The Extent of AS Path Inflation Caused by Routing Policies

Lixin Gao, Feng Wang

Abstract— **Border Gateway Protocol (BGP) is an interdomain routing protocol that allows Autonomous Systems (ASes) to apply local policies for selecting routes and propagating routing information. These routing policies are typically constrained by the contractual commercial agreements between administrative domains. For example, an AS sets its routing policy so that it does not provide transit services between two of its providers. As a result, a route in the Internet may take a longer AS path than the shortest AS path. In this paper, we systematically analyze AS paths and quantify the extent that routing policies inflate AS paths in the Internet. Our results show that AS path inflation in the Internet are more prevalent than expected. We first present the extent of AS path inflation observed from the Route View routing tables. We find that from an ISP, at least 67% of AS paths are inflated by at least one AS hop and AS paths can be inflated by as long as 5 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 given a typical routing policy model. 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 (ASes) 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. ASes interconnect via dedicated links and public network access points, and exchange reachability information using Border Gateway Protocol (BGP) [1], [2]. BGP is an interdomain routing protocol that allows ASes to apply local policies for selecting routes and propagating routing information, without revealing their policies or internal topology to others. These routing policies are typically constrained by the contractual commercial agreements between administrative domains. For example, an AS sets its routing policy so that it does not provide transit services between two of its providers. In addition, an AS may set its policy so as to achieve traffic engineering or load balancing goal. It is well known that an AS may take a longer AS path than the shortest AS path as a result of these routing policies. However, the extent that 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, it is widely believed that routing policies may inflate routes used by the Internet traffic. However, 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 that 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 are 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 size: a tier-1 ISP, a tier-2 ISP, and a tier-3 ISP. From the tier-1 Internet service provider, more than 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 67% of chosen AS paths are longer than the shortest AS paths and AS paths can be inflated by as long as 5 AS hops. From the tier-3 ISP, more than 33% 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 more than 20% of AS pairs take a longer path than the shortest AS path and AS paths can be inflated by as long as 10 AS hops.

The analysis of BGP tables is limited to the view of the publicly available BGP routing tables. In order to understand the extent of AS path inflation for all AS pairs, we present two typical routing policies to show the extent that an AS path can be inflated by the routing policies. The first routing policy assumes that each AS obeys commercial agreements with its neighboring ASes. The commercial agreements between pairs of ASes can be classified into customer-provider, peering, and sibling [3], [4]. A provider does transit traffic for its customers or siblings. However, a customer does not transit traffic between two of its providers and peers. We present an algorithm for computing the AS paths that conform to this routing policy. Our results show that this routing policy inflates the AS path for only 4% AS pairs. The second routing policy

L.Gao and Feng Wang are with the Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA 01002. Email: {lgao, fewang}@ecs.umass.edu.

The authors are 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.

assumes that an AS prefers customer routes over provider or peer routes. We present an algorithm to compute the shortest path between a pair of ASes that conforms the second routing policy, and find that more than 45% of AS pairs use a longer AS path than the shortest AS path. In addition to study the extent of AS path inflation, we investigate how the chosen AS path differs from the AS path inferred by no-valley routing policy and no-valley-and-prefer-customer routing policy. This can give us insight on how to infer AS path used in the Internet through AS relationships and typical routing policies.

The study on the Internet path length [7] assumes that each AS chooses the shortest AS path routing policy. With this assumption, the authors conclude that 20% of Internet paths are inflated 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 BGP routing tables and explore routing policies that conform to commercial relationships. Spring et al.investigate the possible causes of router-level path inflation - topology and routing policy - and quantify the impact of each [14]. Our study focus on the impact of routing policy on AS level path inflation..

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

II. Background on Interdomain Routing

A. AS Graph

The Internet consists of a large collection of hosts interconnected by networks of links and routers, which is partitioned into thousands of autonomous systems (ASes). An AS has its own routers and routing policies, and connects to other ASes to exchange traffic with remote hosts. A router typically has very detailed knowledge of the topology within its AS, and limited reachability information about other ASes. Since we mainly concern ourselves with AS path length in this paper, connectivity between ASes is our focus. We model the connectivity between ASes in the Internet using an AS graph $G = (V, E)$, where the node set V consists of ASes and the edge set E consists of AS pairs that exchange traffic with each other. Note that the edges of AS graph represent logical connections between ASes and do not represent the form of the physical connection. Figure 1 shows an example of an AS graph.

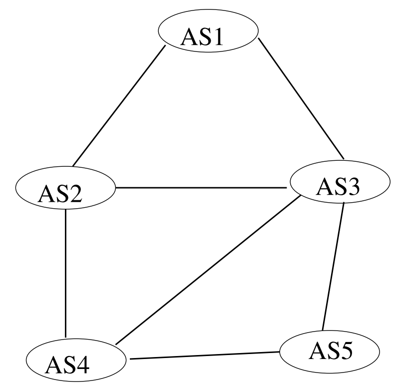

Fig. 1. *An AS graph example*

B. BGP

Each AS has responsibility for carrying traffic to and from a set of customer IP addresses. The scalability of the Internet routing infrastructure depends on the aggregation of IP addresses in contiguous blocks, called *prefixes*, each consisting of a 32-bit IP address and a mask length (e.g., $1.2.3.0/24$). An AS employs an *interdomain* routing protocol (BGP) to advertise the reachability of these prefixes to neighboring ASes. We denote the set of prefixes that are originated from AS u by $O(u)$.

BGP is a path-vector protocol that constructs paths by successively propagating updates between pairs of BGP speaking routers that establish *BGP peering sessions* [1], [2]. Each update r concerns a particular prefix, $r.prefix$, and includes the list of the ASes along the path (the *AS path*), $r.as\_path$. Each BGP speaking router originates updates for one or more prefixes, and can send the updates to the immediate neighbors via BGP sessions. The simplest path-vector protocol would employ shortest-path routing, where each AS selects a route with the shortest AS path. However, BGP allows a much wider range of routing policies such as policies that honor contractual agreements which control the exchange of traffic. Upon receiving an update, a router must decide whether or not to use this path according to *import policies* and, if the path is chosen, whether or not to propagate the update to neighboring ASes according to *export policies*. Let $import(u, v)$ represent import polices that AS v applies to a set of routes received from its neighboring AS u. Let $export(v, u)$ represent export policies that AS u applies to its best route set for sending to a neighboring AS v. Although routing policies are configured on each BGP speaking router, BGP speaking routers in an AS are typically configured with policies that are consistent with the AS's relationship with its neighboring AS. For the simplicity of exposition, we use an AS to represent BGP speaking routers in an AS and use the AS and its BGP speaking routers interchangeably throughout this paper. The routing policies are set by manipulating update attributes including next-hop interface address ($r.next\_hop$), local preference ($r.local\_pref$), multiple-exit discriminator ($r.med$), and community set ($r.c\_set$).

After applying the import policies for route updates from a BGP session, an AS saves all the imported updates in its BGP routing table. The AS then follows a route selection process that picks the best route for each prefix. Let $B(u, d)$ denote the best route selected by u for prefix d. $B(u, d)$ is selected by picking the route with the highest $local\_pref$, breaking ties by selecting the route with the shortest as_path. Note that local preference overrides the AS-path length. Among the remaining routes, the AS picks the one with the smallest med, breaking ties by selecting the route with the smallest intradomain routing cost. If a tie still exists, tie-breaking rules depend on vendor specific implementation.

C. Routing Tables

Each BGP speaking router keeps a *BGP routing table*, which stores a set of candidate routes for the router. We refer to a candidate route as a *routing table entry*, which contains a destination prefix, next-hop, med, local preference and AS path of the route. For the sake of simplicity, we describe the route table entries for a fixed prefix d. The routing table entry in AS u for destination d is a route with empty AS path, denoted as e(u,d), if u originates prefix d. Otherwise, the routing table entries in u for d depend on the best route of its neighboring AS v, $B(v,d)$, as well as the import policies of u from v and the export policies of v to u. Formally, the routing table entries of u for destination prefix d

$$Routing\_Table(u,d) = \begin{cases} e(u,d) & \text{if } d \in O(u) \\ \cup_{(u,v)\in E} import(v,u)[export(u,v)[B(v,d)]] & \\ & \text{otherwise} \end{cases} \quad (1)$$

We show a router's BGP routing table entries for destination prefix 4.2.24.0/21 below. The AS has five candidate routes to 4.2.24.0/21: AS path (1740 1) via next_hop 134.24.127.3, AS path (5459 5413 1) via next_hop 194.68.130.254, etc. Note that the third candidate route has AS path (1849 702 702 1), where AS 702 appears twice consecutively. This is due to AS prepending; AS 702 appends its AS number twice before exporting to AS 1849 to discourage AS 1849 to use the path via AS 702. From Equation (1), we see that a route table entry contains the best path chosen by an AS. For example, if $u_1, u_2, \ldots, u_n$ is a routing table entry for destination prefix d, then u_i chooses $u_{i+1}, \ldots, u_n$ as its best path to destination prefix d. In other words, for some prefix, u_i reaches u_n via AS path $u_{i+1}, \ldots, u_n$. This is the key fact we use in determining the chosen AS path length between an AS pair in the next section.

```
Network      Next Hop       Path
4.2.24.0/21  134.24.127.3   1740 1 i
             194.68.130.254 5459 5413 1 i
             158.43.133.48  1849 702 702 1 i
             193.0.0.242    3333 286 1 i
             144.228.240.93 1239 1 i
```

III. AS Path Inflation

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 ASes. To this end, we use routing tables from the Route View router in Oregon [8], which have 41 peering sessions. The Route Views server collects the BGP routing table once every night [9]. For the detailed description of the Route Views server, see [8]. Although there are other publicly available BGP routing tables, we choose Route View since it has the most peering sessions.

Using the BGP routing table from Route View server, we can construct an AS graph that characterize the connectivity on AS level. The shortest path between a pair of ASes can be derived 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 an overestimation of shortest AS path length.

The chosen AS path length between a pair of ASes is computed by the AS path appeared 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. For instance, if $P = u_1, u_2, \ldots, u_n$ is an AS path in a routing table entry, then the chosen AS path length between u_i and u_n on path P

$$Length(u_i, u_n, P) = |\{u_j | i \le j \le n\}| - 1$$

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. More formally, the chosen AS path length between u_i and u_n

$$ChosenLength(u_i, u_n) = \max_{\text{P is an AS path, ends with } u_n \text{ and contains } u_i} Length(u_i, u_n, P)$$

Note that since we have views from only 41 peering sessions in the routing table, only certain chosen paths from certain AS pairs are visible. For example, if an AS never appears at the end of an AS path – since it does not originate any prefix, we can not derive the chosen AS path length to this AS. If an AS never appears before the end of an AS path, we can not derive the chosen AS path length from this AS. Formally, we can derive chosen AS path length from u_i to u_n if and only if there is an AS path $\ldots, u_i, \ldots, u_n$ for some prefixes in the routing table.

We analyze routing tables during 2000 and 2001, and compute the chosen AS path length as well as the shortest AS path length. We present the results for 2000/12/11 only throughout this paper since the results for other dates are similar. To demonstrate the AS path inflation from different ASes, we choose three ISPs of various sizes: a tier-1 ISP, a tier-2 ISP and a tier-3 ISP. We first illustrate the extent of AS path inflation from these three ISPs. 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 ASes whose paths from the tier-1 ISP are inflated by a fixed number in Figure 2. The plot shows that more than 20% of ASes 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 3 shows the distribution of inflation with respect to the shortest AS path. The most upper curve shows the distribution of AS hops for the shortest AS path. Other curves show the distribution of a certain number of inflation with respect to the shortest AS path. We see that most AS path inflation occurs among 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 (the percentage of shortest paths with 6 AS hops is lower than 0.01% so that is not shown in Figure 3) to any AS in shortest path while it can be as long as 8 hops in chosen path.

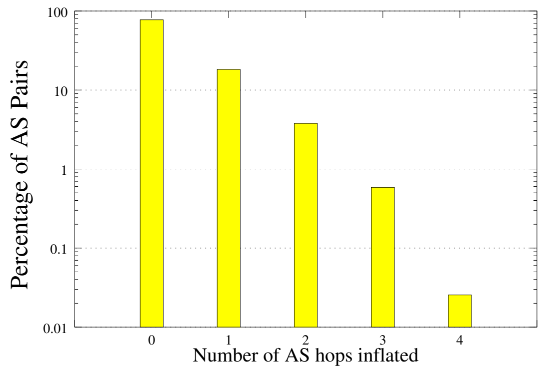

Fig. 2. *Percentage of ASes whose AS paths from tier-1 ISP are inflated by a fixed number*

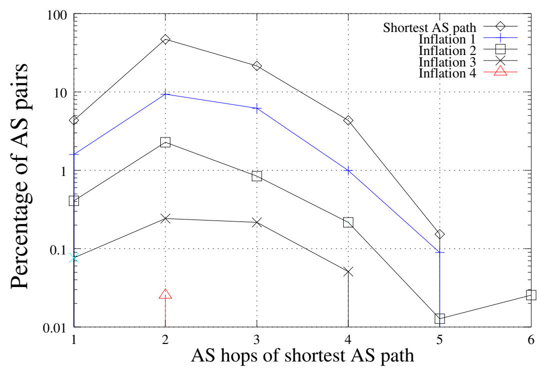

Fig. 3. *Distribution of AS path inflation from tier-1 ISP*

In Figure 4, more than 67% 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 5 shows that from the tier-2 ISP, most inflation occurs among AS path 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 that 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 does.

In Figure 6, more than 30% of AS pairs have a longer chosen AS path than the shortest AS path and AS paths can be inflated by as long as 5 AS hops. Note that although AS path inflation is more severe 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 7 shows that most AS path inflation occurs among 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.

Figure 8 shows more than 20% of AS pairs have a longer chosen AS path than the shortest AS path and inflation is 10 AS hops. In Figure 9. we see that chosen AS paths can be

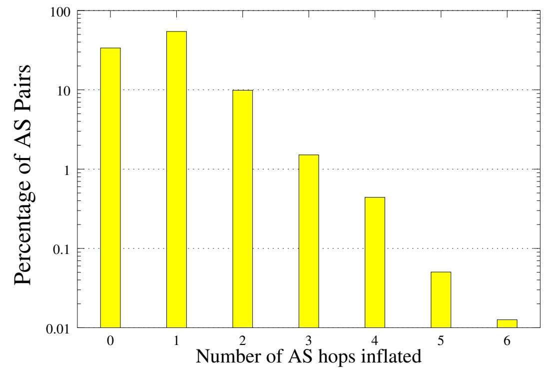

Fig. 4. *Percentage of ASes whose AS paths from tier-2 ISP are inflated by a fixed number*

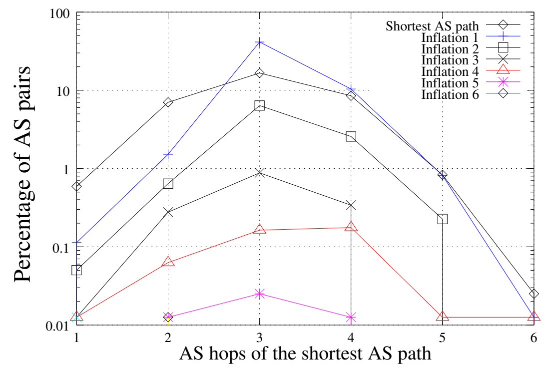

Fig. 5. *Distribution of inflation from tier-2 ISP*

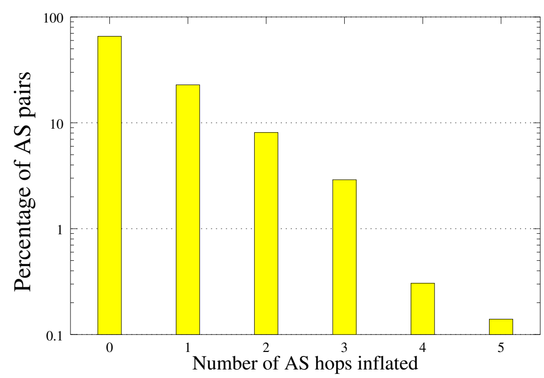

Fig. 6. *Percentage of ASes whose AS paths from tier-3 ISP are inflated by a fixed number*

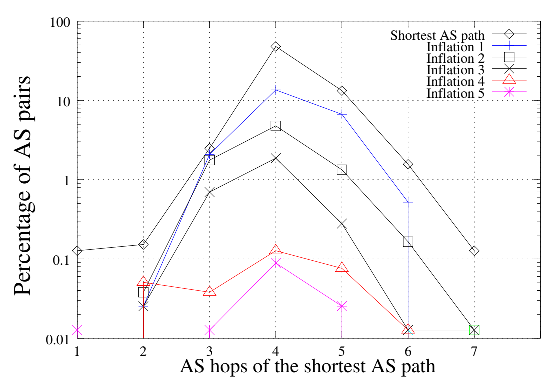

Fig. 7. *Distribution of inflation from tier-3 ISP*

an long as 13 hops while the shortest AS paths are at most 7 hops. Note that Figures 9 and 8 include only AS pairs whose chosen paths are visible from the Route View routing table. 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 as 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.

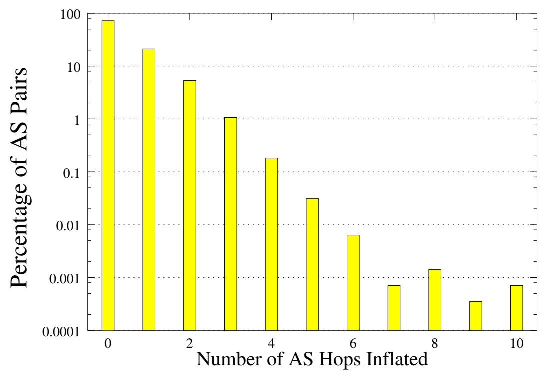

Fig. 8. *Percentage of AS pairs whose AS paths are inflated*

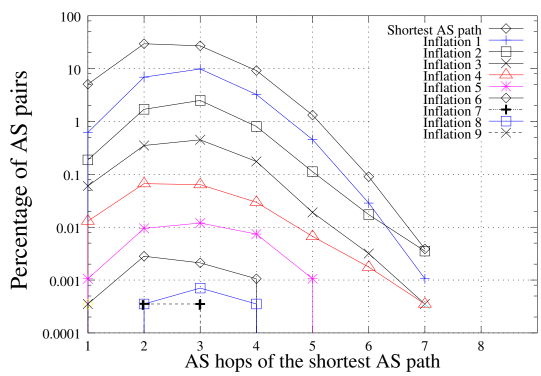

Fig. 9. *Distribution of inflation for all AS pairs whose chosen paths are visible from the Route View server*

IV. AS Path Inflation by Two Routing Policies

We can see that the extent of AS path inflation varies from ISPs to ISPs. To derive AS path inflation from BGP routing tables, we can only get a partial view of the extent of AS path inflation. In order to gain an overall picture of the AS path inflation, one has to know routing polices employed at each AS. Unfortunately, not all ISPs are willing to reveal their routing policies. To quantify the extent of AS path inflation in the Internet, we derive AS path inflation assuming two typical routing policies in this section. We first present the two routing policies that conform to commercial contractual agreements. We then proceed with the algorithms for deriving AS path length given the routing policies, and compare the derived AS path length with the shortest AS path length for all AS pairs. Finally, we compare the chosen AS path with AS path that conform to two typical routing policies to understand how we can infer real AS paths from AS relationships and typical routing policies.

A. *Two Routing Policies*

Routing policies typically conform to the relationships between ASes. AS relationships can be classified into provider-customer, peer-peer and sibling-sibling. 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 they have a mutual-transit agreement to provide connectivity to the rest of the Internet for each other.

The commercial contractual relationships between ASes translate into the export rule that an AS does not transit traffic between two of its providers and peers. Formally, we define $customer(a)$, $peer(a)$, $sibling(a)$, and $provider(a)$ as the set of customers, peers, siblings, and providers of a, respectively. We classify the set of routes in an AS into customer, provider, and peer routes. A route r of AS u is a *customer (provider, or peer) route* if the first non-sibling-sibling consecutive AS pair in $r.as\_path$ has a provider-customer (customer-provider, or peer-peer) relationship. More precisely, let $r.as\_path = (u_1, u_2, \ldots, u_n)$. If (u_i, u_{i+1}) has a sibling-sibling relationship for all $i < j$ and (u_j, u_{j+1}) has a provider-customer (customer-provider or peer-peer) relationship, then r is a customer (provider or peer) route. If (u_i, u_{i+1}) has a sibling-sibling relationship for all $i < n$. then r is defined to be u's own route. An AS selectively provides transit services for its neighboring ASes according to the following rule.

Selective Export Rule **Consider AS** u and AS $v \in provider(u) \cup peer(u)$.
For each best route r of u,
if r is a provider or peer route of u, then $export(v, u)[\{r\}] = \{\}$.

The AS relationships translate into the following rules that govern BGP export policies [12], [4], [13]:

The selective export rule translates into *no-valley* property of an AS path. Intuitively, if we image that provider is always above its customers and two peering ASes are at the same level, then once an AS path goes down or remains at the same level, it does not go up or remain at the same level. Formally,

No-valley: in a no-valley AS path $(u_1, u_2, \ldots, u_n)$, 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 11, 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. Note that the selective export rule ensures that BGP routing table entries contain only no-valley AS paths. For example, if AS path (1, 4, 3) appears in a BGP routing table, then AS 4 exports its provider route (3) to its provider AS 1. This violates the selective export rule.

In addition to the no-valley property, an AS typically chooses a customer route over a route via a provider or peer since an ISP does not have to pay its customer to carry traffic, and tends to avoid the traffic congestion at peering exchange points. Formally, we have a prefer-customer property for the

```
Algorithm for no-valley AS path:
Input: Annotated AS graph G and source node s
Output: no-valley shortest path lengths from node s to all nodes

Phase 0: Initialization
1.  Let node s be the selected node
2.  Assign unconstrained and constrained path length of node s to be 0
3.  Assign unconstrained and constrained path length of all other nodes to be ∞

Phase 1: Repeatedly select nodes with smallest shortest path:
1.  while there is a selected node u,
2.      Update the estimate of the constrained and unconstrained path length for u's neighbors, i.e.,
3.      for each peer of u,
4.          if the node's estimate of unconstrained path length is more than u's unconstrained path length +1
5.              set the constrained path length to be u's unconstrained path length +1
6.      for each provider of u,
7.          if the node's estimate of unconstrained path length is more than u's unconstrained path length +1
8.              set the unconstrained path length to be u's unconstrained path length +1
9.      for each customer of u,
10.         if the node's estimate of unconstrained path length is more than
            u's minimum of unconstrained and constrained path length +1
11.             set the constrained path length to be
                u's minimum of unconstrained and constrained path length +1
12.     for each sibling of u,
13.         if the node's estimate of unconstrained path length is more than u's unconstrained path length +1
14.             set the unconstrained path length to be u's unconstrained path length +1
15.         if the node's estimate of constrained path length is more than u's constrained path length +1
16.             set the constrained path length to be u's constrained path length +1
17.     Find the next selected node, i.e.,
        among unselected nodes, let the node that has the smallest
        minimum of unconstrained and constrained path length be the selected node.
18.     Set the selected node's path length be the minimum of its unconstrained and constrained path length
```

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

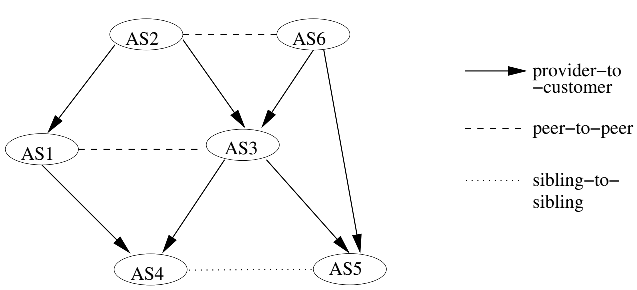

Fig. 11. *AS paths (1, 2, 3) and (1, 2, 6, 3) are no-valley path while AS paths (1, 4, 3) and (1, 4, 5, 3) are not no-valley path.*

import policy of AS a:

Prefer-customer: If $first(r_1.as\_path) \in customer(a)$ and $first(r_2.as\_path) \in priv\_peer(a) \cup provider(a)$, then $r_1.loc\_pref > r_2.loc\_pref$.

Therefore, we have two typical routing policies.

- **No-Valley Routing Policy**: All ASes follow the selective export rule and therefore, all chosen AS paths are no-valley paths. Furthermore, all chosen AS paths are the shortest AS path among all no-valley AS paths.
- **No-Valley-and-Prefer-Customer Routing Policy**: In addition to the no-valley routing policy, all ASes follow the prefer-customer import policy. Furthermore, all chosen AS paths are the shortest AS path among all no-valley-and-prefer-customer AS paths.

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

B. AS relationship Inference

AS relationships are the key in determining AS paths given the two routing policies. In order to define , we need to understand AS relationships in the Internet. To derive AS relationships, we use the algorithm presented in [11]. The main heuristics of the algorithm are that an AS size is proportional to its degree in the AS graph, a provider is larger than its customer, and two peering ASes are of comparable sizes. The result of the heuristic algorithm is an *annotated AS graph* – a partially directed graph whose nodes represent ASes and whose edges are classified into provider-to-customer, customer-to-provider, peer-to-peer and sibling-to-sibling edges. Furthermore, only edges between providers and customers in an annotated AS graph are directed. Figure 11 shows an example of an annotated AS graph.

C. Path Inflation by No-Valley Routing Policy

In this section, we first present an algorithm for deriving AS path length given the no-valley routing policy. We then compare the AS path length with the shortest AS path length.

1) Algorithm for computing AS path lengths given the no-valley routing policy: We have a modified Dijkstra's algorithm for computing the shortest AS path among all no-valley paths. Dijkstra's algorithm computes the single-source shortest-paths problem on a graph. The idea of Dijkstra's algorithm is to maintain a set of vertices, S, whose final shortest path lengths from the source have already been determined. The algorithm also maintains estimates of the shortest path lengths for the rest of vertices given S and repeatedly selects the next vertex with the minimum estimate of the shortest path length. The selected vertex is added to set S and its final shortest path length is its estimate of shortest path length.

To compute the shortest path length among all no-valley paths from a single source, we have the following modified Dijkstra's algorithm. We maintain estimate of two shortest path lengths. One shortest path length is called *constrained path length*, which is the length of a path that has a provider-to-customer edge or peer-to-peer edge and that can only be extended by sibling-to-sibling or provider-to-customer edges. Another shortest path length is called *unconstrained path length*, which is the length of a path that does not contain provider-to-customer or peer-to-peer edges and that can be extended by any edge. To observe the no-valley routing policy, we keep track of both constrained and unconstrained path lengths for further extension and extend a path using the following rule;

- A peer-to-peer edge can extend only unconstrained path. The resulting extended path is a constrained path.
- A customer-to-provider edge can extend only unconstrained path. The resulting extended path is an unconstrained path.
- A provider-to-customer edge can extend both unconstrained and constrained path. The resulting extended path in both cases are constrained paths.
- A sibling-to-sibling edge can extend both unconstrained and constrained paths. An unconstrained path extended by a sibling-to-sibling edge is an unconstrained path. A constrained path extended by a sibling-to-sibling edge is a constrained path.

We estimate the AS path length for each node using the minimum value of the constrained path length and the unconstrained path length. We finalize the AS path length for the node with the smallest AS path length by selecting the node. 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 ASes in the AS graph.

2) AS path inflation resulted from the no-valley routing policy: To see the extent of AS path inflation by the routing policy, we plot the AS path inflation for all AS pairs in Figure 12. Around 4% of AS pairs have longer AS paths than the shortest AS paths. For the sake of comparison, we present the number of AS pairs that are a fixed number of AS hops away from each other – for both the shortest AS paths and the AS paths derived from the no-valley routing policy – in Figure 13. We see that there is small discrepancy between shortest AS paths and derived AS paths. This indicates that ASes typically employ more complicated routing policy than the no-valley routing policy.

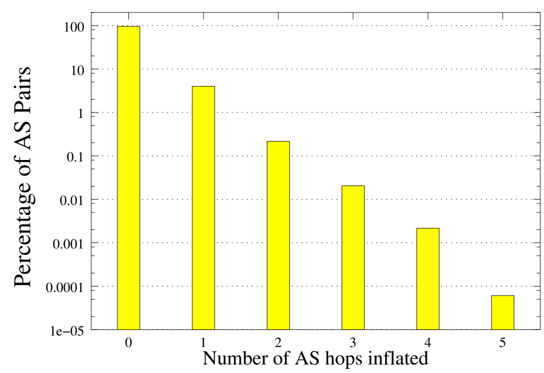

Fig. 12. *Extent that the no-valley routing policy inflates AS paths for all AS pairs*

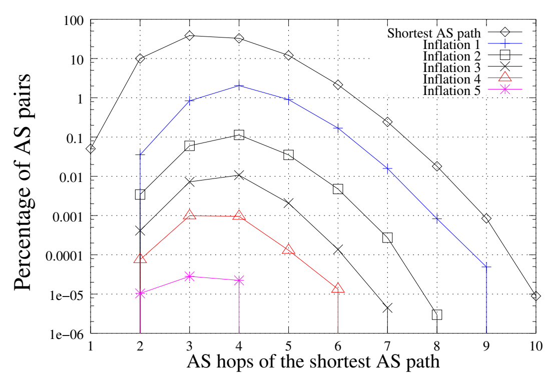

Fig. 13. *A comparison of AS paths derived from the no-valley policy and shortest AS paths for all AS pairs*

D. 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. In this section, we assume that each AS chooses an AS path that conforms to no-valley-and-prefer-customer routing policy.

1) Algorithm for computing AS path length given no-valley-and-prefer-customer routing policy: Figure 14 shows an algorithm for computing AS path length given the no-valley-and-prefer-customer routing policies from all ASes to a destination s. The idea of the algorithm is as follows. First, s propagates routes from customers to providers and from sibling to sibling. Note that this phase derives the customer routes if an AS has a customer route to s. Second, all nodes that derive a route from the first step propagate routes to peers. Third, all nodes that derive a route from the first two phases propagate routes to customers and sibling iteratively. In each phase, we finalize the AS path length for the node with the smallest AS path length by selecting the node. Once a node is selected, its AS path length is finalized and will not be a candidate for selection again.

Note that the no-valley policy ensures that the routes derived from the second and third phases will not propagate from a customer to its provider, and therefore, the routes derived after the three phases are the AS paths that conform to the no-valley-and-prefer-customer routing policy. The algorithm takes $O(NE)$ time to compute AS path length for all AS pairs, where N is the number of ASes 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.

2) AS path inflation resulted from the 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.

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 15. We see that more than 45% of AS pairs have a longer AS path than the shortest AS path, and inflation can be 9 AS hops. In Figure 16, 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.

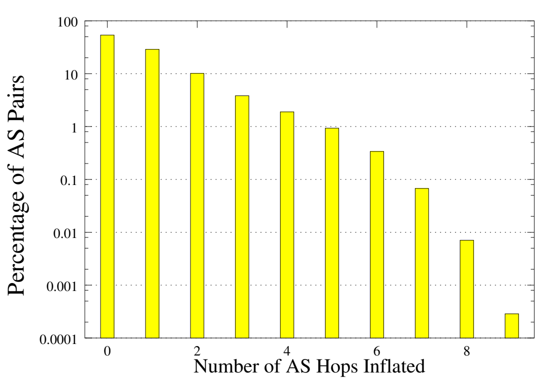

Fig. 15. *AS path inflation resulted from the no-valley-and-prefer-customer routing policy for all AS pairs*

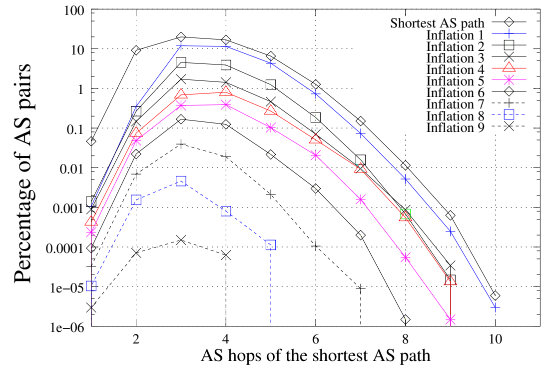

Fig. 16. *A comparison of AS paths derived from the no-valley-and-prefer-customer routing policy and shortest AS paths for all AS pairs*

E. Comparing Chosen AS Paths with AS Paths resulted from routing policy

In this section, we study how the chosen AS paths differs from the AS paths resulted from a no-valley routing policy and no-valley-and-prefer-customer policy. For a pair of ASes, we compare the chosen AS path, which is from BGP routing table, with AS path inferred by no-valley or no-valley-prefer-customer routing policy.

We select 7 tier-1 ASes, 4 tier-2 ASes, and 5 tier-3 ASes, shown in Table I, to compare their chosen paths and inferred path. We find that for all selected ASes, about 70% of chosen AS paths are the same as the paths inferred by no-valley routing policy. However, for tier-1 ASes and tier-2 ASes, only a small number of chosen AS paths are consistent with paths inferred by no-valley-and-prefer-customer routing policy, The reason is that tier-1 and tier-2 ASes may violate prefer-customer routing policies. For example, in Figure 17. AS 2 and AS 3 are tier-1 ASes. If AS 5 advertises its prefix to AS 4 and AS 3, AS 2 will learn two routes to the destination from AS 3 and AS 5, respectively. According to no-valley-and-prefer-customer routing policy, the AS path between AS2 and AS 5 is path (2,4,5). On the contrary, if AS 5 does not advertise the route to AS 4, AS 2 has only one route from AS 3 and this route is chosen as the best one. As a result, the chosen path (2 3 5) at AS 2 is not the same as the path (2 4 5) inferred by routing policy.

For tier-3 or stub ASes, they do not violate prefer-customer policy because they may have only routes from providers. For example, in Figure 17, AS 1 has only one route from AS 2. The chosen path (1 2 3 5) can be inferred by the no-valley-and-prefer-customer routing policy, even though AS 2's chosen path cannot be inferred by the policy. Therefore, for tier-3 ASes and stub ASes, more chosen AS paths are the same as those inferred by no-valley-and-prefer-customer routing policies. Note that the study [15] is similar to our work, but we use well-known routing policy to infer AS path.

V. Conclusions and Future Work

We performed measurement studies on AS path length and observed that AS paths are inflated significantly by interdomain routing policies. This leads us to systematically study the extent of AS path inflation by routing policies. We choose

Optimal algorithm for no-valley-and-prefer-customer AS path:
Input: Annotated AS graph G and destination node s
Output: no-valley-and-prefer-customer AS path lengths from all nodes to node s

Phase 0: Initialization
1. Initialize all path length to be infinity
2. Selected node set $S = \{\}$
3. Let node s be the selected by node $S = \{s\}$ and set node s's path length to be 0

Phase 1: Propagate routes from customers to provider and from sibling to sibling iteratively:
1. while there is a selected node,
2. Update the path length of providers and sibling of the selected node, i.e., for each provider of the selected node,
3. if the node does not have path length set
4. set the path length to be selected node path length $+1$
5. select a node with smallest path length and add the node into S

Phase 2: Propagate routes to peers.
1. for each node u that has been selected in Phase 1,
2. for each peer of node u,
3. if the path length of the node $>$ path length of $u + 1$
4. update the path length to be path length of $u + 1$

Phase 3: Propagate routes from providers to customers and from sibling to sibling iteratively
1. Select the smallest path length node
2. while there is a selected node
3. for each customer or sibling u of the selected node
4. if path length of $u <$ path length of selected node $+1$
5. update the path length to be path length of selected node $+1$
6. select the node with smallest path length and add the node into S

Fig. 14. *Algorithm for computing no-valley-and-prefer-customer path length*

TABLE I
PERCENTAGE OF CHOSEN AS PATHS CONSISTENT WITH AS PATHS INFERRED BY TWO ROUTING POLICIES.

AS number	tier	No-valley	Prefer-customer
AS 1	tier-1	77%	3.3%
AS 293	tier-1	69%	5%
AS 1239	tier-1	79%	9%
AS 3356	tier-1	88%	3%
AS 3549	tier-1	77%	0.4%
AS 3561	tier-1	72%	2%
AS 7018	tier-1	74%	15%
AS 267	tier-2	77%	33%
AS 1740	tier-2	70%	16%
AS 1755	tier-2	60%	11%
AS 2497	tier-2	75%	12%
AS 6395	tier-3	72%	23%
AS 2551	tier-3	72%	33%
AS 715	tier-3	71%	23%
AS 2685	tier-3	80%	26%
AS 234	tier-3	77%	33%

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 routing policies are more prevalent than expected. As a part of our

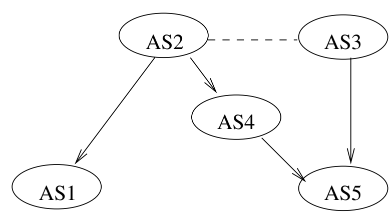

Fig. 17. *The AS path (2,4,5) inferred by no-valley-and-prefer-customer routing policy is not the chosen AS path (2,3,5). The AS path (1,2,4,5) inferred by no-valley-and-prefer-customer routing policy is the same as the chosen AS path.*

future study, we plan to understand how the chosen AS path differs from the AS path resulted from a typical routing policy such as no-valley-and-prefer-customer routing policy.

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.
[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.antc.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] Daniel Golding, Sockeye. Routing Policy Tutorial Audience and Goals. In NANOG24 Meeting, February 2002. http://www.nanog.org/mtg-0202/golding.html
[14] N. Spring, R. Mahajan, and T. Anderson. Quantifying the causes of path inflation. In ACM SIGCOMM, 2003.
[15] Z. M. Mao, L. Qiu, J. Wang, and Y. Zhang. On AS-level Path Inference. In Proceedings of ACM SIGMETRICS, June 2005.