

Analysis of Varying AS Path Lengths from the Edge of the Network

Arif Selcuk Uluagac*

Raheem A. Beyah[†]Roma Kane[‡]Siddharth Joshi[‡]

John A. Copeland*

*CSC Lab, School of ECE
Georgia Institute of Technology
Atlanta, GA 30332, USA
{selcuk,jcopeland}@ece.gatech.edu

[†]CAP Group, Dept. of CS
Georgia State University
Atlanta, GA 30303, USA
{rbeyah}@cs.gsu.edu

[‡]College of Computing
Georgia Institute of Technology
Atlanta, GA 30332, USA
{roma,joshis}@gatech.edu

Abstract—Understanding and analyzing the past and current behavior of the Internet will be instrumental in building tomorrow’s more efficient and scalable networks (e.g., the future Internet). In this paper, we study the impact of Autonomous Systems (ASs) paths’ end-to-end latency. Unfortunately, due to the diverse set of non-disclosed routing policies among ASs, packets belonging to a certain end-to-end connection may traverse different ASs, causing fluctuating AS paths. Fluctuation of AS paths has been studied in the literature directly from the core of the network. In this paper, we take a different approach to the analysis of the fluctuation, solely from the edge of the network. Specifically, from the end user’s perspective, some AS paths may be optimal (or better) and some sub-optimal. Furthermore, there is not a unique definition for sub-optimality as it may be reflected with various measures (e.g., latency) depending on the application requirements and expectations. In this paper we analyze fluctuating AS path lengths (ASPLs) and investigate their impact on the end-to-end latency over the Internet at a greater scale than previous studies. This study was conducted using *Scriptroute* to probe various PlanetLab nodes. Our results show that all of the source nodes experienced some AS path differences and the ASPL values that the sources use greatly vary. At worst, some nodes experienced different paths over 70% of the time during our measurements. We observed that the largest difference in ASPLs on a particular connection was as high as 6 with an average of 2.5. Moreover, we present real cases where ASPL and latency values are related, inversely related, and not related at all. Finally, we provide a simple definition for sub-optimality and analyze the collected data against this definition. We show that overall 82% of the fluctuating paths and 9% of all the traces between source-destination pairs faced sub-optimal AS paths.

Keywords—AS Paths, BGP, end-to-end Internet measurements, PlanetLab, Scriptroute, ASPL

I. INTRODUCTION

There is a significant effort in designing the future Internet (e.g., Future Internet Design (FIND) initiative of NSF [1]) among academia and industry. Understanding and analyzing the past and current behavior of the Internet will be instrumental in building tomorrow’s more efficient and scalable networks. With this in mind, we study the impact of fluctuating Autonomous Systems (ASs) paths on the end-to-end latency.

The Border Gateway Protocol (BGP) in the Internet allows ASs (e.g., ISPs) almost complete freedom in adopting their

routing preferences. This is because each AS is comprised of one or more connected network operators that have their own set of routing rules and policies. Since, connection relationships between ASs (i.e., peering) mostly determine the routing policies, different AS paths may be traversed between the same end points in a given communication setting. For instance, as seen in the illustration in Figure 1, two different AS paths can be traversed between the two end-points: AS1-AS7-AS6-AS5-AS4 and AS1-AS2-AS3-AS4. In some cases, this situation might cause Internet traffic to not traverse the shortest AS paths without affecting latency, whereas in other cases it may. Specifically, from the end user’s perspective, some AS paths may be optimal (or better) and some sub-optimal. In this paper, the impact of such *sub-optimal* AS paths is investigated from end-users’ perspective.




Fig. 1. A possible ASPL fluctuation between the two end-points (hatched circles).

AS path inflation is a well-known problem in the literature [2], [3]. Previous studies on the issue primarily looked at the causes of AS path inflation and focused on observing the problem from either a centralized vantage point [4] or the core of the network at different ASs [2]. In this paper, we take a complementary approach to previous research by analyzing the fluctuating AS paths from the edge of the network. The problem is also investigated over a greater scale than previous studies, spanning more than 100 PlanetLab [5] nodes. Using trace-driven methods, we identify fluctuating AS path lengths (ASPLs) between the same source and destination pairs. We note that the ASPL is basically the measure for the number of different ASs traversed between two endpoints in the Internet.

For instance, in Figure 1 two exemplified AS paths AS1-AS7-AS6-AS5-AS4 and AS1-AS2-AS3-AS4 between the endpoints have ASPLs of 5 and 4, respectively. Further, the impact of such varying paths over the users' measured end-to-end performance is studied. This involves the analysis of latency with respect to ASPLs.

In this paper, we label sub-optimal AS paths as the paths that have higher latency values between the same source and destination pairs. For simplicity, we assume that traces reflect the actual BGP behavior as our purpose is not to directly study BGP behavior. Thus, independent of the policies adapted by ASs, we provide this simple definition for sub-optimal paths and analyze our collected data against our definition to determine how frequent the paths between the same end hosts experience sub-optimal AS paths.

Our results show that the PlanetLab nodes that have oscillating ASPLs on a given connection exhibit a varying behavior. ASPL values showed fluctuations 30% and 11% of the times based on the two sets of data collected at different times. As a worst case, some nodes experienced different paths over 70% of the time during our measurements. We observed the largest ASPL difference as high as 6 with an average of 2.5. Furthermore, we witnessed that every source node took sub-optimal AS paths and that overall 82% of the fluctuating paths and 9% of all the traces between source-destination pairs faced sub-optimal AS paths.

The rest of this paper is organized as follows. In Section 2 related work is given. Section 3 explains our technical approach when collecting data from PlanetLab nodes. Analysis of the AS paths is the focus of Section 4. We then provide a simple definition for sub-optimality in Section 5. In Section 6 we articulate other possible definitions of sub-optimality. Section 7 explores the frequency of sub-optimal paths. Finally, the paper is concluded in Section 8.

II. RELATED WORK

In this section, we present previous related research efforts placing each in one of two categories. The first category involves studies focusing on AS path inflation, and the second category addresses BGP behavior.

A. AS Path Inflation

AS path inflation is a well-known problem that has been studied well by previous works. Gao and Wang [2] focused on the extent of AS path inflation by routing policies using RouteView route tables from the perspective of the core of the network (ASs). They showed two typical routing policies influencing the inflation. In [3], Spring et al. quantified the causes of path inflation. The study proposed a way to create the overall topology picture and to identify the routing policies between ISPs. Finally, the study presented the impact of the routing policies and the topology on the AS paths over 65 operational ISPs and 45 PlanetLab nodes. An earlier study by Savage et al. [4] focused on showing the existence of the perfect routing paths that enabled maximum performance for

end-users. They measured the performance from a centralized location to other nodes (between 20 to 40 nodes).

The theme of the aforementioned works and ours is similar. However, in our work, we solely aim to see the sub-optimality of the AS paths in relation to latency. We observe the frequency of such paths and their consequences to end-users. However, our intention is not to study the causes of oscillating AS paths. We complement the previous research in several ways. First, some of the previous studies primarily looked at the issue of AS path inflation from either a centralized vantage point [4] or the core of the network at different ASs [2] in a one-to-many fashion between the nodes. In contrast, we analyze the fluctuating AS paths from the edge of the network in a many-to-many relationship. Second, we look at the oscillation of the AS paths with more recent data compared to earlier efforts. The availability of the PlanetLab research platform made it possible to reach more nodes compared to earlier studies. Last, collecting data from PlanetLab nodes with Scriptroute tool allowed us to easily synchronize the AS paths and the latency values. Also, our study ignores the failures in the Internet and only focuses on the paths that are accessible on an end-to-end basis in order to see the effect of sub-optimal ASPLs.

B. BGP Behavior

There are several other studies related to the behavior of BGP and the effects of route selections on the performance seen by the end-users. Gao et al. defines the current BGP path selection problem and propose a solution to the BGP convergence problem [6]. Later, they extended their research and focus on the routing policies between ASs [7]. They investigated the import and export policies applied by ASs. Another relevant study by Mahajan et al. [8] focused on the mis-configurations of the BGP and its effects on Internet availability, suggesting a way to recognize them on the BGP announcements. In addition, Feamster et al. described the design and implementation of a tool called rcc, a router configuration checker, which finds BGP configuration faults using static analysis [9]. Furthermore, McPherson et al. focused on the route oscillations in an AS that were caused by IBGP and route reflectors [10]. They proposed two solutions to this problem: One involves using the MED attribute, whereas the other proposes creating a new protocol that does not contain a route oscillation problem in its behavior. Similarly, Basu et al. focused on route oscillation problems [11]. They presented some cases where the previous solutions may fail and proposed a better resolution. Likewise, we address the same oscillation issues in this paper. However, we focus mainly on the AS fluctuations. Lastly, in a recent paper [12], the authors presented measurements identifying slow convergence and path exploration in the global Internet routing system and proposed a new path-ranking method.

III. DATA COLLECTION METHODOLOGY

In this section, we describe the method used for obtaining and analyzing the data for this study.

To quantify the impact of changing AS path lengths (i.e., sub-optimal paths) we must first identify the presence of such paths in the Internet. The large number of ASs today in the Internet and the fact that their inter-connections are not publicly available make this objective a daunting task. Accepting this as the premise of our study, we used a trace-driven methodology to collect the data. Specifically, we ran scripts implemented in Ruby on PlanetLab [5] nodes. The PlanetLab nodes are a collection of machines distributed over the entire globe spanning over 30 countries and including 753 machines hosted by 363 sites. *Scriptroute* [13] is one of the measurement tools freely available on these nodes. We traced the route towards certain destinations via the implementation of *traceroute* in *scriptroute* followed by an AS-lookup via *whois* assuming that the traceroute reflected the actual BGP behavior. Also, we were able to remotely run our Ruby scripts with a cookie placed over PlanetLab nodes. The method of data collection is shown in Figure 2.




Fig. 2. Indirect measurements using Scriptroute.

Two sets of data were collected at different times for this study. The purpose of the first set of data was to observe the general AS path behavior (Section IV-A) whereas the aim of the second one was tailored more for studying the individual end-to-end pair of nodes (Section IV-B). For the first dataset, hereafter called *DataSet-I*, we selected the most stable 101 PlanetLab nodes from 19 countries, which were mostly educational sites. From 101 diverse points, we probed 100 other nodes simultaneously. We repeated the measurement cycle 27 times over a week in April-2006. In total, we collected data from 272700 traceroute probes. As for the second set of data, hereafter *DataSet-II*, it was comprised of 130 nodes spread over 27 countries in different continents from mostly educational sites. In *DataSet-II*, we collected 55 successful runs of data. This resulted in a total of 922337 traceroute values. The data was collected from March 2007 to April 2007. Both *DataSet-I* and *DataSet-II* are available for interested parties via the Communications Systems Center Lab webpage at Georgia Institute of Technology [14].

One important remark is on the number of available PlanetLab nodes for our study. Although PlanetLab has more than

700 nodes connected to the system, we realized that not all had a *scriptroute* daemon running on them. To address this problem, we pruned the nodes that did not have a *scriptroute* daemon running on them. This was done via another utility called *sr-listservers* script that is available inside the *scriptroute* tool.

IV. ANALYSIS OF THE AS PATHS

The aim of this section is first to exhibit the overall behavior of oscillating ASPLs. Then, we investigate ASPL and latency results from individual node pairs spanning different continents.

A. Overall Behavior

We first identify paths that use different AS paths between the same end-to-end pairs.

As there is no specific definition for sub-optimality of the AS paths in the Internet, it is reasonable to first analyze how paths between certain source and destination pairs behave. Given a certain source-destination pair, some paths may experience longer ASPLs in their successive transmissions to the same destinations due to many reasons¹. Therefore, investigation of how many of the probes from certain PlanetLab nodes had changing ASPLs towards their destinations is presented in this subsection. Without loss of generality, we assume that the paths that have varying ASPLs would have the potential of exhibiting sub-optimal paths. Figure 3

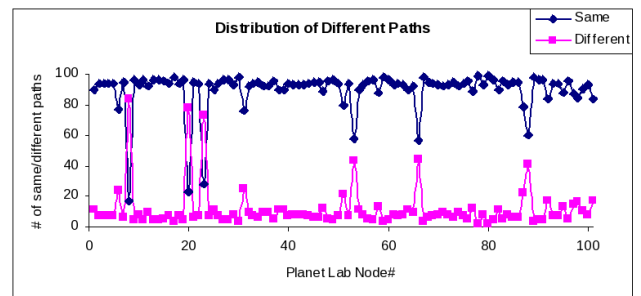


Fig. 3. Number of different and same paths towards their destinations from a given PlanetLab node.

shows the number of different and same paths towards their destinations from PlanetLab nodes. As seen in the figure, the number of PlanetLab nodes with oscillating ASPLs exhibit a varying behavior. At worst, some nodes experienced different paths during our measurements over 70% of the time. We also provide how much a certain ASPL towards a destination differ from its subsequent probes to the same destination in Figure 4. The largest difference in ASPL to a destination has been observed to be as high as 6 and as low as 1. Overall the average difference of ASPLs in the same figure is 2.5. It is our observation that all of the source nodes experienced some

¹Please note that there are several studies primarily investigating the cause of AS path inflation in the Internet. Thus, studying the reasons is outside of the scope of this paper.

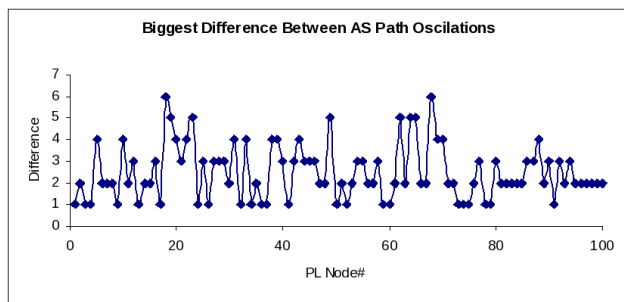


Fig. 4. Biggest difference between varying AS Path lengths towards their destinations from a given PlanetLab node.

path differences. Overall 11%² of all probes from sources to destinations took different ASPLs. There were 25 nodes out of 101 PlanetLab nodes that experienced more than or equal to a 10% differences in ASPLs (from their lowest ASPL values) with the average ASPL of almost 3. On the other hand, there are 37 PlanetLab nodes that experienced more than 2 differences in ASPLs with the average of almost 4.

B. Results Based on Intercontinental Traffic and Geographical Location

In this part of the study, we look at individual node pairs from different continents. For this, we determined the following nodes pairings from the continents using DataSet-II: US-US, US-Asia, US-Europe, Europe-Asia, Europe-Europe. Moreover, under each category, we tried to choose different countries where possible. Due to space limitations, we only include a subset of pairings. In all the figures, the forward and reverse path ASPL and latency results are plotted together to allow one to make meaningful comparisons.




Fig. 5. Individual pair: One node (14) is in the US and the other (121) is in China.

²The overall fluctuation rate was 30% for DataSet-II.

In Figure 5, one node is located in the US and the other node is located in China. Several odd behaviors can be observed in analyzing this figure. For instance, a big downward jump in the ASPL may increase the latency, similarly an upward increase in the ASPL may decrease the latency, and finally an increase in the ASPL may not cause any distinguishable effect on the latency. In the same figure (Figure 5), one can observe that a steady ASPL may experience fluctuating latency values. Depending on the user application, the value of the latency difference may cause some performance degradation. Moreover, in Figure 5, we observe that a higher value of ASPL on one path may not correspond to a higher value of latency in the same path in comparison to the opposite path. The latency may be the same as or lower than that of the other path.

Finally, small ASPL fluctuations may not disturb the end-to-end latency as can be seen in Figure 6. In this figure, node 62 is located in Korea and the other node (35) is in Italy.




Fig. 6. Individual pair: One node (62) is in Korea and the other (35) is in Italy.

V. SUB-OPTIMALITY ANALYSIS

In the previous section, we identified the presence of oscillating ASPLs, which have great potential of exhibiting sub-optimal paths. We also presented our observation from the collected data that a smaller ASPL value could exhibit a poorer end-to-end latency than that of a big ASPL value among the same end-points. In this section, we analyze the data to infer sub-optimal ASPLs.

One can come up with many definitions for sub-optimal AS paths. For a certain source and destination pair A and B , we simply define *sub-optimality* as an AS path that exhibits inferior end-to-end latency to its counterparts. Specifically, in a set of n RTT and ASPL observations, we can find the optimal path among different paths by identifying the one that has the lowest mean latency (RTT_{AB}^{min}), compare this value with the other ASPLs' mean latencies. If the separation among mean values is larger than a certain threshold, γ ,³ we define the i^{th}

³This value is just taken as a lower bound with 2% for now, further analysis can be done with other values.

AS path as sub-optimal, \overline{RTT}_{ψ}^i for this pairing:.

$$\overline{RTT}_{\psi}^i \geq \overline{RTT}_{AB}^{min} * \gamma \quad (1)$$

where $\overline{RTT}_{AB}^{min} = \min(\overline{RTT}_{AB}^1, \overline{RTT}_{AB}^2, \dots, \overline{RTT}_{AB}^n)$ and i is the particular ASPL value with $(1, 2, \dots, n)$ and $i \leq n$.

Albeit a simple definition, with our definition for sub-optimality we can shed light on the collected data to identify sub-optimal ASPLs from optimal ones. Thus, Figures 7-9 illustrate sub-optimal and optimal ASPLs for nodes that have 4, 3 and 2 ASPLs towards their destinations, respectively. We should also note that we focus on the same destination-source pairings because different destinations from the same node may normally be exposed to different paths.

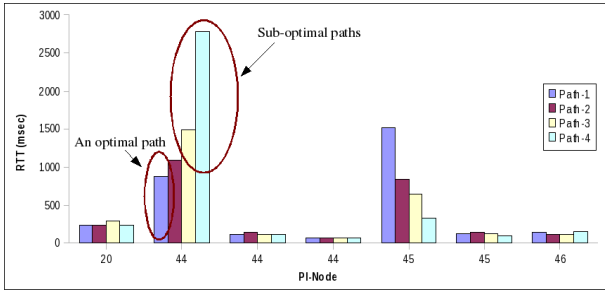


Fig. 7. End-to-end average latency of sub-optimal AS paths from nodes that used 4 different AS paths toward their destinations.

In these figures, we are able to see which nodes experienced sub-optimal AS paths. The same numbers on the x-axis represents AS paths towards different destinations from the same node. This reflects the fact that different destinations from the same node may be experiencing sub-optimal paths as explained above.

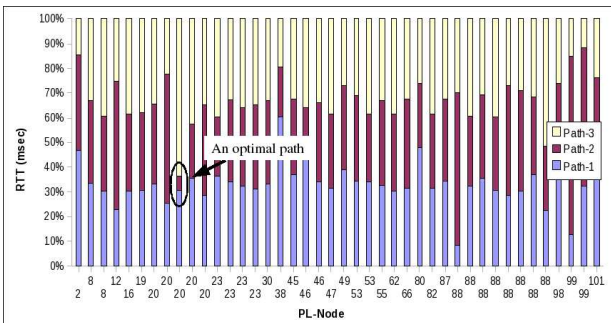


Fig. 8. End-to-end average latency of sub-optimal AS paths from nodes that used 3 different AS paths toward their destinations.

VI. OTHER POSSIBLE DEFINITIONS

In our observations, we saw cases where different ASPLs on the same connection had huge latency differences. Therefore, the basic definition we provided in the previous section could also be extended easily to consider cases where the end-to-end latencies from nodes exhibit huge RTT variability or huge deviations from the mean. Since the variance and the closely-related standard deviation are measures of how spread out a

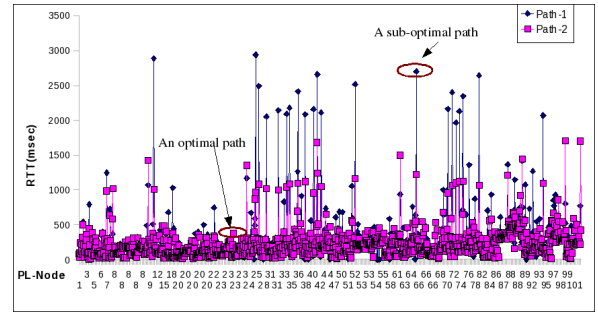


Fig. 9. End-to-end average latency of sub-optimal AS paths from nodes that used 2 different AS paths toward their destinations.

distribution is, we would be able to present a fine-tuned sub-optimality definition for the cases where source and destination pairs experience variable statistical measures. To illustrate this, in Figure 10, we present sub-optimal and optimal ASPLs for nodes that use 4 ASPL towards their destinations, this time with RTT variability. Due to space limitations, RTT variability of nodes that use 3 and 2 ASPLs towards their destinations were not shown.




Fig. 10. End-to-end latency variability of sub-optimal AS Paths from nodes that used 4 different AS-Paths toward their destinations.

VII. FREQUENCY OF SUB-OPTIMAL AS PATHS

Based on the definition given in the previous section, this section quantifies how frequent, (which we refer to as *sub-optimality ratio*), the nodes experienced sub-optimal AS paths during our measurements.

As we have described earlier, for a given source-destination pair and set of observations, we identify sub-optimal paths and compare them with the optimal one. If a path is not optimal, it contributes to the sub-optimal paths. We computed the percentage distribution of sub-optimal AS paths from nodes that use 4, 3, and 2 different AS paths towards their destinations using this approach.

Sub-optimality ratios for paths that use different ASPLs are shown in Figures 11-13. The figures also depict the average sub-optimality ratio for these cases. The packets that were transmitted over 4 different ASPLs had the average sub-optimality ratio of 79%. Similarly, the ones that are carried over 3 and 2 different ASPLs had average sub-optimality ratios of 57% and 41%, respectively. In other words, source-destination pairs that have varying ASPLs experienced at worst

almost 80% sub-optimality. Moreover, it is important to note that overall 82% of the fluctuating paths and 9% of all the traces between source-destination pairs faced sub-optimal AS paths.




Fig. 11. Percentage distribution of sub-optimal AS paths from nodes that used 4 different AS paths toward their destinations, average = 79%.

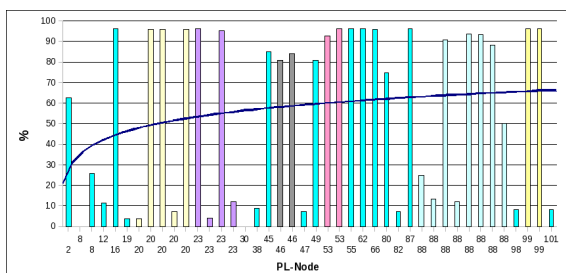


Fig. 12. Percentage distribution of sub-optimal AS paths from nodes that used 3 different AS paths toward their destinations, average = 57%.

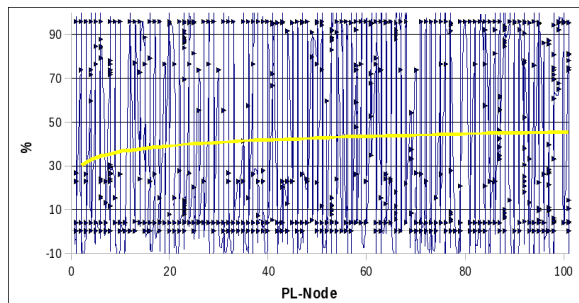


Fig. 13. Percentage distribution of sub-optimal AS paths from nodes that used 2 different AS paths toward their destinations, average = 41%.

VIII. CONCLUSION AND FUTURE WORK

In this paper the impact of sub-optimal AS paths has been investigated from the end-users' perspective. We collected data from more than 100 PlanetLab nodes with a trace-driven method. Assuming that the traces reflected the actual BGP behavior, we then analyzed the data to identify how many ASPLs the sources used towards their destinations.

We have shown that ASPLs vary greatly. All of the source nodes experienced some path differences. Overall, ASPL values showed fluctuations between 11% and 30% of the times

based on the two sets of data collected at different times. In the worst case, during our measurements, some nodes experienced different AS paths 70% of the time. We observed the largest difference in ASPL was as high as 6 with an average of 2.5. We have presented real cases where the ASPL and latency values are related, inversely related, and not related. For instance, a lower ASPL does not yield a lower latency always from analysis of the individual end-to-end pairings. In all the cases, latency fluctuates both in the reverse and the forward paths, however the average values of the latency on the forward and reverse paths between the end points are close.

Finally, as there is not a unique definition for sub-optimality, we provided a simple definition for sub-optimal ASPLs. We later quantified our measurements against our definition to determine how frequent the paths between the same end hosts experienced were sub-optimal AS paths. Our results showed that overall 82% of the fluctuating paths and 9% of all the traces between source-destination pairs faced sub-optimal AS paths. Furthermore, how the definition of sub-optimality could also be extended easily to consider cases where the end-to-end latencies from nodes exhibit huge RTT variability was discussed.

As future work, we plan to extend the basic definition of sub-optimality to include various other metrics. Another interesting future study from our data would be to identify where the AS path differences occur most of the time; in other words, how far the source is usually away from the point the AS path varies over sub-optimal paths. Although this would primarily depend on how ASs are interconnected, we expect some meaningful results.

REFERENCES

- [1] Nsf nets find initiative. [Online]. Available: www.nets-find.net
- [2] L. Gao and F. Wang, "The extent of as path inflation by routing policies," Nov. 2002.
- [3] N. Spring, R. Mahajan, and T. Anderson, "Quantifying the causes of path inflation," in *ACM SIGCOMM*, 2003.
- [4] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson, "The end-to-end effects of internet path selection," in *Proc. 1999 ACM SIGCOMM Conference*, August 1999, pp. 289–299.
- [5] PlanetLab. Planetlab web site. [Online]. Available: <http://www.planetlab.org>
- [6] L. Gao and J. Rexford, "Stable internet routing without global coordination," in *Proceedings of ACM SIGMETRICS*, 2001.
- [7] F. Wang and L. Gao, "Inferring and characterizing internet routing policies," in *ACM SIGCOMM IMC*, 2003.
- [8] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding bgp misconfigurations," in *Proceedings of SIGCOMM*, 2002.
- [9] N. Feamster and H. Balakrishnan, "Detecting bgp configuration faults with static analysis," in *Proc. 2nd Symp. Networked Syst. Design and Implementation*, May 2005.
- [10] D. McPherson, V. Gill, D. Walton, and A. Retana, "Border gateway protocol (bgp) persistent route oscillation condition," in *Internet Information RFC 3345*, August 2002.
- [11] A. Basu, L. Ong, B. Shepherd, A. Rasala, and G. Wilfong, "Route oscillations in i-bgp with route reflection," in *ACM SIGCOMM*, 2002.
- [12] R. Oliveira, B. Zhang, D. P. Izhak-Ratzin, and L. Zhang, "Quantifying path exploration in the internet," in *Internet Measurement Conference*, Rio de Janeiro, Brazil, October 25–27 2006.
- [13] A. facility for distributed Internet debugging and measurement. (2007) Scriptroute web site. [Online]. Available: www.scriptroute.org
- [14] Communications systems center (csc) web site. [Online]. Available: <http://www.csc.gatech.edu>