. Extent that the no-valley routing policy inflates AS paths for all AS pairs




Fig. 12. A comparison of AS paths derived from the no-valley policy and shortest AS paths for all AS pairs

C. Path Inflation by No-Valley-and-Prefer-Customer Routing Policy

We see that the no-valley routing policy does not inflate AS path significantly. This leads us to study a more sophisticated routing policy: no-valley-and-prefer-customer routing policy. Figure 13 shows an algorithm for computing AS path length given the no-valley-and-prefer-customer routing policies from all ASs to a destination s . For detailed description of the algorithm, see [13]. The algorithm takes $O(NE)$ time to compute AS path length for all AS pairs, where N is the number of ASs and E is the number of edges in the AS relationship graph. Note that this algorithm has lower complexity than the algorithm for no-valley path since it explores less paths than the algorithm for no-valley path due to the prefer customer policy.

For the sake of comparison, we compare the shortest AS paths and the AS paths derived from the no-valley-and-prefer-customer routing policy in Figure 14. We see that more than 45% of AS pairs have a longer AS path than the shortest AS path, and AS path can be inflated by as long as 9 AS hops. In Figure 15, we see that the AS paths derived from the no-valley-and-prefer-customer routing policy can be as long as 14 AS hops while the shortest AS paths are at most 10 AS hops.

IV. CONCLUSIONS AND FUTURE WORK

We performed measurement studies on AS path length and observed that AS paths are inflated significantly by interdomain

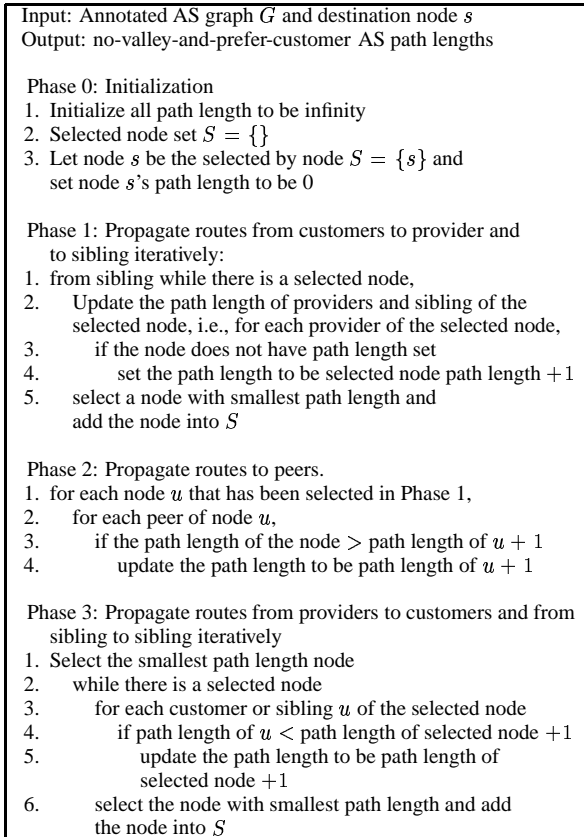


Fig. 13. Algorithm for computing no-valley-and-prefer-customer path length

routing policies. This leads us to systematically study the extent of AS path inflation by routing policies. We choose two typical routing policies to estimate the extent of AS path inflation for all AS pairs. We found that at least 45% of AS pairs choose a longer AS path than the shortest AS path. This study shows that the shortest AS path routing policy are not typical routing policies used in the current Internet, and AS path inflation is more prevalent than expected. As a part of our future study, we plan to understand how the chosen AS path differs from the AS path resulting from a typical routing policy such as no-valley-and-prefer-customer routing policy. This can give us insight into the routing policies configured in the Internet and the extent of traffic engineering performed in the Internet.

REFERENCES

- [1] Y. Rekhter and T. Li, "A border gateway protocol 4 (BGP-4)." Request for Comments 1771, March 1995.
- [2] J. W. Stewart, *BGP4: Inter-Domain Routing in the Internet*. Addison-Wesley, 1999.
- [3] G. Huston, "Interconnection, peering and settlements—Part I," in *Internet Protocol Journal*, March 1999.
- [4] G. Huston, "Interconnection, peering and settlements—Part II," in *Internet Protocol Journal*, June 1999.
- [5] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet routing convergence," in *Proc. ACM SIGCOMM*, August 2000.
- [6] C. Labovitz, R. Wattenhofer, S. Venkatachary, and A. Ahuja, "The impact of Internet policy and topology on delayed routing convergence," in *Proc. IEEE INFOCOM*, April 2001.




Fig. 14. AS path inflation resulted from the no-valley-and-prefer-customer routing policy for all AS pairs




Fig. 15. A comparison of AS paths derived from the no-valley-and-prefer-customer routing policy and shortest AS paths for all AS pairs

- [7] H. Tangmunarunkit, R. Govindan, D. Estrin, and S. Shenker, "The impact of routing policy on Internet paths," in *Proc. IEEE INFOCOM*, April 2001.
- [8] <http://www.anc.uoregon.edu/route-views/>.
- [9] National Laboratory for Applied Network Research, <http://moat.nlanr.net/AS/>.
- [10] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*. McGrawHill, 1990.
- [11] L. Gao, "On inferring autonomous system relationships in the Internet," in *Proc. IEEE GLOBAL INTERNET*, November 2000.
- [12] C. Alaettinoglu, "Scalable router configuration for the Internet," in *Proc. IEEE IC3N*, October 1996.
- [13] Lixin Gao, Feng Wang, "Technical Report", unpublished. <http://www-unix.ecs.umass.edu/lgao/globalinternet.ps>.
- [14] Daniel Golding, Sockeye, "Routing Policy Tutorial Audience and Goals," in NANOG24 Meeting, February 2002. <http://www.nanog.org/mtg-0202/golding.html>.