

On Routing Asymmetry in the Internet

Yihua He Michalis Faloutsos Srikanth Krishnamurthy
yhe@cs.ucr.edu michalis@cs.ucr.edu krish@cs.ucr.edu
Department of Computer Science and Engineering
University of California, Riverside
Riverside, CA 92521 USA
Tel. (951) 827-2348

Bradley Huffaker
bhuffake@caida.org
CAIDA/SDSC
University of California, San Diego
La Jolla, CA 92093 USA
Tel. (858) 822-0932

Abstract

Routing asymmetry in the Internet can significantly affect the manner in which we model and simulate its behavior. In this paper, we study routing asymmetry in the Internet and present quantitative evaluations on the extent of such asymmetry today. Our quantitative evaluations provide a measure of the *difference* between the forward and reverse paths between two end points. Routing asymmetry has not been studied extensively before; this is primarily due to the lack of a systematic approach for quantifying asymmetry except for simply computing the difference between the forward and reverse path lengths. By applying our framework for representing asymmetry, we quantify routing asymmetry for both US higher education academic networks and general commercial networks at two different levels: the Autonomous System (AS) level and the router (or link) level. We take into consideration, not only the difference in the forward and reverse path lengths, but also the AS and link identities and the sequence in which these entities appear on the paths. We measure the AS level routing asymmetry, and provide upper lower bounds on *link level* routing asymmetry. Our studies show that academic networks appear to be more symmetric than general commercially deployed networks. Furthermore, our studies demonstrate that routing asymmetry exhibits a skewed distribution i.e., a few end-points seem to display a higher extent of participation on asymmetric routes.

Keywords

ASYMMETRY, ROUTING, INTERNET, MEASUREMENT, FRAMEWORK, AUTONOMOUS SYSTEM

The paper is intended for Autonomic Networks Symposium in Globecom 2005

On Routing Asymmetry in the Internet

Yihua He Michalis Faloutsos Srikanth Krishnamurthy
yhe@cs.ucr.edu michalis@cs.ucr.edu krish@cs.ucr.edu
Department of Computer Science and Engineering
University of California, Riverside

Bradley Huffaker
bhuffake@caida.org
CAIDA/SDSC
University of California, San Diego

Abstract—Routing asymmetry in the Internet can significantly affect the manner in which we model and simulate its behavior. In this paper, we study routing asymmetry in the Internet and present quantitative evaluations on the extent of such asymmetry today. Our quantitative evaluations provide a measure of the difference between the forward and reverse paths between two end points. Routing asymmetry has not been studied extensively before; this is primarily due to the lack of a systematic approach for quantifying asymmetry except for simply computing the difference between the forward and reverse path lengths. By applying our framework for representing asymmetry, we quantify routing asymmetry for both US higher education academic networks and general commercial networks at two different levels: the Autonomous System (AS) level and the router (or link) level. We take into consideration, not only the difference in the forward and reverse path lengths, but also the AS and link identities and the sequence in which these entities appear on the paths. We measure the AS level routing asymmetry, and provide upper lower bounds on *link level* routing asymmetry. Our studies show that academic networks appear to be more symmetric than general commercially deployed networks. Furthermore, our studies demonstrate that routing asymmetry exhibits a skewed distribution i.e., a few end-points seem to display a higher extent of participation on asymmetric routes.

I. INTRODUCTION

Routing asymmetry is an important routing phenomenon, which can influence the manner in which, we model and simulate the Internet. For a pair of hosts A and B, if the path from A to B (forward direction) is different from the path from B to A (reverse direction), we say that the pair of paths between A and B exhibit routing asymmetry. This asymmetry in the Internet can appear on both AS level and router level paths. Asymmetry at the AS level also leads to router level asymmetry. However, router level asymmetry can exist even if the AS paths are symmetric; this would be due to asymmetry within an AS.

It is of interest to examine the reasons for routing asymmetry. Long term asymmetric routes are mainly created due to routing policies and traffic engineering. A packet in a provider’s network but destined for a different network, is moved out of the provider’s network as per the provider’s policies. This is done even if it would cause the packet to experience a longer path or higher levels of congestion than what exists within the provider’s network. This practice often results in the hot potato routing or an “early exit” [12]; this policy causes link level asymmetry, and potentially could also result in AS level asymmetry. Load balancing[3] could also lead to routing asymmetry; a router may attempt to *shift* traffic from a highly loaded link to a lightly loaded link. In other words, with traffic engineering, the network may potentially alter routing to avoid congested regions. At the same time, the Internet may exhibit transient effects of asymmetry due to

the delayed BGP convergence. In addition, the absence of a unique shortest path [9] [10] between a pair of hosts could also lead to asymmetric routing, since routing protocols can randomly choose any of the possible shortest paths.

Routing asymmetry is important on the way we measure, model, and manage the network. Identifying topological and performance properties of paths becomes more challenging in the presence of asymmetry. First, modeling paths becomes more difficult, since we need to identify two directions instead of one, between any pair of hosts. Second, asymmetry introduces significant problems in estimating one way latency between hosts. Measuring and modeling one-way delay is important for many delay-sensitive applications. Currently, the most common practice is to estimate the one way delay by the equating it to half of the easier-to-compute round-trip time delay[13]. This estimate becomes worse as the routing asymmetry increases. Third, monitoring the routing asymmetry may potentially be an important indicator of the state of the Internet. For example, higher than normal asymmetry may suggest changes or even errors in current routing practices. Finally, high levels of asymmetry may require us to revisit the design philosophies of many of currently employed higher layer protocols.

An in-depth study of routing asymmetry can undoubtedly increase our understanding of the Internet. However, there have been few studies on routing asymmetry. One may attribute this to the absence of a systematic approach for measuring asymmetry. Most previous work focuses on asymmetry in terms of the difference in lengths between the forward and reverse paths [1][18] between pairs of nodes. Typically, such efforts simply classify a path as either asymmetric or symmetric without quantifying the magnitude of the asymmetry.

The goal of this paper is to quantify the asymmetry in the Internet at the AS and router (or link) levels. Note that our objective is not just to identify the existence of routing asymmetry but, more importantly, *measure* its magnitude. Since routing asymmetry can appear at both the AS level and on router level paths, we examine both in our study. We also try to look at the differences in asymmetry between two kinds of networks — the academic networks among US higher education and research institutes, and the commercial Internet. In summary, we present a novel framework for quantifying asymmetry and apply it to study the aforementioned two networks. We quantify the differences between the asymmetry displayed in the two networks. In more detail, our contributions can be summarized as follows:

- We apply our “framework for quantifying asymmetry” to capture the magnitude and the relative significance of the routing asymmetry at both the AS and the router (link)

levels.

- We propose a heuristic to determine if two IP addresses correspond to the interfaces of the same link. We then use this heuristic to estimate the range of link-level asymmetry.
- We compare the distribution of asymmetry in US academic networks with that in commercial networks. We find that the commercial Internet not only has a higher percentage of asymmetric pairs, but also displays a higher magnitude of the asymmetry: about 10% of the routes in the commercial Internet exhibit high asymmetry at the AS level and almost none of the routes in academic networks exhibit such high asymmetry.
- We find that the spatial distribution of the asymmetry is skewed; a few end-points consistently participate on asymmetric paths. This implies that the asymmetry may be directly related to the way a host connects to the Internet, i.e., the specific provider associated with the host. In particular, whether or not a host perceives asymmetry may depend on the routing policies that are used by the *provider* of the host.

The rest of this paper is organized as follows: In section 2, we briefly present background and related work on the topic of routing asymmetry. We describe our approach to quantify asymmetry at both the AS and router levels in section 3. A presentation on the analysis of the distribution of routing asymmetry will be in section 4. In section 5, we conclude our work and discuss possible future research directions.

II. BACKGROUND AND RELATED WORK

The seminal work in this area by Paxson [1] studies routing pathologies and defines the problem of routing asymmetry. Subsequently, other studies have addressed this issue partially, looking primarily at the difference in the path lengths (or the difference in delay) between the forward and reverse paths[4][18]. There has been some work on path inflation due to policies [7][8][12]. Although these topics are related to our work, they are fundamentally different for two reasons: (1) asymmetry does not necessarily result in a difference in path length or delay incurred on the forward and reverse paths; and (2) path inflation does not necessarily result in routing asymmetry (both the forward and reverse paths may be the same; but they may be inflated in comparison with the shortest paths between the considered end-hosts).

Our previous work[20] builds a framework to quantify routing asymmetry in a systematic way; some properties of Internet asymmetry are studied. However, the work is limited to the study of asymmetry at the AS level and the measurement data is limited to the US higher education and research network. In this paper, we study router level topologies and consider both the US education and research network and the commercial Internet.

Several tools have been built to study the end-to-end routing properties in the Internet. Mercator [5] is a routing map collection tool run from a single host. Skitter [6] monitors, probe the network from about 20 different locations worldwide. Rocketfuel [11] uses a list of public traceroute servers to

probe ISP maps. However, these efforts have not addressed the quantification of routing asymmetry.

III. METHODOLOGY

In this section, we first review our framework for quantifying the routing asymmetry between two end hosts. Then, we explain how the data is collected and our methodology for mapping IP addresses to AS numbers. Finally, we discuss a new heuristic for evaluating the link level asymmetry.

A. Framework to quantifying routing asymmetry

The first challenge in systematically quantifying asymmetry is the design of an appropriate algorithm and the requirement of a set of metrics that can effectively capture various notions of asymmetry, such as path length, differences in the identities of the entities (AS/router) on the forward and reverse routes and the sequence in which these entities appear on the routes. We propose a framework to quantify the routing asymmetry between two end hosts by computing the *dissimilarity* between the entities on the two routes, considered in sequence. Our framework consists of two steps:

Step 1: For any two entities¹ x and y in a pair of paths, we define

a non-negative *base dissimilarity value* $\omega[x, y]$, which represents the magnitude of “how much x is different from y ”. At the AS level, the value should be set to 0 if $x = y$, and would be greater than zero if $x \neq y$. At the link level, the criteria may vary, as we will discuss in section III-C. The set of such values for all entity pairs forms a *base dissimilarity matrix* $\Omega(S)$, where S represents the set of all entities on the path. If a base dissimilarity matrix only contains values of 0 and 1, it is called a *simple base dissimilarity matrix*.

Step 2: We compute the *minimal composite dissimilarity* of each pair of routes for a given base dissimilarity matrix. By means of a dynamic programming strategy, the algorithm runs in $O(mn)$ time, where m and n are the path lengths in the forward and the reverse direction, respectively.

We define two asymmetry metrics based on the minimal composite dissimilarity.

Absolute Asymmetry (AA) is the *minimal composite dissimilarity* between a pair of forward and backward paths u and v .

$$AA(u, v) = C_{min}$$

length-based Normalized Asymmetry (NA) is the Absolute Asymmetry normalized by the computed round-trip path length.

$$NA(u, v) = \frac{AA(u, v)}{length(u) + length(v)}$$

The absolute asymmetry captures the absolute magnitude of asymmetry inherent in a pair of forward and backward paths, and the length-based normalized asymmetry indicates the extent to which a forward route between a pair of nodes is “off” from its reverse counterpart; in other words, a pair

¹Here, entities could be either ASes or links.

TABLE I
GEOGRAPHIC DISTRIBUTION OF AMP MONITORS

Region	#of AMP sites	Percentage
North America	122	90.4%
Pacific and Asia	6	4.4%
Europe	5	3.7%
Latin America	2	1.5%
Total:	135	100%




Fig. 1. Geographic locations of RETRO vantage points

of long routes are considered to be less significantly deviant from symmetry as compared with a pair of shorter routes with the same absolute asymmetry. From the discussion above, we immediately note that for any pair of routes u and v , $0 \leq NA(u, v) \leq 1$ if $\forall \omega(x, y), 0 \leq \omega(x, y) \leq 1$. In the two extreme cases when (a) u and v are exactly the same, $NA(u, v) = 0$; and when (b) u and v contain completely different entities, $NA(u, v) = 1$.²

B. A New Tool and Data collection

We use two sets of data, representing two different types of networks that often appear in computer network research literature. The first data set is collected by the Active Measurement Project (AMP) from NLANR [2] on Jan 22, 2004. The AMP architecture uses a near-full mesh (each monitor sends probes to almost all of the other monitors) that interconnects approximately 135 active monitors deployed at spread out sites. These sites are mostly distributed among North American higher education and research institutes. However, sites are also in the Asian Pacific, Latin America, and Europe (see Table I).

The second set of data was collected by our REverse TraceRO tool (RETRO) on Feb 5, 2005. RETRO is a convenient tool that can collect traceroute server’s configurations, send out traceroute requests, and collect/parse traceroute results dynamically. Currently, we mainly use public traceroute servers listed in [21]. These traceroute servers have about 1200 vantage points, which are located in about 350 different ASes, covering 55 different countries. Fig. 1 shows their locations on a world map. From among these vantage points, we randomly pick a vantage point for each AS (if there are more than one vantage point for that AS) and perform a full mesh traceroute between each pair from among these points, as per the scheduling algorithm we will describe in the next paragraph.

²In fact, if we use the simple basic dissimilarity matrix, NA can never be equal to 1 since, in all cases, $AA(u, v) < length(u) + length(v)$, due to the fact that the end host pair on the forward and reverse paths are identical.




Fig. 2. Traceroute only probes the router interfaces

For the AMP dataset, we have extracted traceroute results conducted between 00:05am and 00:15am PST on Jan 22, 2004. We wish our RETRO tool could allow us to capture a “snapshot” of the Internet routing at any specified time. However, there is an upper limit imposed on the frequency (or a lower limit on the time between consecutive trace route queries) with which trace route probes may be generated from public trace route servers. Thus, within the small time window specified above, it is hard to set forth multiple probes from a single server³ To solve the problem, we schedule the traceroutes in such a way⁴ that for any two monitors A and B, the probe from A to B and the probe from B to A are triggered simultaneously. Thus, we ensure that each pair of paths that we compare are approximately computed at the same time. These efforts are intended to minimize the temporal impact of changes [14] on the traceroute results.

There are two levels of routing asymmetry that can be quantified by our framework: the AS level asymmetry and the router (link) level asymmetry. To identify the ASes that a path passes through, we map the IP addresses from the traceroute data to AS numbers as per existing IP-AS mapping policies [16][17]. After this conversion, consecutive entities that reflect the same AS number are collapsed into a single entity. Our methodology is similar to the technique used in [7]. Thus, we obtain AS level data paths, and for these paths we compute the corresponding AA and NA metrics.

C. A heuristic towards measuring the link level asymmetry

The challenge in measuring router/link level asymmetry is the determination of whether two IP addresses belong to the same link. Note that IP addresses refer to *link interfaces* at routers. To perform router level analysis, we need to understand some of the intricacies of the data that we obtain from traceroute-like probes. The traceroute data returns IP addresses, one address per router. This address, normally would correspond to the interface on the router via which the packet entered the router⁵. Since each link has at least two interfaces, and each interface could belong to a different router, the interface that is indicated due to the traversal of the link in the forward path will be *different* from that which is

³Most traceroute servers do not explicitly post a limit of probing rate. But these traceroute servers should not be abused. In practice, we enforce a 900 seconds time period between each traceroute request for each traceroute server.

⁴The problem is formally known as “Edge Coloring Problem” in graph theory.

⁵Normally, the ICMP TIME_EXCEEDED message is sent from the address that corresponds to the *incoming* interface, i.e., the interface via which the packet entered the router. However, RFC792 (on ICMP) does not explicitly specify the interface using which a router should send out an ICMP message. A more detailed discussion can be found in [16].

indicated by the traceroute on the reverse path (although the path might be identical).

For example, in Fig 2, if we send a traceroute from node $R1$ to $R4$, and the route is $R1 \rightarrow R2 \rightarrow R3 \rightarrow R4$, the traceroute is going to return the IP addresses corresponding to the network interfaces (b, d, f) , in sequence. If the path is symmetric (the forward and the reverse paths are the same), the traceroute would return IP addresses corresponding to the interfaces (e, c, a) , in sequence. On an alternate “asymmetric” reverse path (shown by dotted lines) ($R4 \rightarrow R5 \rightarrow R2 \rightarrow R1$) the traceroute returns the IP addresses corresponding to the interfaces (j, h, a) , in sequence. The challenge then is to detect that interfaces c and d belong to the same link whereas interfaces h and d belong to different links.

The case is relatively easy for a point-to-point link since the link would form a network by itself and the interfaces on the link should share a common IP prefix. Traditionally, numbered Internet subnets do not use subnet masks longer than 30 bits; this requires four addresses per link: two host addresses, one network address, and one broadcast address. In the binary representation of the IP addresses, the last two bits of those two host addresses are 10 and 01, and the last two bits of the network address and the broadcast address are 00 and 11, respectively. Since the network address and the broadcast address are generally not used in point-to-point links, RFC 3021 proposed the use of 31-bit subnet masks in point-to-point links in year 2000 in order to further conserve IP address space in the Internet. Nowadays, both 30-bit and 31-bit subnet masks are used for point-to-point links. Theoretically, ISPs could use subnet masks that are shorter than 30 bits to represent a point-to-point link. But since such a practice would obviously waste the address space, we believe that shorter-than-30-bit subnet masks are rarely used with point-to-point links.

Assuming that a point-to-point link could only have a 30-bit or 31-bit subnet mask, we propose the following heuristic to determine whether or not two IP interfaces belong to the same point-to-point link:

Two IP addresses, one taken from a forward IP path and the other from its reverse IP path, belong to the same point-to-point link, if any of the following two conditions holds:

- (1) the two IP interface addresses differ only in the last bit (31-bit subnet mask)
- (2) the two IP interface addresses only differ in the last two bits (30-bit subnet mask) and their last two bits are 01 and 10; in addition, neither their network address (last two bits: 00) nor their broadcast address (last two bits: 11) appears in any traceroute induced dataset.

However, not all links on the paths are point-to-point links. This is particularly true near the edge of the Internet, where Ethernet based links are very commonly used. Further more, inside the Internet, IXPs (Internet Exchange Points) allow publicly peering ISPs to exchange traffic via either an FDDI ring, Ethernet or ATM. The IXP addresses, which are visible in route probing data, often use 24-bit subnet masks[19]. These factors potentially result in causing “overestimates” of routing asymmetry, since IP interfaces from the same *multiple-access* link may be incorrectly classified as “not belonging to the

same link”, since the interfaces do not share a 30-bit or a 31-bit common prefix. In the next section, we overcome these issues by relaxing the manner in which we look for matching IP addresses (thus far they can differ only in the last two bits) and examine how the link level asymmetry distribution would vary with varying levels of relaxation.

IV. EXPERIMENTAL EVALUATIONS

We first examine the asymmetry at the AS level for both the AMP dataset and the RETRO dataset. Fig 3 depicts the cumulative distribution of Absolute Asymmetry (AA). In this plot, we first compute the minimal composite dissimilarity of all the routes in the AMP and the RETRO datasets, respectively; we use the simple base dissimilarity matrix. Next we sort them in the descending order of their AA. The ordinate represents the value of AA and on the abscissa, we have the percentage of nodes that have at least, the value of AA that is indicated by the corresponding ordinate value. We observe that about 65% of the total pairs of routes have some degree of asymmetry at AS level (i.e., the forward path and the reverse path pass through at least one different AS), with our RETRO dataset. This is consistent with the result observed in [18]. However, the proportion of asymmetric route pairs is much lower if the AMP dataset is considered than if the RETRO dataset is considered: only about 14% pairs of routes are asymmetric as per the AMP dataset. Furthermore, at each accumulative fraction level, the RETRO dataset displays a higher magnitude of asymmetry than the AMP dataset. For example, the 10% of the routes that display the highest levels of asymmetry in commercial networks have an NA⁶ larger than 0.5; however, the 10% of the routes that display the maximum AA in academic networks only have an NA larger than 0.1. This implies that the RETRO dataset not only has a higher percentage of asymmetric routes, but also displays a higher level of asymmetry than the AMP dataset.

Our initial suspicion was that the RETRO dataset displays a higher level of asymmetry than the AMP dataset for the following reason: RETRO monitors are much more spread out with respect to each other than AMP monitors, most of which are in the United States. This could lead to longer AS paths in the RETRO dataset, and longer paths are more likely to be asymmetric. However, the initial suspicion does not seem to be true. In Fig. 4, we plot the length-based Normalized Asymmetry (NA) versus the AS path length. NA provides an intuitive view of the significance of asymmetry between a pair of routes based on its round-trip path length. We see that the NA level in the RETRO dataset is still higher than that in the AMP dataset. This suggests that path length is not a significant significant factor for difference in the asymmetry levels between the RETRO and the AMP datasets.

One possible explanation for the difference in the level of asymmetry displayed by the two datasets is that a majority of the AMP monitors are connected by a single top tier provider namely, Abilene; this appears to probably provide the shortest or the most preferred paths of the immediate ISPs of the AMP monitors at any two chosen ends. On the contrary,

⁶Recall that NA represents the length-based Normalized Asymmetry.




Fig. 3. AA Distribution at AS level




Fig. 4. NA Distribution at AS level

RETRO monitors have a plurality of top level providers; these providers typically peer with each other. The most preferred path selected by the immediate ISP at one end may not be the same as the one that is chosen at the other end. This is a direct consequence of the well known policy of hot-potato routing: an ISP will try to push a packet out of its network as soon as possible.




Fig. 5. NA Distribution for AMP data at link level

Second, we examine the routing asymmetry at the link level (or router level). However, as discussed in the previous section, we do not have an accurate way of knowing whether or not the two *different* network interfaces produced by traceroute probes correspond to the same link. Thus, we can not exactly tell if or not the routes are symmetric. To work around this obstacle, we postulate four different criteria to categorize paths as asymmetric and examine the base dissimilarity matrix with each of the four criteria. Using the results, we establish the




Fig. 6. NA Distribution for RETRO data at link level

distributions of the upper and lower bounds on asymmetry.

In each of the criteria, we compare two IP addresses, one chosen from the ones on the forward path and the other from the ones on the reverse path; every such candidate pair is compared. The criteria are as follows:

Criterion 1: If the pair of IP addresses are categorized to form a point-to-point link as per the heuristic introduced in the previous section, we assume that these two IP addresses belong to the same link⁷.

Criterion 2: If the two IP addresses are within the same /24 network, we assume that the two IP addresses belong to the same link.

Criterion 3: If the two IP addresses are within the same /16 network, we assume that the two IP addresses belong to the same link.

Criterion 4: If two IP addresses are within the same AS, we assume that the two IP addresses belong to the same link.

We wish to point out that criterion 1 and criterion 4 provide the upper bound and lower bounds, respectively, on the observed Internet routing asymmetry at the link level. Criterion 1 is the least forgiving of address discrepancy, while Criterion 4 is the most forgiving⁸. Criterion 2 and criterion 3 are used for the purposes of comparison. Fig 5 shows link level NA distribution for the AMP dataset with the above four criteria; the plots corresponding to the four criteria are labeled as “/31-/30”, “/24”, “/16” and “same AS” in the figure. We see that at least (shown by Criterion 4) 38% of the total pairs of routes show some degree of asymmetry at the link level. Furthermore, (shown by Criterion 1) 84% of the pairs of routes, at the most, display some level of asymmetry. The real distribution of asymmetry should fall between the plots represented by “/31-/30” and “same AS”. With the RETRO dataset (shown in Fig 6), as per the criteria used, 88% to 99% of the routes, display some level of asymmetry. In addition to the percentage of asymmetric routes being higher with the RETRO dataset than with the AMP dataset, the magnitude of asymmetry seen with the the RETRO dataset is also much higher. For example,

⁷In the base similarity matrix, if two IP addresses belong to the same link, we set $\omega[ip1, ip2] = 0$; otherwise, we set $\omega[ip1, ip2] = 1$.

⁸Note that Criterion 4 is not the same as comparing paths at the AS level. With “Criterion 4, a path still consists of a number of *redefined* nodes equal to the number of routers, but they are now “categorized” into classes that correspond to the AS that they belong to: for example, a path may be represented as (AS1, AS1, AS1, AS2, AS2). In contrast, an AS level path would be simply represented as (AS1, AS2).

the median pair of paths⁹ in the RETRO dataset has an NA ranging from 0.30 to 0.62, while with the AMP dataset, the value varies from 0 to 0.17.




Fig. 7. Spatial NA Distribution for AMP data




Fig. 8. Spatial NA Distribution for RETRO data

Next, we examine the spatial distribution of the asymmetry, i.e., we quantify the level of asymmetry that is experienced by a particular monitor, on average. For each AMP or RETRO monitor, we compute the average NA of all the routes that start from or end at that location. We sort the computed average NA in the ascending order and plot the distributions at the AS level in Fig. 7 and at the link level in Fig. 8. These two plots suggest that the average level of normalized asymmetry is not the same for all monitors. Some monitors consistently experience a higher average asymmetry than others. This is true for both the AMP dataset and the RETRO dataset. This seems to suggest that *the asymmetry is most likely related to the way in which a host connects to the Internet; in particular it suggests that it may depend on the routing practices of its Internet providers*. The rationale is that if the asymmetry was introduced at the middle, we would observe a more uniform distribution of the asymmetry, especially in terms of the length normalized asymmetry NA.

V. CONCLUSIONS

In this work, we present our framework for quantifying asymmetry and apply it to real Internet measurements. Thus,

⁹The median pair of paths refers to the pair of paths that has the rank of 50% from among all of the paths in terms of NA. Strictly speaking, the pair of paths that ranks at 50% as per criterion 1 may not rank exactly at the same level (50%) as per criterion 4. However, we observe that the difference in its rank as per the different criteria, is negligibly small.

we compute the distribution of the routing asymmetry in the Internet at both the AS and the link/router levels. At the AS level, we find that US academic networks display a low level of asymmetry while commercial networks show a much higher level of asymmetry. We discuss the intricacies that arise due to the fact that routers possess multiple interfaces and therefore multiple IP addresses. At the link level, we propose a heuristic to identify whether two IP addresses correspond to interfaces that span the same link. We construct a set of criteria that can help us compute the distributions for upper and lower bounds on the link level asymmetry in the Internet. As with our studies on AS level asymmetry, we find that at the link level, commercial networks seem to display a much higher percentage of asymmetric routes and with higher magnitudes of asymmetry than in US academic networks. We interpret this result to suggest that the commercial network differs significantly from the education network in terms of asymmetry and therefore its network properties.

REFERENCES

- [1] V. Paxson. End to end behavior in the Internet. In Proceeding of the ACM SIGCOMM, Volume 26, number 4, page 25-38, August 1996.
- [2] NLANR active measurement, available at <http://watt.nlanr.net>
- [3] R. Teixeira, K. Marzullo, S. Savage, G. Voelker, Characterizing and Measuring Path Diversity of Internet Topologies, ACM SIGMETRIC, San Diego, CA 2003
- [4] S. Savage, A. Collins, E. Hoffman, J. Snell and T. Anderson, The End-to-end Effects of Internet Path Selection, ACM SIGCOMM Conference, Cambridge, MA, September 1999.
- [5] R. Govindan and H. Tangmunarunkit, Heuristics for Internet Map Discovery, IEEE Infocom 2000
- [6] B. Huffaker, D. Plummer, D. Moore, K. Claffy, Topology Discovery by Active Probing, Symposium on Applications and the Internet (SAINT) 2002
- [7] H. Tangmunarunkit, R. Govindan, S. Shenker, D. Estrin, The Impact of Routing Policy on Internet Paths, IEEE Infocom, Anchorage, Alaska. 2001
- [8] L. Gao and F. Wang, The Extent of AS Path Inflation by Routing Policies, Global Internet 2002
- [9] R. Ghosh and G. Varghese, Symmetrical routes and reverse path congestion control. Technical Report TR-97-37, Department of Computer Science, Washington University, St. Louis, September 1997.
- [10] G. Varghese, R. Ghosh, Modifying shortest path routing protocols to create symmetrical routes routing information protocol. UCSD technical report CS2001-0685, September 2001.
- [11] N. Spring, R. Mahajan, D. Wetherall, Measuring ISP Topologies with Rocketfuel, ACM SIGCOMM, Pittsburgh, USA 2002
- [12] N. Spring, R. Mahajan, T. Anderson, Quantifying the Causes of Path Inflation, ACM SIGCOMM, Karlsruhe, Germany, 2003
- [13] Q. Zhang, W. Zhu, Y. Zhang, "Resource allocation for multimedia streaming over the Internet", IEEE Trans. on Multimedia, Sept. 2001
- [14] C. Labovitz, G. Malan and F. Jahanian, Internet Routing Instability, ACM SIGCOMM, Cannes, France, 1997
- [15] V. Jacobson, "Traceroute", <ftp://ftp.ee.lbl.gov/traceroute.tar.gz>
- [16] Z. Mao, J. Rexford, J. Wang, R. Katz, Towards an Accurate AS-level Traceroute Tool, ACM SIGCOMM, Karlsruhe, Germany, Aug 2003
- [17] Z. Mao, D. Johnson, J. Rexford, J. Wang, R. Katz, Scalable and Accurate identification of AS-level Forwarding Paths, IEEE INFOCOM, Hongkong, China, 2004
- [18] L. Amini, H. Schulzrinne, A. Lazar, Observations from Router-level Internet Traces, DIMACS Workshop on Internet and WWW Measurement, Mapping and Modeling, February 13-15, 2002.
- [19] Packet Cleaning House, <http://www.pch.net/resources/data/exchange-points/ep-in-addr.txt>
- [20] Y. He, M. Faloutsos, S. V. Krishnamurthy. "Quantifying Routing Asymmetry in the Internet at the AS Level", IEEE GLOBECOM 2004 - Global Internet and Next Generation Networks, Dallas, Texas, November, 2004.
- [21] <http://www.traceroute.org>