On Inferring and Characterizing Internet Routing Policies

Feng Wang and Lixin Gao

Abstract: **Border gateway protocol allows autonomous systems (ASes) to apply diverse routing policies for selecting routes and for propagating reachability information to other ASes. Although a significant number of studies have been focused on the Internet topology, little is known about what routing policies network operators employ to configure their networks. In this paper, we infer and characterize routing policies employed in the Internet. We find that routes learned from customers are preferred over those from peers and providers, and those from peers are typically preferred over those from providers. We present an algorithm for inferring and characterizing export policies. We show that ASes announce their prefixes to a selected subset of providers to perform traffic engineering for incoming traffic. We find that the selective announcement routing policies imply that there are much less available paths in the Internet than shown in the AS connectivity graph, and can make the Internet extremely sensitive to failure events. We hope that our findings will help network operators in designing routing policies.**

Index Terms: **Export routing policies, import routing policies, inter-domain routing protocol, routing policies.**

I. INTRODUCTION

The Internet connects thousands of autonomous systems (ASes) operated by many different administrative domains such as Internet service providers (ISPs), companies and universities. Routing information is exchanged using the interdomain routing protocol, border gateway protocol (BGP) [1], and routing within an AS is performed using an intradomain routing protocol such as IS-IS or OSPF. A key feature of BGP is that it allows ASes to adopt diverse routing policies to control the selection of routes and propagate reachability information to other ASes. For example, a multihomed AS can control the in-bound traffic link by propagating prefixes to a subset of its providers only. Therefore, the prefix can be reached only through the subset of its providers. As a result, connectivity does not mean reachability in the Internet and the extent of the reachability is determined by both connectivity and routing policies. Although a significant number of studies have been focused on the Internet topology [2], [3]–[5], little is known about what routing policies network operators employ to configure their networks.

Understanding routing policies applied in the Internet has several implications. First, it is important to have a global view of the routing policies applied. Clearly, each ISP has information about its own routing policies. However, many ASes are unwilling to reveal their routing policies to others. Furthermore, the routing information stored in Internet routing registry (IRR) [6], which maintains ASes' routing information in several public databases, is either incomplete or out-of-date. Therefore, there is no global view of the typical routing polices configured in an AS. Second, the global view of routing policies might have implications on important properties of Internet. The connectivity in the Internet does not mean reachability since routing polices might lead to less available paths. Moreover, this can lead to implications on robustness of the Internet. Third, being able to infer routing polices of other ASes might allow an AS to perform traffic engineering effectively. To control traffic flow, network operators can change their routing policies to shift traffic load among multiple candidate routing paths. This task can be performed if candidate routing paths can be predicted by inferring routing policies of ASes involved.

In this paper, we first infer and characterize import routing policies. In particular, we infer the route preference setting among routes learned from providers, customers and peers. From a large collection of routing tables, we find that in most cases, route preference conforms to AS relationships. That is, routes learned from customers are typically preferred over those from providers or peers, and routes learned from peers are typically preferred over those from providers. Second, we present an algorithm for inferring export policies and characterize the export policies. We infer how an AS announces its routes to its (direct or indirect) providers. Our results show that a significant number of ASes announce their prefixes to a selected subset of providers. The major reason is due to the traffic engineering practice for controlling incoming traffic. That is, an AS may announce its prefixes only to a subset of its direct providers, or to its direct providers. Although the selective announcement routing policies are not surprising, the impact of these routing policies might be significant. The selective announcement routing policies imply that there are much less available paths in the Internet than shown in the AS connectivity graph. Therefore, the adoption of the selective announcement routing policies can make the Internet extremely sensitive to failure events. We hope that our findings will help network operators in designing routing policies.

The rest of this paper is organized as follows. In Section II, we describe routing policies. Section III presents our data source. In Section IV, we describe our methodology for inferring import policies and the characteristics of the import policies. Then, in Section V, we present algorithms for inferring export policies and characterize the export policies. We conclude the paper with a summary in Section VI.

Manuscript received May 16, 2007.

This paper was adapted from "On inferring and characterizing Internet routing policies" in ACM SIGCOMM IMC 2003. This work was partially supported by NSF grants CNS-0626617, CNS-0208116, CNS-0325868.

F. Wang is with School of Engineering Computational & Sciences at Liberty University, Lynchburg, VA 24502 USA, email: fwang@liberty.edu.

L. Gao is with Department of ECE, University of Massachusetts, Amherst, MA 01003 USA, email: lgao@ecs.umass.edu.

1229-2370/07/$10.00 © 2007 KICS

II. ROUTING POLICIES

Routing policies are a set of rules that are configured by network operators. Routing policies include *import policies* and *export policies.*

A. Import Routing Policies

From each neighbor, a router receives a set of route announcements. In order to distinguish routes from different neighbors, we define a route received from a customer as *customer route*, and the AS path the route traversed as *customer path.* We define a route received from a provider as *provider route*, and the AS path the route traversed as *provider path.* A route received from a peer is defined as *peer route*, and the AS path the route traversed is defined as *peer path.* We will use those terms to help us infer routing policies throughout this paper.

After receiving a route announcement, a BGP router discards the route if its own AS number is present in the AS path to avoid a loop in AS path. The router then applies import policies to this route, which include denying, or permitting a route, and assigning a local preference to indicate how favorable the route is. The preference is a value used to rank routes received.

B. Export Routing Policies

After selecting the best route, a BGP router will propagate only the best route to its neighboring ASes. Export policies allow the router to determine whether to advertise the best route to a neighbor. Export policies include permitting or denying a route, assigning MED to control the inbound traffic, tagging a BGP community to indicate what preference a neighboring AS should assign to the route, or prepending AS paths.

The following rules are well known BGP export policies [7]:

- **Exporting to provider**: A customer can export to its providers its routes and the routes learned from its own customers, but cannot export routes learned from other providers or peers.
- **Exporting to customer**: A provider can export to its customers its routes, the routes learned from the other customers, its providers, and its peers.
- **Exporting to peer**: A peer can export to another peer its routes, the routes learned from its customers, but cannot export the routes learned from its providers and other peers.

The constrain that a customer cannot export routes learned from other providers or peers and a peer cannot export the routes learned from its providers and other peers is called *no-valley* routing policy. We note that route selection of an AS depends on the routes coming from its neighboring ASes, and on the import policies of the AS.

III. DATA SOURCES

In order to infer routing policies, we need to analyze BGP routing tables, which provide us with routes and associated attributes such as local preference value and AS paths. However, routing policies applied by one AS may be different from routing policies applied by another AS. Therefore, we conduct our analysis using BGP routing tables from routeview servers, looking glass servers, and routing information from the IRR.

Oregon RouteView provides a view of the global routing system from the perspectives of several different backbones and locations around the Internet [8]. On Nov. 2002, it peered with 56 ASes that announce their routes to it. Those ASes include nearly all Tier-1 ASes in the Internet, such as AS 1239 (Sprint) and AS 7018 (AT&T). Looking Glass servers from 15 ASes [9] provide BGP routing information, including 3 Tier-1 ASes. Through those Looking Glass servers, we can retrieve fine-grained routing information, such as Local Preference, and BGP community. All BGP tables from RouteView servers and Looking Glass servers are downloaded from Nov. 11, 2002 to Nov. 18, 2002. Overall, we capture 71 BGP tables, 45 located in North America, 21 located in Europe, 3 located in Australia, and 2 in Asia.

The IRR maintains ASes' routing information in several public databases to coordinate global routing policies. We downloaded public IRR database files mirrored at [10] on Nov. 25, 2002, which include routing information for 62 ASes.

Our overall data comes from 133 ASes, and includes routing tables from most Tier-1 ASes. We believe that those data sources are sufficient for our study, and we will discuss the representativeness of our data for inferring routing policy in Section IV and Section V .

Our study relies on AS relationships. There are several novel algorithms which can be used to infer AS relationships [11]–[13]. Here, we choose the one described in [11]. In Section IV-B, we show that the potential error introduced by inferred AS relationships is small.

IV. INFERRING IMPORT POLICIES

One of the most important aspects of import policies is to set local preference. In this section, we first describe two methods to infer import policies. And then, we investigate the potential errors of our method and the representativeness of our data.

A. Route Preference Among Different Routes

Local preference can be used to influence the selection of the best route among a set of routes, and control outgoing traffic. Network operators usually assign different local preference values to customer, provider, and peer routes. We use BGP routing tables from RouteView servers and Looking Glass servers, and information from IRR to infer import policies.

A.1 Inferring Import Policies from BGP Routing Tables

Given a routing table from AS u, we infer the local preference between the given AS u and its neighbors according to the following steps:

1. Extract local preference value for each route. We use $r_i.lpf$ to represent the local preference value for route r_i.
2. Cluster local preference values according to the nexthop attribute of each route (represented by $r.nexthop$), $L = (r_1.lpf, r_2.lpf, \cdots, r_k.lpf)$ and $r_i.nexthop = v$.
3. Infer AS relationship between the given AS u and v, $REL(u, v) = (customer, peer, provider)$.
4. Compare local preference setting among AS u's neighbors.

Hence, after knowing AS relationships between an AS and its neighbors and deriving local preference, we associate each

Table 1. Typical local preference assignment for 15 ASes. It shows the prevalence of typical local preference.

AS number	% of typical local preference	AS number	% of typical local preference
AS 577	94.3	AS 2578	99.9982
AS 5511	96.5	AS 513	100
AS 3549	99.7	AS 6762	100
AS 6667	99.94	AS 559	100
AS 7474	99.955	AS 12859	100
AS 12359	99.98	AS 8262	100
AS 7018	99.99	AS 6539	100
AS 1	99.994		

neighbor, or customer, provider and peer, with one or more local preference values. To compare local preference values among different routes, we define:

- **Typical local preference**: Routes from customers have higher local preference than those from peer and provider, and routes from peer have higher local preference than routes from provider.
- **Atypical local preference**: The local preference of routes from peer or provider is not lower than that of customer routes, or the local preference of routes from provider is not lower than that of peer routes.

Table 1 shows the percentage of prefixes which have typical local preference for each AS. Our result implies that the percentage of atypical local preference for each AS is very small. Those 15 ASes include 3 Tier-1 ASes (AS1, AS3549, and AS7018), 2 Tier-2 ASes (AS5511 and AS7474), and other 10 ASes.

A.2 Inferring Import Policies from IRR

In order to get a more complete view of local preference setting, we resort to the IRR to infer import policies. The motivation of IRR is to coordinate global routing policies, but the IRR database may not be complete and some part of it can be out-of-date. We check each AS's last update time and discard those ASes which are older than 2002.

The IRR database expresses routing information at various levels (e.g., individual prefix or AS, etc.). The following example shows how the import policy is expressed in routing policy specification language (RPSL).

```
aut-num: AS1
import: from AS2 action pref = 1; accept ANY
```

Policy *actions* in RPSL can assign a preference to a route. This example states that all routes are accepted from AS 2 with preference 1. In router configuration the preference can be done by setting a local preference.[1]

However, we cannot infer AS relationships for some ASes shown in IRR because they do not appear in Oregon BGP routing table, which we use to infer AS relationship. Therefore, we only consider those ASes that their AS relationships can be inferred and have more than 50 neighbors. We infer the typical local preference for 62 ASes from IRR, shown in Table 2. Even though those ASes (15 ASes from BGP tables, and 62 ASes from IRR) are a small fraction of ASes in the Internet, we believe that the chosen ASes are representative for studying import policies in the Internet. Therefore, we conclude that local preference value for a customer is typically higher than for a provider and peer, and that local preference for a peer is higher than that for a provider.

B. Potential Error Introduced by Inferred AS Relationships

Since studying routing policies relies on AS relationships, a large number of ASes with incorrectly inferred AS relationships will affect our conclusion about import policies. We use BGP community to verify inferred AS relationships. One of the most common usages of community values is to tag the routes received from specific neighbor ASes. In this case, an AS defines different community values for its customers, peers, and providers. When border routers of the AS receive a route from its neighbors, they tag the route with a community indicating the relationship with those neighboring ASes. Details about this method are described in work [14].

Table 3 shows that the AS relationships between 9 ASes and their neighboring ASes are verified. As shown in the table, for those 9 ASes, most of their AS relationships are correctly inferred. Since we use BGP community to verify AS relationships, some neighbors of those 9 ASes do not have BGP community so that their relationships cannot be verified. Therefore, in Table 3, AS relationships that cannot be verified does not mean the inferred AS relationships incorrect. Hence, the potential error introduced by inferred AS relationship is so small that it will not affect our results.

C. Representativeness of Date Sources for Inferring Import Policies

To understand the representativeness of our data for inferring import policies, we investigate which ASes are likely to violate typical local preference. ASes can be classified into different tiers based on their position in the Internet hierarchy. Until Nov. 18, 2002, there are 14202 ASes derived from Oregon RouteView. About 83% of ASes are *stub* ASes, and 17% are *transit* ASes, such as Tier-1 and Tier-2 ASes,.

From the aspect of reachability to the Internet, Tier-1 ASes, Tier-2 ASes, and stub ASes have different tendency to violate typical local preference. That is, it is *possible* for Tier-1 and Tier-2 ASes to violate typical local preference without affecting their neighbors' reachability to the destination. On the contrary, it is *impossible* for stub ASes to violate typical local preference.

Tier-1 ASes are those ASes that do not have providers so that the available routes in Tier-1 ASes' routing tables are from customers or peers. For example, if Tier-1 AS A prefers the route learned from its peer (Tier-1 AS B), instead of the route from its customer, AS A cannot advertise the route to another Tier-1 AS C. On the other hand, Tier-1 ASes are fully connected. That means that AS C will have one route learned from AS B even though AS C does not have a route from AS A. In this case, the violation of typical local preference does not affect reachability of Tier-1 AS C.

Let's look at Tier-2 ASes that have routes from customers, peers and providers. We give an instance to illustrate that the violation of typical local preference at Tier-2 ASes does not affect reachability of the Internet. Suppose a Tier-2 AS A has two Tier-1 providers, $T1$ and $T2$, and one customer, C. AS A learns

---

[1] Preference is opposite to local preference in that the smaller values are preferred over larger values.

Table 2. Typical local preference assignment for 62 ASes (ASes are sorted according to their AS degree in non-decreasing order) that are selected from IRR. It shows the prevalence of typical local preference.

AS number	% of typical local preference	AS number	% of typical local preference
AS 12635	100	AS 5611	98
AS 15498	100	AS 8608	100
AS 4004	99.86	AS 12306	93.5
AS 6863	99.90	AS 5400	100
AS 12322	99.92	AS 3215	100
AS 12779	100	AS 3300	94.7
AS 12626	99.94	AS 1740	100
AS 2518	100	AS 8341	100
AS 8650	91.66	AS 293	83.2
AS 20646	100	AS 6705	80
AS 5539	89	AS 8434	100
AS 5615	100	AS 12390	98
AS 12573	96	AS 5607	95
AS 1140	98	AS 5427	99
AS 6873	100	AS 4000	100
AS 12781	98	AS 1901	97
AS 8365	100	AS 15290	100
AS 852	100	AS 3320	83
AS 8527	100	AS 13127	93
AS 5551	100	AS 9191	100
AS 3313	97.8	AS 5466	94
AS 12731	97.8	AS 5597	98
AS 15435	98.9	AS 6453	100
AS 3216	100	AS 12868	99
AS 2118	88.6	AS 5594	96
AS 1103	88.9	AS 13129	99.23
AS 21392	100	AS 6830	100
AS 9013	96.9	AS 1299	99.1
AS 5571	98	AS 3292	86
AS 3344	90.4	AS 4513	100
AS 5503	98	AS 3561	99.46

Table 3. The AS relationships between 9 ASes listed below and their neighbors are verified.

AS number	# of neighbors	Percentage of AS relationships between AS and its neighbors verified
AS 1	599	95.65
AS 577	89	98.9
AS 3549	558	96.28
AS 5511	168	99.4
AS 6539	157	96.45
AS 6667	26	97.46
AS 7018	1330	99.55
AS 12359	31	94.1
AS 12859	109	98.2

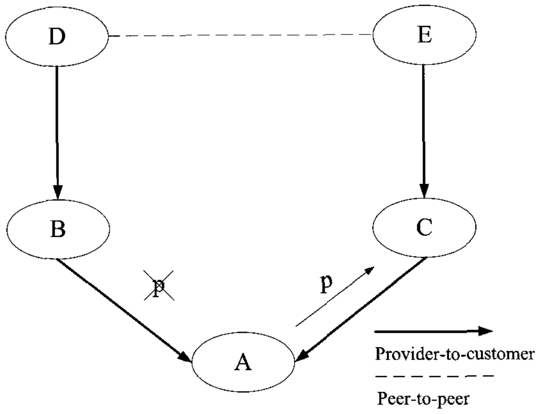

Fig. 1.The selective announcement routing policies employed by its customers can be observed at provider D. Customer A announces prefix p to provider C but not to B. In the BGP table of provider D, prefix p is received from its peer E.

two routes to a destination in C from provider $T1$ and customer C, respectively. If AS A prefers the route from provider $T1$, it cannot transit the route to another provider $T2$ due to no-valley policy. However, Tier-1 ASes $T1$ and $T2$ are fully connected. $T2$ still has one route from $T1$ even though $T2$ does not have the route from Tier-2 AS A. As a result, $T2$ still can reach the destination.

Stub ASes that have single or multiple providers cannot violate typical local preference. The reason is because stub ASes have routes only from its providers to reach the rest of the Internet. Therefore, the import policies for stub ASes is to select the route from one of providers as the best route and advertise it, so that the Internet still can reach the destination.

From our data, majority of ASes are stub ASes. That means, we do not need to infer those stub ASes' import routing polices. Our data sources come from Tier-1, Tier-2, and Tier-3 ASes, representing 3.2% of transit ASes. However, we infer import policies applied by most Tier-1 ASes' and many Tier-2 or Tier-3 ASes. We believe that the chosen ASes are representative for studying import policies in the Internet.

V. INFERRING EXPORT POLICIES

One important component of export policies is whether an AS announces its prefixes to its customers, peers, or providers. A provider has obligation to announce all of its prefixes, or default routes to its customers according to their agreements. However, a customer may advertise its prefixes to either all of its providers, or a subset of providers. In the latter case, customers can control their inbound traffic on a heavy traffic link by switching announcements of some prefixes away from the link. Peers also have control over their prefix announcements to other peers. In this paper, we focus the export policies for provider, i.e., strategies a customer uses to export prefixes to its providers. Measurement result for peers' export policies is shown in work [14].

A. Algorithm for Inferring Export Policies to Providers

The direct way to infer a customer's export policies is to use the BGP table from its provider. After searching prefixes originated by the customer in the table, if those prefixes have customer routes, which we defined above, we know that the customer exports those prefixes to the provider. On the contrary, if those prefixes do not exist or do not have customer routes, it implies that the customer does not export them to the provider directly.

Therefore, we infer the export policies from the viewpoint of a provider. For a given provider, if it receives a prefix originated by a customer via a peer path instead of a customer path, we call this prefix a *selective announced prefix* (*SA* prefix) with respect to the provider. Here, we use our analysis result that customer's selective announcement policies give rise to *SA* prefixes. That is, a customer exports prefixes to a subset of providers instead of all providers. As a result, the selective announcement used by customers can be observed from the viewpoint of a provider.

```
Algorithm for inferring export policy
Input:
    Annotated AS graph $G$
    AS $o$ which originates prefixes $P$
    routing table from the viewpoint of AS $u$
Output:
    Whether $P$ contains SA prefixes from the point of
    view of AS $u$
Phase1: Initiation
    Selected AS set $S = \{u\}$

Phase2: Determine if AS $o$ is a customer of AS $u$
    while there is a selected AS
      for each AS $v$ that is a customer of the selected AS
        if $v$ is AS $o$
          $o$ is a customer of AS $u$
          go to Phase3
        else add $v$ into $S$
    AS $o$ is not a customer of $u$
    return

Phase3: Determine if $P$ contains SA prefixes
    for each next hop AS $w$ of the best route to $p_i$, $p_i \in P$
      if $u$ is not a provider of $w$
        $p_i$ is a SA prefix
      else $p_i$ is not a SA prefix
    if any prefix in $P$ is a SA prefix
      then $P$ contains SA prefixes
    else $P$ does not contain SA prefixes
```

Fig. 2. Algorithm for inferring exporting policy.

Table 4. Percentage of *SA* prefixes for 16 ASes.

AS number	% of *SA* prefixes	AS number	% of *SA* prefixes
AS 1	32	AS 7018	22
AS 3549	23	AS 701	27.8
AS 6453	48.6	AS 6461	4
AS 1239	29.4	AS 3561	5.2
AS 2914	14	AS 209	38
AS 5511	18	AS 577	17
AS 6538	11	AS 6667	13
AS 12359	0	AS 12859	0

Table 5. Percentage of prefixes from each customer inferred as *SA* prefixes for AS 1, AS 3549, and AS 7018.

Customer	# of prefixes	# of *SA* prefixes for AS 1, AS 3549 and AS 7018
AS 376	344	205 (60%)
AS 6280	33	32 (97%)
AS 10910	51	17 (33%)
AS 11647	28	24 (86%)
AS 14743	22	15 (68%)
AS 15087	65	11 (17%)
AS 19024	30	13 (43%)
AS 19916	25	24 (96%)

For example, in Fig. 1, customer A exports prefix p to a selected subset of providers, provider C. In D's BGP routing table, prefix p is received from D's peer, E. No customer route to p is received from customer B.

Note that from the point of view of a provider, the best routes to customers' prefixes, instead of all available routes, are sufficient to infer the selective announcement policies. From Section IV-A, we know that a customer route is typically preferred over other routes. In a provider's BGP table, if a customer route to a prefix exists, the route is the best route as well. Otherwise, if a customer route does not exist, the best routes are peer routes or provider routes.

The first step of the algorithm for inferring export policies to the provider for a given AS is to find if the AS is a customer of the monitored provider. This can be solved by using *depth first search* (DFS) algorithm in a directed graph to find a customer path from the provider to the AS. If there is a customer path between the AS and the provider, the AS is a customer of the provider. Not all paths found by DFS can be customer paths, however, those paths should obey export rules described in Section II. That is, from the direction of provider down to customer, each pair of ASes in the path should have provider-to-customer relationship. In an annotated AS graph $G = (V, E)$, we use modified *DFS* which satisfies path relationship constrains to find a customer path between a pair of ASes.

The next step is to investigate if the best routes to the customer's prefixes are peer or provider routes. If the best routes are peer or provider routes, those prefixes are not exported from the customer to the provider, or some intermediate customers who receive those prefixes do not export them.

Fig. 2 shows the algorithm in detail. Given an AS, we use this algorithm repeatedly for the AS's customers to infer those customers' export policies.

B. Prevalence of SA Prefixes

Here, we present experimental results of inferring export policies using the algorithm. We first use dataset described in Section III to construct the annotated AS graph which is used to find all direct or indirect customers of a given provider. We then use the routes from Oregon or ASes' BGP tables to derive the best routes to customers' prefixes. *SA* prefixes for 10 Tier-1 ASes can be inferred by using Oregon RouteView and 3 Tier-1 ASes' BGP tables.

Table 4 shows the percentage of customers' prefixes that are *SA* prefixes for 16 ASes. We find that Tier-1 ASes, such as AS 1, AS 3549, and AS 7018, have a significant number of *SA* prefixes. Those Tier-1 ASes reach their (direct or indirect) customers via their peers instead of customers. For example, AS 6280 is a customer of AS 1. However, AS 1 does not receive a prefix p originated by AS 6280. It receives p from its peer, AS3549. Note that *SA* prefixes for a provider may be due to the selective announcement policies of originating ASes or intermediate ASes.

Next, we examine *SA* prefixes from the viewpoint of a set of customers. We consider those customers which all have 3 direct or indirect providers: AS 1, AS 3549, and AS 7018. From those customers, we select 8 ASes which originate a significant number of prefixes as shown in Table 5. Table 5 shows that those 3 providers cannot access some of customers' prefixes directly via their customer paths.

Applying the selective announcement policies, a customer can balance its inbound traffic but its inbound and outbound traffic might be asymmetric. From the point of view of a provider, it may find that traffic between its customers has to forward to the rest of Internet via its peer links. This strategy may affect traffic engineering practice of the provider.

C. Representativeness of Data Sources for Inferring Export Policies

From the viewpoint of tier-1 ASes, we can infer export policies applied by stub ASes and non-tier-1 ASes. Our algorithm can infer how an AS exports its prefixes to its providers if the AS is a multi-homed AS. On Nov. 18, 2002, 65% of 14202 ASes derived from Oregon RouteView are multi-homed ASes. Using the algorithm described in Fig. 2, we can infer about 80% of multi-homed ASes' export policies.

VI. CONCLUSION

In this paper, we demonstrate how to infer the routing policies and characterize the routing policies. We first infer the import policy. We find that for most ASes routing preference conforms to AS relationships. Routes learned from customers are typically preferred over those from providers and peers, and routes received from peers are typically preferred over those from providers. Moreover, route preference assignment is based on next hop ASes. Second, we present an algorithm for inferring export policies, and characterize export policies. Customers can export their prefixes to a selected subset of providers. For 3 Tier-1 ASes, we find a large number of prefixes are exported to a selected subset of providers. From our study, customers announce their prefixes to a subset of providers because of load balancing. We find that selective announcement is harmful to the resilience and stability of the Internet. We hope that our findings will caution network operators in designing routing policies.

REFERENCES

[1] Y. Rekhter and T. Li, "Border gateway protocol 4 (BGP-4)," RFC 1771, T. J. Watson Research Center, IBM Corp., Cisco Systems, Mar. 1995.

[2] H. Tangmunarunkit, R.Govindan, S. Jamin, S. Shenker, and W. Willinger, "Network topology generators: Degree based vs. structural," in *Proc. ACM SIGCOMM*, Aug. 2002.

[3] R. Govindan and H. Tangmunarunkit, "Heuristics for Internet map discovery," in *Proc. IEEE INFOCOM*, 2000.

[4] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in *Proc. ACM SIGCOMM*, Cambridge, MA, USA, Sept. 1999.

[5] P. Barford, A. Bestavros, J. Byers, and M. Crovella, "On the marginal utility of network topology measurements," in *Proc. SIGCOMM Internet Measurement Workshop (IMW 2001)*, Nov. 2001.

[6] Meirt Network. [Online]. Available: http://www.irr.net

[7] G. Huston, "Interconnection, peering, and settlements," in *Proc. INET*, June 1999.

[8] University of Oregon Route Views Project. [Online]. Available: http://www.antc.uoregon.edu/route-views/

[9] Trace route. [Online]. Available: http://www.traceroute.org

[10] Merit Network. [Online]. Available: ftp://ftp.radb.net/routing.arbiter/radb/dbase/

[11] L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE Trans. Networking*, Dec, 2001.

[12] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *Proc. IEEE INFOCOM*, June 2002.

[13] G. D. Battista, M. Patrignani, and M. Pizzonia, "Computing the types of the relationships between autonomous systems," Tech. Rep. RT-DIA-73-2002, Dipartimento di Informatica e Automazione, University Roma Tre, Rome, 2002.

[14] F. Wang and L. Gao, "On inferring and characterizing Internet routing policies," in *Proc. ACM Internet Measurement Conf.*, Miami, USA, Oct. 2003.

Feng Wang is an assistant professor of School of Engineering Computational & Sciences at Liberty University. He received his Ph.D. degree in Electrical and Computer Engineering at the University of Massachusetts, Amherst. He received the B.A. degree from Zhejiang University in China, and M.S. degree from Yanshan University in China. His research interests include Internet routing, Internet measurement, network security, and embedded computer system.

Lixin Gao is an associate professor of Electrical and Computer Engineering at the University of Massachusetts, Amherst. She received her Ph.D. degree in computer science from the University of Massachusetts at Amherst in 1996. Her research interests include multimedia networking and Internet routing and security. Between May 1999 and January 2000, she was a visiting researcher at AT&T Research Labs and DIMACS. She is an Alfred P. Sloan fellow and received an NSF CAREER award in 1999. She has served on number of technical program committees including SIGCOMM2006, SIGCOMM2004, SIGMETRICS2003, and INFOCOM2004, and is on the editorial board of IEEE Transactions on Networking.