On the Impact of Route Monitor Selection

Ying Zhang *Univ. of Michigan*
Zheng Zhang *Purdue Univ.*
Z. Morley Mao *Univ. of Michigan*
Y. Charlie Hu *Purdue Univ.*
Bruce M. Maggs *Carnegie Mellon and Akamai Tech.*

ABSTRACT

Several route monitoring systems have been set up to help understand the Internet routing system. They operate by gathering real-time BGP updates from different networks. Many studies have relied on such data sources by assuming reasonably good coverage and thus representative visibility into the Internet routing system. However, different deployment strategies of route monitors directly impact the accuracy and generality of conclusions.

Our work is the first to critically examine the visibility constraints imposed by the deployment of route monitors on various applications. We study the difference due to diverse deployment schemes on three important classes of applications: (1) discovery of relatively stable Internet properties such as the AS topology and prefix to origin AS mappings, (2) discovery of dynamic routing behavior such as IP prefix hijack attacks and routing instability, and (3) inference of important network properties such as AS relationships and AS-level paths. We study several simple schemes of route monitor selection and provide insights on improving monitor placement.

Categories and Subject Descriptors: C.2.2 COMPUTER-COMMUNICATION NETWORKS: Network Protocols

General Terms: Measurement, Experimentation

Keywords: BGP, Internet measurement

1. INTRODUCTION

There exist several public route monitoring systems, such as Route Views [1] and RIPE [2], which have been deployed to help understand and monitor the Internet routing system. These monitoring systems operate by gathering real-time BGP updates and periodic BGP table snapshots from various ISP backbones and network locations to discover dynamic changes of the global Internet routing system. Various research studies have been conducted relying on these data, including network topology discovery [3], AS relationship inference [4, 5, 6, 7, 8], AS-level path prediction [9, 10], BGP root cause analysis [11], and several routing anomaly detection schemes. Most of them process the routing updates from the route monitoring system to study the dynamic routing behavior.

IMC'07, October 24-26, 2007, San Diego, California, USA.

These studies relying on BGP routing data usually assume that data from the route monitoring systems is reasonably representative of the global Internet. However, no existing work has studied the limitations of route monitoring systems and the visibility constraint of different deployment scenarios. For example, recent work using these data to detect malicious routing activities, such as address hijacking [12, 13, 14, 15] could potentially suffer from evasion attacks similar to those affecting traffic monitoring systems [16]. The accuracy of such anomaly detection schemes depend heavily on the coverage of the route monitoring system. The limitation of the route monitor system is critical for any system relying on BGP data from multiple vantage points.

It is usually impossible to obtain routing data in real time from every network due to the scalability issue and privacy concern. Obtaining one feed from one AS often provides a restricted view given there are many routers in an AS, each with a potentially different view of routing dynamics. Additional BGP feeds are useful for detecting routing anomalies, traffic engineering, topology discovery and other applications. But adding an additional feed usually requires interacting with a particular ISP to set up the monitoring session. Therefore, an urgent question is to understand the generality and representativeness of the given monitor system, and to understand how to select monitor locations to maximize the overall effectiveness of the route monitoring system.

Some existing work [17, 18, 19] studied the limitation of existing monitor placement and monitor placement algorithms [20] in terms of topology discovery. In this work, we study the impact of monitor network location constraints on various research work in the Internet routing community. We are the first to examine the visibility constraints imposed by the deployment of route monitors, impacting a diverse set of applications. To understand the difference among current deployment settings, we analyze three deployment scenarios: all Tier-1 ISPs only, Route Views and RIPE setup, and a setup combining many public and private vantage points. We further study four simple schemes of network monitor selection and the resulting impact on multiple metrics based on the applications using the data. Our analysis shows that current public monitors already provide good coverage in various applications we study.

The paper is organized as follows. In Section 2 we introduce the methodology of our study, followed by a short discussion comparing three deployment scenarios in Section 3. We study in detail several different monitor selection schemes in Section 4 and conclude in Section 5.

2. METHODOLOGY

In this section, we describe the methodology of our study, including the data we used and various metrics for comparing monitor selection schemes motivated by several common but important

applications using BGP data.

2.1 Route monitor locations

The BGP data we used in our study are collected from around 1000 monitoring feeds, including public data sources such as Route Views [1] and RIPE [2], feeds from the local ISP, and data from private peering sessions with many other networks, covering more than 200 distinct ASes, which are not in the public feeds. In the remainder of the paper, we use the term *monitoring feed* to refer to a BGP data source from a particular router. We define a *vantage point* to be a distinct AS from which we collect BGP data from. Note that feeds from different routers in the same AS may provide different information, and we leave the study of the difference between feeds in the same AS for future work. We use one monitoring feed from one vantage point. For ease of comparison across vantage points, we only choose feeds with default-free routing tables (with entries for all prefixes), and create a data set called *LargeSet* consisting of data from 156 ASes for our subsequent analysis.

The BGP updates are collected from a set of route monitors, each of which establishes peering session with one router in each network being monitored. Note that our study is inherently limited by the BGP data we have access to and we attempt to draw general conclusions independent of the data limitation. Although the BGP data from all available monitors is still not the ground truth for the whole network, we study different applications using data from different sampling strategies and compare with this *LargeSet*. Developing more intelligent monitor placement algorithms is part of future work.

To understand static network properties, instead of using a single table snapshot from each feed, we combine multiple snapshots taken at different times with routing updates from each feed whenever available. This helps improve the topology completeness as many backup links are only observable during transient routing changes. We use two snapshots of tables from each monitoring feed including feeds from about 100 ASes, along with six months of updates and tables from Route Views, RIPE and a local ISP from May 2006 to Oct. 2006. The resulting network topology contains 25,876 nodes(ASes) and 71,941 links. We list the properties of current peers that Route Views and RIPE have in Table 2.

To compare different deployment strategies, we construct three sets of realistic deployment scenarios. First, to understand the visibility of the core of the Internet, we select only 9 well-known Tier-1 ISPs to be monitors, including AS numbers: 1239, 174, 209, 2914, 3356, 3549, 3561, 701, and 7018. Second, we use only feeds from commonly used Route Views and RIPE. Third, we use LargeSet to obtain the most complete topology from all available data. We denote the three deployment scenarios as Tier-1, Route Views, and LargeSet, respectively.

We focus on three types of applications relying on BGP data, namely (1) discovery of relatively stable Internet properties such as the AS topology and prefix to origin AS mappings, (2) discovery of dynamic routing behavior such as IP prefix hijack attacks and routing instability, and (3) inference of important network properties such as AS relationships and AS-level paths. Note that the first two applications simply extract properties directly from the routing data. The performance of the third one depends not only on the data but also the algorithm used for inference. We describe these applications in more detail below.

2.2 Discovery of static network properties

BGP data is an important information source for understanding the Internet topology. Very basic network properties are critical for understanding the Internet routing system. These properties include AS connectivity, IP prefix to origin AS mappings, identifying stub AS information and its provider's information, multihomed ASes, and AS path information. Intuitively, including vantage points from the core is more beneficial as a larger number of network paths traverse the core networks. Previous work [21, 22, 3] has shown the influence of data sources besides BGP table data, *e.g.*, traceroute data and routing registries, on the completeness of inferred AS topology. We extend this analysis to two other properties: (1) multihomed stub ASes to understand edge network resilience and potentially increased churn in updates, and (2) AS paths, which are difficult to infer.

2.3 Discovery of dynamic network properties

Dynamic properties of the routing system are of strong interest for studying routing instabilities, *e.g.*, due to misconfigurations, and detecting anomalies. Understanding such properties is useful for troubleshooting and identifying possible mitigation to improve routing performance. We focus on two representative applications here: monitoring routing instability and IP prefix hijack attack detection.

Routing instability monitoring: Routing updates are a result of routing decision changes in some networks caused by events such as configuration modifications, network failures, and dynamic traffic engineering. Comprehensively capturing Internet routing changes is useful for important applications like troubleshooting, routing health monitoring, and improved path selection.

IP prefix hijacking detection: One of the original goals of the public route monitoring systems in Route Views and RIPE is troubleshooting. Nowadays they are increasingly used for the timely detection of malicious routing activities such as prefix hijacking attacks. Current hijack detection systems in control plane [13, 12] rely on detecting inconsistency in observed BGP updates across vantage points. However, the detection system may not detect all attacks due to limited visibility. In this work, we study the impact of different monitoring deployment setups on the detection coverage.

Intuitively, an attack is missed if no vantage point of the monitoring system adopts the malicious route. Thus, we define attack evasion as follows. For a monitoring system $SM = m_1, m_2, \ldots, m_n$ with n monitors, given an attacker A, a victim V, and the hijacked prefix p, if $\forall i$, $Pref^A_{m_i}(p) < Pref^V_{m_i}(p)$, where $Pref^A_{m_i}(p)$ is the route preference value for p announced from A as observed by m_i, then attacker A can hijack V's p without being detected.

2.4 Inference of network properties

The third class of application studied relates to properties inferred from the above basic properties from BGP data.

AS relationship inference: There is much work [6, 7, 4, 8, 5] on inferring AS relationships from BGP AS paths. Knowing commercial relationships among ASes reveals network structure and is important for inferring AS paths. In this work, we study the commonly-used, Gao's degree-based relationship inference algorithm [8].

AS-level path prediction: Accurately predicting AS paths is important for applications such as network provisioning. In this work, we compare two path prediction algorithms under various monitor deployment settings. We use the recent algorithm [9] which makes use of the inferred AS relationships, and study both profit-driven and shortest-path-based route selection. For the profit-driven policy, the route selection prefers customer routes to peering routes and over to provider routes. Note that predicted paths for both approaches need to conform to relationship constraints [8]. We also study the recent work [23] which does not use AS relationships but

Category	Tier-1	Route Views	LargeSet
Number of ASes	25732	25801	25876
Number of AS links	51223	56000	71941
Profit-driven prediction	34%	39%	43%
Length-based prediction	67%	76%	73%

Table 1: Comparison among three deployment scenarios.

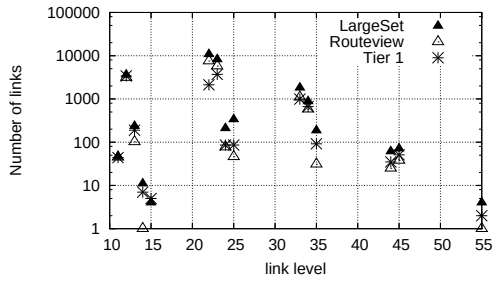

Figure 1: Distribution of observed links across tiers.

instead exactly matches observed paths.

To improve scalability, we eliminate *stub AS nodes*, or customer ASes that do not provide any transit to other ASes. The graph without stub nodes contains only 4426 (16% of all nodes) and 25849 links (15% of all links). For completeness, we also simulate the path prediction to 50 randomly sampled stub ASes. We include these 50 stub ASes and their links.

3. DEPLOYMENT SCENARIO ANALYSIS

We first analyze the three deployment scenarios, Tier-1, Ruter Views, and LargeSet defined in Section 2.1. We study the impact of these three settings on applications of AS topology discovery, AS relationship inference, and AS-level path prediction.

Table 1 summarizes the comparison across the three setups. Confirming previous studies [17, 19], we find that the largest monitor set, LargeSet, observes much more links but only slightly more non-private ASes. The additional ASes in the LargeSet are mostly at the edge. Using Gao's degree-based relationship inference algorithm, we compare the accuracy of inferred paths comparing with paths in BGP data in terms of path length. Note that the improvement is small for path prediction with increasing vantage points. Interestingly, using the largest data set lowers the length-based prediction accuracy. These results imply that Gao's algorithm is reasonably stable with changes in the BGP data.

We list the network properties of current peers of both Route Views and LargeSet in Table 2. We use the tier definition specified in previous work [7]: Tier-1 means closest to the core Internet and Tier-5 is associated with stub or pure customer ASes. We also analyze each AS in the aspects of geographic location, the number of IP addresses it announces, its degree and its customers. The additional ASes in LargeSet are mainly Tier-2 ASes in US, with large number of addresses and degree.

To understand which links are identified using a larger data set, we plot in Figure 1 the topological location of links in each data set. The X-axis indicates the link level, defined by the tier value of the two ASes associated with the link sorted in increasing order. For example, there are 10 links observed from LargeSet between nodes in tier-1 and tier-4 at the X value of 14. The hierarchy level for each node is assigned according to the relationship inferred using all the data available. As expected, the additional benefit of observed links

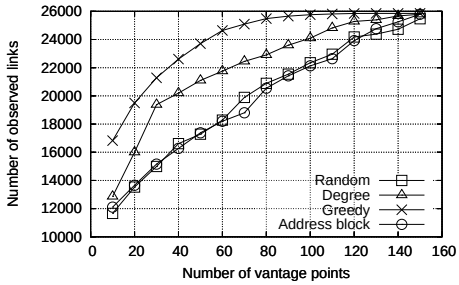

Figure 3: Number of observed links

are mostly at the edge.

4. MONITOR SELECTION ANALYSIS

In the previous section, we have observed some differences and similarities among the three realistic deployment settings. To delve deeper, we apply four simple schemes to identify the incremental benefit and even possible negative effects of adding monitors for a wider set of applications.

4.1 Monitor selection schemes

Our candidate set of monitors consists of all BGP feeds we have access to. We study the following four ways of adding monitors.

Random based: monitor nodes are selected randomly.

Degree based: monitors with the largest node degree are selected first based on the entire data set. Node degree means the number of neighbors each AS has.

Greedy link based: at any time, the next monitor is selected with the largest number of unobserved links, given the set of already selected monitors.

Address block based: without relying on all the data, monitors in the ASes that originate the largest number of IP addresses are selected with random tie breaking.

4.2 Discovery of static network properties

To fully understand how each scheme works, we study the topological distribution of the monitors selected based on the tier classification, with the first three tiers shown in Figure 2. We observe that as expected the address-block-based scheme always selects the Tier-1 nodes first as they usually announce largest number of addresses. For Tier 2 and Tier 3, there is little difference among the schemes.

We first show that the observed link count increases with vantage point in Figure 3. Confirming previous studies [20], the increase of links from 80 vantage points can be twice as the links observed from one. The greedy-based scheme performs best as expected, followed by the the degree-based one. Interestingly, the address block based scheme is no better than random selection. This is likely due to the fact that most ASes in our candidate set contribute a similar number of links.

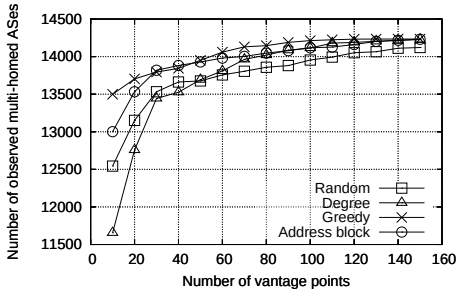

Figure 4: Number of observed multi-homing stub ASes

Data	Tier				Geographic location				Address			Degree			Customer		
	1	2	3	4	Europe	Asia	Africa	America	Min	Avg	Max	Min	Avg	Max	Min	Avg	Max
Route Views	9	40	58	12	37	4	1	77	156	65313	1561473	3	247	2922	0	112	2899
LargeSet	9	82	60	5	46	4	1	105	156	116989	1561473	1	344	2922	0	177	2899

Table 2: Statistics of the monitors.

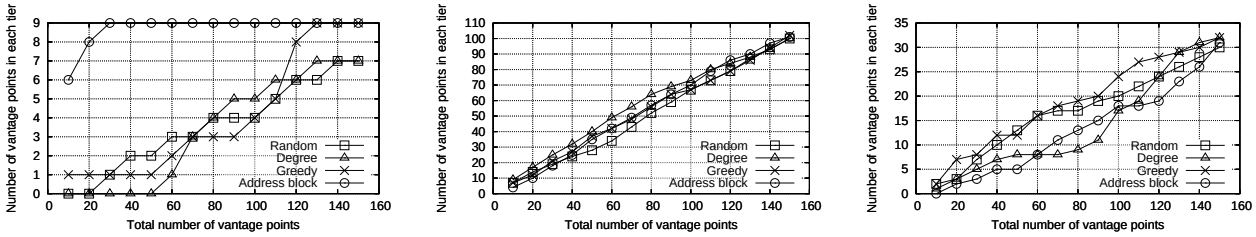

(a) Tier 1 (b) Tier 2 (c) Tier 3

Figure 2: Monitor distribution in each tier for different monitor selection schemes.

Next, we study the prevalence of multi-homing at edge networks for network redundancy as shown in Figure 4. The greedy-based selection again performs best as additional edge links for multi-homed stub ASes are more likely discovered. The difference between random and greedy can be up to several hundred, indicating that we may not have a complete set of multi-homed customer ASes.

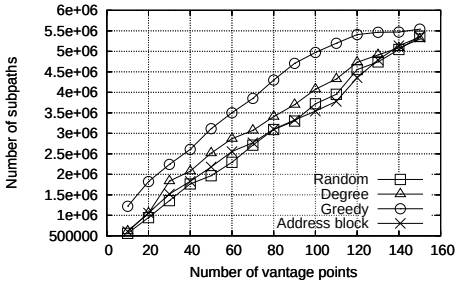

Figure 5: Observed AS path count (including subpaths)

As we have shown, accurate AS path prediction is still quite challenging. One way to lower the difficulty is to collect as many empirically observed AS paths as possible, as depicted in Figure 5. Greedy performs the best, followed by the degree-based scheme. Note that the absolute difference in observed paths for the same number of vantage points among various schemes can be as large as one million.

4.3 Discovery of dynamic network properties

We study two applications relying on monitoring of dynamic routing events.

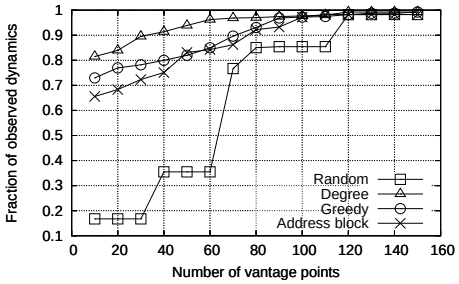

Figure 6: Fraction of observed routing events

Routing instability monitoring: A single network event such as link failure can trigger routing updates from many networks. We study how to monitor as many routing events occurring on the Internet as possible. Figure 6 shows the fraction of BGP routing events observed by the set of vantage points selected. Notice there is a huge difference between random selection and the other three schemes, indicating that vantage points associated with core networks (*i.e.*, with high degree and many links, and originating many addresses) are more likely to observe network instabilities.

IP prefix hijacking detection: Intuitively, more monitors enable more diverse paths to be observed. Therefore, the IP prefix hijacking detection system has a higher chance of detecting all hijacks. However, based on our simulations, we observe there still exist attacker-victim pairs that can evade detection even using all the monitoring feeds we have access to. Studying to what extent attackers can evade detection is important for knowing the limitation of current detection systems due to visibility constraints.

The main metric we study is the number of attacker-victim pairs that can evade detection. As shown in Figure 7(a), with 10 nodes deployed in the random scheme, 0.35% of all possible attacker-victim pairs can evade the detection, which is the worst case we observe from our simulation. We also show changes in the average number of evading attackers for each victim in Figure 7(b), and in the average number of victims an attacker can attack without being detected in Figure 7(c). Overall, address block scheme performs similar to the random scheme, while greedy performs the best in most cases.

4.4 Inference of network properties

In the following we analyze the effect of vantage point selection on inference of AS relationships and AS-level paths. We study two algorithms for path inference.

4.4.1 AS relationship inference and path prediction

We first study commonly used path inference algorithms relying on AS relationships as indicated in Table 1. In particular, we apply Gao's degree-based relationship inference scheme [8] and then predict paths enforcing the AS relationships. Figure 8 shows that, surprisingly, as the number of monitors increases, the accuracy may decrease compared with observed AS paths.

We conjecture this may be caused by the nature of the degree-based relationship inference algorithm. The algorithm determines the AS relationships based on the relative degree values of AS nodes within an AS path. The topology obtained from the vantage points tends to be quite complete already in terms of relative degree

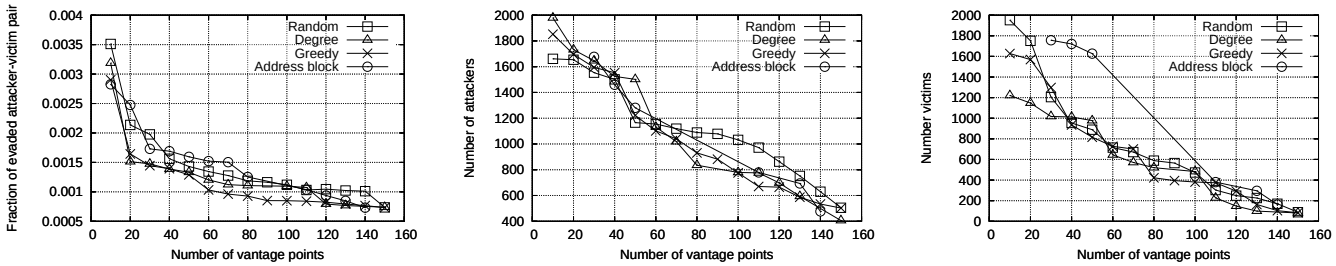

(a) Number of attacker-victim pairs (b) Number of attackers per victim (c) Number of victims per attacker

Figure 7: IP prefix hijacking evasion under different monitor selection schemes.

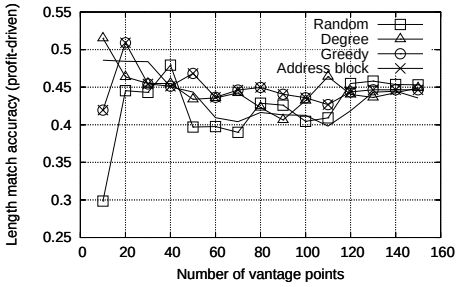

Figure 8: Profit-driven path prediction accuracy (length match).

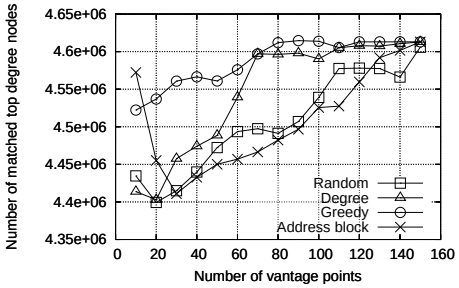

Figure 9: Number of matched top degree AS in all observed AS paths

information. As more vantage points are added, more noise may be introduced causing inaccuracies in inferred AS relationships.

To further understand this, we analyze the changes in the top degree node per path to explain why the increase in number of vantage points does not always result in increased accuracy. Based on the degree of each node observed in the topology using all data available, we identify the top degree AS for each observed AS path. From each set of vantage points we also locate the top-degree node. We then examine for each monitor data set, the fraction of matched top ASes for all AS paths compared with the case for the complete topology, as shown in Figure 9. The fluctuation in the graph indicates that additional BGP data does not consistently improve the estimation of the top-degree nodes in each path.

We emphasize that we have made an important observation: BGP data from more vantage points may not necessarily increase the accuracy of inferred network properties. The inference algorithm [8] is based on degree, which may vary in different selection of monitors: the further away an AS is to the monitor, the more incomplete the observed degree is. We point out that developing inference algorithms that are less sensitive to available data feeds but also more fully utilize the data available is important in this area. We also observe that profit-driven path prediction as shown in Figure 8 actually performs worse than length-driven prediction. This can be possibly explained by the fact that profit-driven path

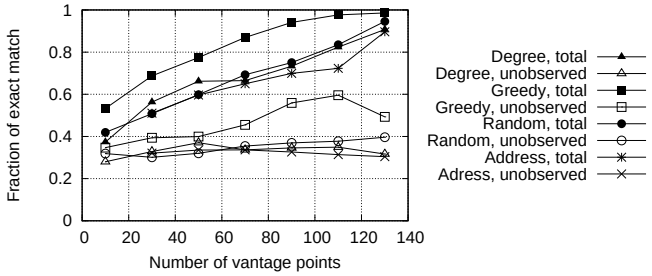

Figure 10: Sampled path prediction accuracy: exact matching (new algorithm)

selection is more sensitive to the impact caused by inaccurate AS relationship inference.

Besides accuracy, we also perform other sanity checks for inferred relationships. Two metrics are used: first, some observed paths are considered as invalid based on the inferred relationships. The fraction of such invalid paths can be used as an indication of inaccurate AS relationship inference. We found that the number of observed invalid paths slightly decreases as the number of vantage point increases. Second, for some node pairs no valid paths are predicted. Such disconnected node pairs can be used as another metric of relationship inference inaccuracy. The number of invalid paths generally decreases with more vantage points as expected; similarly, the number of AS pairs with valid paths increases with vantage point. Greedy is again observed to be the best for identifying valid paths.

4.4.2 AS-relationship-independent path prediction

In the following, we study the behavior of a recent proposed path prediction algorithm [23] that does not rely on AS relationships for prediction. For each deployment scenario, we use all observed AS paths to construct an initial topology model, and then use observed AS paths of 50 random prefixes to iteratively train the topology model using the refinement algorithm specified. The trained model is used to predict the paths from any AS to the same 50 prefixes.

To evaluate the accuracy of the predicted paths, we consider three sets of paths. The first set, $total$, is the AS paths to the 50 prefixes observed from the total default-free 165 vantage point ASes. The second set $observed$ is the AS paths to the 50 prefixes observed from all the monitors a particular deployment scenario. The third set $unobserved$ is the complementary set of $observed$ in $total$. The algorithm always produces a perfect match on the $observed$ set. Therefore, we use the other two sets for evaluation. Note that the path prediction in Section 4.4.1 is evaluated on $observed$ instead.

Figure 10 shows the fraction of paths in $total$ and $unobserved$ that match the predicted paths. Overall, all schemes accurately predict $28\% \sim 60\%$ of the unobserved paths in all scenarios. This

number is lower than those in [10] because we do not include suffix subpaths in the evaluation sets, and hence do not give partial credits to the paths that partially match the prediction. The match percentage on $unobserved$ generally does not increase with more monitors. The above observations show the difficulty of path prediction: predicting an unobserved path does not benefit much from observing its subpaths or its reverse path. The figure also shows that the accuracy on the $total$ set improves with more monitors, which is a result of more paths being observed. Greedy performs best on the $total$ set because this scheme observes most paths.

5. CONCLUSIONS

In this work we illustrate the importance of route monitor selection on various applications relying on BGP data: discovery of static network properties, discovery of dynamic network properties and inference of network properties. For the first class, more vantage points generally improve completeness and accuracy of the topological properties studied. We show that it is important to take into consideration possibility of evasion due to visibility constraints for detecting routing attacks. The coverage of routing instability monitoring varies significantly across different monitor selection. Finally, we take the first step at analyzing how various AS path inference algorithms and a commonly used AS relationship inference algorithm are impacted. Our work motivates future work in the area of identifying algorithms less sensitive to the input routing data set.

Acknowledgment

We would like to thank the anonymous reviewers for their helpful comments. We are grateful to our shepherd Renata Teixeira for her detailed and insightful comments. Our work was supported by the National Science Foundation under grants CAREER-0643612, CAREER-0238379, CNF–0435382 and by the Department of Homeland Security under grant W911NF-05-1-0415.

6. REFERENCES

[1] "University of Oregon Route Views Archive Project." `www.routeviews.org`.

[2] "Ripe NCC." `http://www.ripe.net/ripencc/pub-services/np/ris/`.

[3] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy, "A systematic framework for unearthing the missing links: Measurements and Impact," in *Proc. of NSDI*, 2007.

[4] X. Dimitropoulos and G. Riley, "Modeling Autonomous System Relationships," in *Proc. of PADS*, 2006.

[5] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley, "AS Relationships: Inference and Validation," *ACM Computer Communication Review*, vol. 37, no. 1, 2007.

[6] G. Battista, M. Patrignani, and M. Pizzonia, "Computing the Types of the Relationships Between Autonomous Systems," in *Proc. IEEE INFOCOM*, March 2003.

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

[8] L. Gao, "On Inferring Autonomous System Relationships in the Internet," in *Proc. IEEE Global Internet Symposium*, 2000.

[9] Z. M. Mao, L. Qiu, J. Wang, and Y. Zhang, "On AS-Level Path Inference," in *Proc. ACM SIGMETRICS*, 2005.

[10] W. Muhlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig, "Building an AS-Topology Model," in *Proc. of ACM SIGCOMM*, 2006.

[11] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs, "Locating Internet Routing Instabilities," in *Proc. ACM SIGCOMM*, 2004.

[12] X. Hu and Z. M. Mao, "Accurate Real-time Identification of IP Prefix Hijacking," in *Proc. of IEEE Security and Privacy*, 2007.

[13] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A Prefix Hijack Alert System," in *Proc. of USENIX Security Symposium*, 2006.

[14] M. Lad, R. Oliveira, B. Zhang, and L. Zhang, "Understanding resiliency of internet topology against prefix hijack attacks," in *Proc. of DSN*, 2007.

[15] J. Karlin, J. Karlin, S. Forrest, and J. Rexford, "Pretty Good BGP: Improving BGP by Cautiously Adopting Routes," in *Proc. of ICNP*, 2006.

[16] Y. Shinoda, K. Ikai, and M. Itoh, "Vulnerabilities of passive internet threat monitors," in *14th USENIX Security Symposium*, August 2005.

[17] P. Barford, A. Bestavros, J. Byers, and M. Crovella, "On the marginal utility of network topology measurements," in *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, 2001.

[18] D. Raz and R. Cohen, "The internet dark matter: on the missing links in the as connectivity map," in *Proc. IEEE INFOCOM*, 2006.

[19] B. Zhang, R. Liu, D. Massey, and L. Zhang, "Collecting the internet as-level topology," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 1, pp. 53–61, 2005.

[20] R. Oliveira, M. Lad, B. Zhjang, D. Pei, D. Massey, and L. Zhang, "Placing BGP Monitors in the Internet." Technical Report, UCLA CS Department, TR 060017, May 2006.

[21] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Towards capturing representative AS-level Internet topologies," *Computer Networks*, 2004.

[22] B. Zhang, R. Liu, D. Massey, and L. Zhang, "Collecting the Internet AS-level Topology," *ACM SIGCOMM Computer Communication Review, special issue on Internet Vital Statistics*, 2005.

[23] W. Muhlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig, " Building an AS-topology model that captures route diversity," in *Proc. ACM SIGCOMM*, 2006.