

On Detection of Anomalous Routing Dynamics in BGP

Ke Zhang¹, Amy Yen¹, Xiaoliang Zhao³, Dan Massey³, S. Felix Wu¹, and Lixia Zhang⁴

¹ University of California, Davis, CA, U.S.A
{kezhang, ahyen, sfwu}@ucdavis.edu

² USC/ISI, Arlington, VA, U.S.A
{xzhaol, masseyd}@isi.edu

³ University of California, Los Angeles, CA, U.S.A
lixia@cs.ucla.edu

Abstract. BGP, the de facto inter-domain routing protocol, is the core component of current Internet infrastructure. BGP traffic deserves thorough exploration, since abnormal BGP routing dynamics could impair global Internet connectivity and stability. In this paper, two methods, signature-based detection and statistics-based detection, are designed and implemented to detect BGP anomalous routing dynamics in BGP UPDATEs. Signature-based detection utilizes a set of fixed patterns to search and identify routing anomalies. For the statistics-based detection, we devise five measures to model BGP UPDATEs traffic. In the training phase, the detector is trained to learn the expected behaviors of BGP from the historical long-term BGP UPDATEs dataset. It then examines the test dataset to detect "anomalies" in the testing phase. An anomaly is flagged when the tested behavior significantly differs from the expected behaviors. We have applied these two approaches to examine the BGP data collected by RIPE-NCC servers for a number of IP prefixes. Through manual analysis, we specify possible causes of some detected anomalies. Finally, comparing the two approaches, we highlight the advantages and limitations of each. While our evaluation is still preliminary, we have demonstrated that, by combining both signature-based and statistics-based anomaly detection approaches, our system can effectively and accurately identify certain BGP events that are worthy of further investigation.

1 Introduction

As the size, complexity, and connectivity of the Internet increase, the analysis of operational BGP dynamics becomes more and more challenging. First, because a huge amount of BGP UPDATE traffic is generated in a single domain everyday, operators are not able to conduct thorough analysis on the whole logged BGP dataset. Second, even for a single BGP event, the root cause analysis could be extremely hard. Sometimes, an experienced network administrator needs to, if possible, access the information in the core of the service networks, even from

different administrative domains, in order to identify potential faults or configuration problems. Since the process of problem and fault analysis can be highly expensive, it is critical to put our focus on a small set of valuable network events. In other words, given a large set of BGP update messages, can we accurately categorize them as "normal" or "abnormal"? With this categorization, we can then spend our precious resources mostly on the "abnormal" ones. Two criteria jointly define an anomaly. One criterion is related to BGP performance. For example, slow convergence for a router to reach a stable view of the Internet's available routes [1] belongs to this type, because the router announces many invalid routes to downstream BGP routers during the convergence process. The other criterion refers to a statistical anomaly (also called "relative anomaly")—the significant deviations of current routing behavior from expected routing behavior.

However, in practice, to our best knowledge, we do not have a systematic approach to consistently label a set of BGP events as normal or abnormal. Borrowing the techniques from intrusion detection area, we develop two approaches to detect BGP anomalous routing dynamics—signature-based detection and statistics-based anomaly detection. For signature-based detection, we devise a set of anomalous BGP routing update patterns to search for matching incidents in BGP UPDATEs data. For statistics-based anomaly detection, the long-term historical BGP UPDATEs datasets are used to train the detector to learn the statistical properties. Thereafter, we perform the anomaly detection on the short-term testing UPDATEs datasets. Following detection, we examine the anomalous routing incidents and explain why they should be categorized as anomalies and specify possible root causes of some incidents.

Analyzing the root causes for BGP dynamics is a very challenging task. Our work moves the first step towards this problem by providing the approaches to automatically locate the anomalous routing updates. Due to limited routing information we can acquire, the anomalies discussed in the paper are still speculative ones. However, results from manual examination show that these anomalies are worthy of further investigation. Thus, we believe the approaches are valuable in that they drastically reduce the search space from a large amount of BGP data to a small set of "abnormal" BGP events. Moreover, the signatures and statistics developed in these approaches can be used to analyze BGP data and quantitatively evaluate the "normality" of each BGP UPDATE.

The rest of the paper is organized as follows. Section 2 introduces the concepts of signature-based detection and statistics-based detection, and briefly reviews related work. Section 3 describes the BGP UPDATE dataset that we have used in the experiments. Section 4 and 5 present signature-based detection and statistics based detection respectively. Section 6 compares these two approaches, followed by conclusions in section 7.

2 Related Work

Signature-based detection and statistics-based detection are two major approaches in modern intrusion detection area. Signature based detection systems,

such as the Snort IDS, report an attack when a set of symptoms corresponding to a predefined attack signature is observed. Statistics-based intrusion detection flags as attacks any traffic that is unusual for that system. Several research has been carried out on BGP routing behavior. Labovitz et al. [2] showed that around 1996 unstable and pathological routing behaviors dominated the Internet. Later, they presented possible explanations for these anomalies [3]. Other BGP routing problems, such as slow convergence [1], persistent MED oscillation [4,5], have also been well examined. Also concerning abnormal BGP route changes, Wang et al. [6] proposed a path-filtering approach to validate the correct route changes for DNS prefixes. Teoh et al. developed an interactive visualization process to explore BGP data [7]. These works are complementary to our approaches described in this paper.

3 Dataset

The dataset we examine in the paper consists of BGP UPDATE messages collected by the Routing Information Service of RIPE [8]. The collector has multi-hop BGP sessions with 9 peer ASes located at different countries.

We examine BGP updates for a set of networks (IP prefixes) in different peers. Different prefixes may show different behaviors. Even for a single prefix, routing behaviors may be different from different observation points. As an initial step, we choose the following prefixes as samples. We choose 8 prefixes of either root DNS servers or gTLD servers because of the critical role of DNS service. 4 prefixes from Korean and China are selected because we attempt to examine the impact of SQL worm attack. Similar to previous work [9,10], we also select 4 prefixes of popular destinations and 4 prefixes from Department of Defense (DoD).

In addition, we removed duplicated updates due to known implementation problem in some vender's router [3], although duplicate updates are anomalous.

4 Signature-Based Detection

4.1 Patterns of Anomalous BGP Dynamics

A route announced by a BGP router is generally the best route at that moment. Comparing the consecutively announced routes, we can infer the route changes in that router's BGP routing table. In order to compare the consecutively announced routes, we assign a value corresponding to the preference of each route based on BGP route selection process, which is described in [11]. Based on the relative preference values of two consecutive routes, we define four terms.¹

¹ Since the BGP updates are collected through the EBGP sessions, we cannot acquire information on the following four aspects: LOCAL_PREFERENCE, source of the route (EBGP or IBGP), the IGP cost to the NEXT_HOP, and router ID. We can assign relative preference value by comparing the AS_PATH length, origin type and MED value of each consecutive route announcement.

Table 1. Signatures of BGP Update Burst

TYPE	Pattern	Examples	Indication
B	A sequence of updates with WD in the middle	<D,D,W,U,U> <D,W,U,W,U>	Transient failure followed by fast fail-over
C	A sequence of updates with only one preference fluctuation	<U,U,D,D,F> <D,D,U,F,U> <D,D,U,U,D>	Transient failure followed by fast fail-over OR Normal route changes
F	A sequence of updates with same preference	<F,F,F,F,F>	Anomaly in community attributes or aggregation or same length AS path oscillation OR Normal route changes

UP: if the second route is more preferable than the previous one, we label the second route as UP.

DOWN: if the second route is less preferable than the previous one, we label the second route as DOWN.

FLAT: if two routes have the same preference value, we label the second route as FLAT.

WD: if the second announcement is a route withdrawal, we label the second route as WD.

We define BGP update burst as a sequence of updates within a short time window. Formally, BGP update burst is K consecutive updates for the same prefix that space close together. The time interval between update messages is less than T and the average update rate $> \alpha$. In the experiments, we empirically set $K = 4$, $T = 240s$, and $\alpha = 1/90$. Examining different parameters is part of our future study.

Given a BGP update burst, we map the updates into a {UP, DOWN, FLAT, WD} sequence. We define 7 patterns for update sequences. Due to the limited space, we only show 3 types in the following. Other types are described in [12].

Type-B: If the update burst has a WD in the middle, it indicates a transient failure followed by a fast fail-over.

Type-C: If the update burst does not consist of WD and has only one <UP, DOWN> or <DOWN, UP> in the middle (the preference fluctuation only happens once in the sequence of updates), it indicates either a transient failure (or congestion) followed by a fast fail-over, or normal route changes.

Type-F: If the update burst consists of all the routes with the same preferences, it might be anomalous. These routes have the same length AS_PATH, the same origin types and the same MED values. Difference might lies in the content of AS_PATH or other attributes, such as community attributes, ATOMIC_AGGREGATE and AGGREGATOR. Since we cannot get the local preference of each route, these route changes might be legitimate. However, from discovered incidents(presented later in this paper), we believe this type of incident is speculative.

4.2 Experiment Results for Signature Detection

We have performed signature detection on 20 prefixes over the period from Feb. 2002 to Jan. 2003. We demonstrate here detected anomalies for four represen-




Fig. 1. Anomalies detected by Signature Detector

tative prefixes. Other prefixes have similar results. For each prefix, we choose to present results from 5 ASes (AS2914, AS3333, AS13129, AS3549, AS3257) out of 9 observed ASes. These four prefixes are the prefix for Yahoo.com, the DNS root Server-A prefix, a prefix for Department of Defense and a prefix from a University of China.

Figure 1 plots the number of the each type of incidents at five ASes. We notice that B,C,F are three major types of abnormal incidents. Through further analysis, we identify three different speculative incidents in type F.

We notice that the DoD prefixes has a total of 69 type-F incidents observed from 5 ASes. The pattern capturing these special sequences is frequent substitutions of AGGREGATOR. Due to limited information, we cannot verify whether or not this special behavior is normal BGP operation. However, the high rate of AGGREGATOR substitution, once per minute on average, deserves more attention from the operators. Since local AGGREGATOR changes should be restricted in the local area, and not be propagated to the outer networks. These anomalies indicate BGP operation in DoD networks violates this desired property.

AS3549 and AS3257 have more type-F sequences than other three ASes. We find that some Type-F update bursts for prefix 166.111/16 are due to community attribute changes. The change rate is very high. For example, in one case, AS 3549 changed the community attribute 4 times in 6 seconds. These frequent changes of community attribute generate a lot of BGP updates in a short period of time. If the downstream BGP router performs BGP route flap damping, the route announced by AS3549 would be suppressed. We performed BGP route flap damping (using default CISCO router's damping parameters) on the Jan. 2003 updates data. We find that 174 out of 250 effective updates would be suppressed, and the total suppression time in that month is 7.3 hours.

In addition, from AS3257, we observe that the DNS root server-A prefix has 17 type-F incidents, 7 out of which are oscillations of two routes. The two routes, {AS3257, AS1, AS10913, AS19836} and {AS3257, AS3356, AS10913, AS19836} have the same AS path length, the same origin type and the same MED value. They replace each other at least three times in a short time window. The possible root cause might be link flap, or transient link congestion or even other unknown reasons. Although we do not know the root cause, we believe

Table 2. Five Measures

Intensity Measure	BGP Updates Message Arrival Frequency Number of AS paths in a period
Categorical Measure	BGP Updates Type AS path Occurrence Frequency
Counting Measure	AS path Difference

this kind of incident should be anomalous because the frequent route changes can degrade packet forwarding.

5 Statistics-Based Anomaly Detection

We apply a statistics-based anomaly detection method, NIDES/STAT [13]. The NIDES/STAT algorithm monitors a subject's behavior on a computer system, and raises alarm when the subject's current (short-term) behavior deviates significantly from its expected behavior, which is described by its long-term profile. A subject's behavior is described by a set of detection measures. For each individual measure, there is a corresponding Q statistic. The historical profile records the frequency distribution of Q . For each measure, the corresponding S , derived from Q , is indicative of the degree of abnormality of the behavior with respect to that measure. T^2 summarizes the abnormality of many measures, reflecting the degree to which recent behavior is similar to the historical profile. Large values indicate abnormal behavior.

5.1 Measures

Like NIDES/STAT, we define 3 types of measures listed in Table 2.

BGP updates message arrival frequency (M1). This measure is one of activity intensity measures. It measures the inter-arrival time of BGP update messages sent by a router for a single prefix. We devise this measure to detect BGP update burst. BGP update burst most likely indicates abnormal operations. Moreover, the burst itself may impair the network because a huge number of update messages can occupy the overall resource of a BGP router, or even cause a router crash.

For this measure, the Q value corresponding to the current update message represents the number of update messages that have arrived in the recent past. In exponentially weighed sums scheme, whenever a new update arrives, the system will assign a Q value based on the following formula.

$$Q_n = 1 + Q_{n-1} * 2^{-r*\Delta t}$$

where r is the decay factor, Δt is the inter-arrival time between the current and the previous update.




Fig. 2. BGP Update Class Hierarchy

Number of AS paths (M2). This measure is another intensity measure. Due to link failure or router crash, BGP will suffer slow convergence problem. During convergence process, BGP router may receive a number of potential AS paths that are seldom seen in the past. Therefore, the number of AS path in that period may drastically increase. This measure is devised to monitor the variation of the number of AS paths. The Q value is calculated by the following formula

$$Q_n = N_{new_aspaths} + Q_{n-1} * 2^{-r * \Delta t}$$

where current Q is the number of new AS paths detected in the current audit record plus decayed previous Q .

BGP update type (M3). Similar to [14], we classify BGP update messages into 7 types in a hierarchical structure (Fig 2). At the top of the class hierarchy are two major classes: announcements and withdraws. Announcement is further classified into three sub-classes. Duplicate announcement indicates that the consecutive updates contain exactly the same information. (However, note that because we remove all the duplicate updates, this category does not exist in the experiments.) If the new route contains the same AS path as the current route, it is labeled as SPATH. Due to MED oscillation problem, we further distinguish SPATH by checking if the MED value is different. DPATH indicates that the current route is replaced by a different AS path. Because the length of AS path is a key factor in the BGP route selection process, we divide this sub-class into 3 more specific groups: same length AS path, longer AS path, Shorter AS path. The leaf nodes in this classification tree are the types of BGP update messages. Currently, we use these six types. If necessary, we can still sub-classify these types.

AS path occurrence frequency (M4). According to the observation that only a small number of different AS paths are announced, we define a categorical measure to capture the frequency distribution of AS paths occurrence. Each individual category within this measure is a different AS path. We calculate the frequency of each AS path occurrence. Since a new AS path will appear in the future, we utilize the "new path" category to denote the new path.

The Q computation for the categorical measure is:

$$Q_n = \sum_{m=1}^M [(g_{m,n} - f_m)^2 / V_m]$$

where

f_m = the relative frequency with the m^{th} AS path has occurred in the history.

$g_{m,n}$ = the relative frequency with which the m^{th} AS path has occurred in the recent past (which ends at the n^{th} received UPDATE message).

V_m = the approximate variance of the $g_{m,n}$

For detailed computation of these variables, please refer [13].

AS path difference (M5). In order to compare the current AS path with historical dominant AS path, we employ this measure. We use a *Simi*(*path1*, *path2*) function to calculate the difference between two AS paths. First, we define the AS path as a string in which each character is an AS number. Then we calculate the edit distance of two strings. In *Simi* function, *path1* is the current AS path, *path2* is the historical dominant AS path which is usually the most stable path. The edit distance of two paths denotes their difference. The larger the distance, the greater the difference.

Combination of five measures. The NIDES/STAT algorithm defines another variable S which is "normalizing" transformation of Q statistics so that the degree of abnormality for different measures can be added on a comparable basis. S has a half-normal distribution. Since each individual measure has a S value for each BGP update message, the anomaly detector can generate a single score value T^2 by the following formula:

$$T^2 = (S_1^2 + S_2^2 + \dots + S_n^2)/n$$

The details of transformation from Q to S can be found in [12].

Because the value of S ranges from 0 to 3.9, T^2 can range from 0 to 15.2 theoretically. In practice, we set the threshold of T^2 to be 2.5, since chances are very small for T^2 to have greater values based on our past experience.

5.2 Experiments for Statistics-Based Anomaly Detection

Experiments overview. Our experiments consist of two major parts, historical profile training and detecting process. Long term historical profile training is the process by which the anomaly detector learns the past behaviors of a subject. Detecting process examines the testing data by comparing current routing behaviors with the historical behaviors. If the deviation score is above the predefined threshold, a warning will be flagged. Otherwise, the data will be considered normal and incorporated into historical profile.

Experiments parameters. The decay factor has a significant impact on our detector. According to [1], most convergence time is about 3 minutes. Thus, in the case of inter-arrival time measure, the decay factor is set to be 1/300, which corresponds to the half-life of 300 seconds. Please note that convergence time is a function of the topology, MRAI timer, route flap damping, and routing policy. We cannot prove that this decay value is optimal. However, based on the distribution of the inter-arrival time for each prefixes we observed, most of inter-arrival time is less than 300 seconds or greater than 3000 seconds. We




Fig. 3. 166

Time	AS_PATH
01:51:37	3333 3356 1239 9405 4538
02:08:07	3333 3356 1239 9407 9407 4538
02:09:55	3333 286 209 1239 9407 9407 4538
02:10:22	3333 12859 13237 1299 701 1239 9407 9407 4538
02:11:17	Path Withdrawal

Fig. 4. Anomalous Update Sequence

choose 300 seconds as half-life value to capture the frequent route changes. For the categorical measure, we set the half-life decay to 20 BGP update messages, corresponding to an $r = 0.05$.

Experiment results. We test our statistical anomaly detector on the BGP UPDATE data to see whether it is able to effectively detect the BGP routing anomalies and whether it can help users analyze BGP routing dynamics. We conduct our test in two directions. On one hand, we perform the test on the BGP UPDATES data during SQL worm attack. Although SQL worm does not intent to attack BGP protocol, BGP has been impacted during worm attack. We test our detector to see if the detector can find out the anomalies in that period. On the other hand, in order to compare the results with those from signature detection, we perform the test on the same prefixes.

Table 3. S and T^2 values

Prefixes	S_{M1}	S_{M2}	S_{M3}	S_{M4}	S_{M5}	T^2
166.111.0.0/16	1.032996	2.809773	2.610673	1.220799	2.659919	4.868608
203.250.84.0/24	1.607696	1.967813	2.554678	0.212015	2.046884	3.235723
199.226.96.0/20	2.497571	1.907853	2.366445	0.773497	2.313003	4.285221

Detecting SQL worm. SQL worm attacked the Internet on Jan 25, 2003. Although SQL worm did not intent to attack the Internet routing architecture, a large increase of the number of BGP routing updates have been observed during that period. We apply the anomaly detection on the BGP updates and find that the warnings have been flagged for some prefixes from DoD, Korean and China, while the prefixes for popular destination and root server appear normal. Through the comparison of the prefixes from DoD, Korean and China, we can infer that their abnormal behaviors are similar in essence, while their BGP update sequences are different. Table 3 lists the S and T^2 values of the most abnormal update for each prefix. From the table, we observe some similarities in S distribution among these prefixes— all three assigned S values (S_{M2} , S_{M3} , and S_{M5}) are abnormally large (greater than 1.96, indicating that the probabil-

ity for the update message occurrence is less than 5%). This similarity leads us to further manually examine the BGP traffic of the three prefixes.

Figure 3 plots the T^2 values for the prefix 166.111/16 that is the address block for a university in China. The X-axis denotes the index of each update. All updates were recorded from Sept 1, 2002 to Jan 31, 2003. The highlighted updates (between the arrow “begin” and “end” in the figure) were recorded on Jan 25, 2003 when SQL worm attacked the whole Internet. Observing many UPDATES with large T^2 values (statistically significant deviation) on that day, we can easily infer that BGP routing has produced many highly abnormal behaviors. One of BGP update sequences was shown in Fig 4. The last BGP withdrawal message flags a warning. The corresponding S and T^2 values are listed in the first row of table 3. The large S_{M2} and S_{M5} value were due to the previous two updates. Large S_{M2} value indicates the arrival of new pathes (third and fourth AS_PATH). large S_{M5} value indicates that the pathes are significantly different from the dominant path. This is a classical case of slow convergence [1]. Possible root cause could be that huge amount of traffic generated by SQL worm congested the link between AS1239 and AS9407 and tore down the BGP session. This example demonstrates that our detector can detect BGP slow convergence effectively, and the statistical information learned from the detector can help analyze what causes a anomaly.

Comparison with signature detection. We apply statistics-based anomaly detection on the same prefixes examined by signature-based anomaly detection. Compared with signature detection, the number of anomalous incidents identified by statistics-based detection is much smaller.

For example, we examined BGP updates for the prefix 166.111/16 from AS3257 in the dataset of Jan. 2003. The detector did not flag warning during the SQL worm attack, because of the imperfection of our training dataset. However, it did flag a warning for one type B incident that is worthy of more investigation than the incidents caused by the worm. The abnormal in that type B incident is a brand new fail-over path, {3257, 3356, 12013, 3681, 20080, 11537, 9405, 4538}. This path has never appeared before and remains for only 500 seconds. In addition, the new path is very different from the primary path {3257,1239,9405,4538}. Moreover, the three transit ASes (3681,12013 and 20080) are ASes of two universities in Florida. It is highly abnormal that the networks of two American universities provide transit service for a university in China. Thus, we believe this incident is more interesting than other type B incidents because it may indicate some routing policy misconfiguration.

6 Discussion

For signature-based BGP detection, if the patterns of anomalies are well defined and persistently updated, this method should be very efficient in terms of false rate. However, in the BGP scenario, it is very hard to accurately define signatures. In this paper, the parameters ($K = 4, T = 240s$) are set empirically.

We may miss some anomalous incidents with only three close updates, or may incorrectly treat two consecutive updates as two different events.

Statistics-based BGP anomaly detection does not require knowledge about patterns of anomalies in advance. It can assign BGP updates with different deviation scores, providing an objective measure telling which incident is more abnormal and deserves more attention. In addition, statistics-based detector can provide information on the detected anomalies and help network operators investigate what may have triggered a warning. In our detection system, whenever an alarm is raised, detector provides both expected distributions and observed distributions (an example is shown in [12]). With this additional information, operators might be able to speculate what could have accounted for the statistically significant deviation.

Limitations also exist in statistics-based method. In our experiments, anomaly detector appears to have a relatively higher false rate compared to signature-based detector, because we do not have a clean training dataset in advance. The expected behaviors learned by the detector may have included problematic BGP UPDATE sequences.

Through comparison, we find that while both approaches are capable of identifying BGP routing anomalies to some extent, the list of detected anomalies is not exhaustive. Experiments demonstrate that combination of the two approaches can generate more comprehensive results.

At the current stage, we are unable to evaluate the identified anomalies, because evaluation is based on root cause analysis which is still an open question. The major barrier for root cause analysis is that we cannot acquire the necessary information from real operational network. Further, root cause analysis for BGP anomalies may need cooperations among ASes, because under some circumstances, identification of certain causes is extremely hard for an individual AS. Thus, the goal of our current work is not to provide accurate root causes analysis for speculative anomalies. In stead, we aim to devise an approach to identify possible anomalies, which is the first step towards solving the root cause analysis problem.

7 Conclusion and Future Work

In this paper, two approaches, signature-based detection and statistics-based detection, are proposed to search for anomalous BGP routing dynamics. The value of our work lies in the following aspects: First, we develop two systematic approaches to detect abnormal BGP UPDATE traffic. In current network management, they can help operators and researchers to filter out the trivial events and focus mainly on the most important BGP events. Second, through our experiments, we identify advantages and limitations of both methods. A feasible way to overcome the weakness of each is through combination of both. Third, these two detection approaches can be further used in monitoring and analyzing the real-time BGP traffic. In particular, statistics-based approach can quantitatively measure the “abnormality” of each BGP UPDATE.

The limitation of our work lies in the lack of information from real BGP run-time environment. At this stage, we are not able to thoroughly evaluate the identified anomalies. Root causes for most of the anomalies are still our conjectures. However, in our future work, we plan to attack this problem by building a large scale BGP testbed [15]. In this simulated BGP operational environment, we can generate various fully-controlled network failures and attacks. Applying the detection approaches to examine the simulated BGP traffic, we can provide a more extensive evaluation of the detectors' performance.

Acknowledgment. We would like to thank Chen-Nee Chuah for valuable comments. We also thank Yifei Zhu for proofreading.

References

1. C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed Internet Routing Convergence. In *Proceedings of ACM Sigcomm*, August 2000.
2. C. Labovitz, G. Malan, and F. Jahanian. Internet Routing Instability. In *Proceedings of ACM Sigcomm*, September 1997.
3. C. Labovitz, F. Jahanian, and G.R.Manlan. Origin of Internet Routing Stability. In *Proceedings of the IEEE INFOCOM*, June 1999.
4. A. Basu, C. Ong, A Rasala, F. Shepherd, and G. Wilfong. Route Oscillations in I-BGP with Route Reflection. In *Proceedings of ACM Sigcomm*, August 2002.
5. T. Griffin and G. Wilfong. Analysis of the MED Oscillation Problem in BGP. In *Proceedings of ICNP*, November 2002.
6. L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. Wu, and L. Zhang. Protecting BGP Routes to Top Level DNS Servers. In *Proceedings of the ICDCS 2003*, 2003.
7. S.T Teoh, K.L. Ma, and S.F. Wu. A Visual Exploration Process for the Analysis of Internet Routing Data. In *Proceedings of IEEE Visualization*, 2003.
8. The RIPE Routing Information Services. <http://www.ris.ripe.net>.
9. J. Rexford, J. Wang, Z. Xiao, and Y. Zhang. BGP Routing Stability of Popular Destinations. In *Proceeds of Internet Measurement Workshop*, November 2002.
10. X. Zhao, M. Lad, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang. Understanding BGP Behavior through a Study of DoD Prefixes. In *Proceedings of the IEEE DISCEX III*, 2003.
11. B. Halabi. *Internet Routing Architectures*. Cisco Press, second edition, 2001.
12. Ke Zhang, Amy Yen, Xiaoliang Zhao, Dan Massey, S.Felix Wu, and Lixia Zhang. On Detection of Anomalous Routing Dynamics in BGP. Technical Report CSE-2004-2, UCDAVIS, 2004.
13. H.S. Javitz and A. Valdes. The NIDES Statistical Components: Description and Justification. Technical report, SRI Network Information Center, March 1993.
14. L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. Wu, and L. Zhang. Observation and Analysis of BGP Behavior under Stress. In *Proceedings of the ACM IMW 2002*, October 2002.
15. R. Bajjcsy, T. Benzell, M. Bishop, B. Braden, C. Brodley, S. Fahmy, S. Floyd, W. Hardaker, G. Kesidis, K. Levitt, B. Lindell, P. Liu, D. Miller, R. Mundy, C. Neuman, R. Ostrenga, V. Paxson, P. Porras, C. Rosenberg, S. Sastry, D. Sterne, and S.F. Wu. Cyber Defense Technology: Experimental Research Network and Evaluation Methods. under submission.