

On Detection of Anomalous Routing Dynamics in BGP

Ke Zhang, Amy Yen, S.Felix Wu	Xiaoliang Zhao, Dan Massey	Lixia Zhang
Department of Computer Science	USC/ISI	University of California
University of California	Arlington, VA, U.S.A	Los Angeles, CA, U.S.A
Davis, CA, U.S.A		

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.

I. 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], [2], [3] 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 the intrusion detection area, we develop two approaches to detect BGP anomalous routing dynamics—the signature-based detection and the statistics-based anomaly detection. For signature-based detection, we devise a set of anomalous routing patterns to search for matching incidences 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 incidences and explain why they should be categorized as anomalies and specify possible root causes of some incidences.

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 the 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 "nomality" of each BGP UPDATE.

The rest of the paper is organized as follows. Section II introduces the concepts of signature-based detection and statistics-based detection, and briefly reviews related work. Section III describes the BGP UPDATE dataset that we have used in the experiments. Section IV and V presents signature-based detection and statistics based detection respectively. Section VI compares these two approaches, followed by conclusions in section VII.

II. RELATED WORK

Signature-based detection and anomaly-based detection are two major approaches in modern intrusion detection area. Signature based detection systems, such as the Snort IDS [4], report an attack when a set of symptoms corresponding to a predefined attack signature is observed. Anomaly-based intrusion detection flags as attacks any traffic that is unusual for that system. Statistical based anomaly detection systems are trained on some dataset; thereafter, any traffic that statistically differs from the training data is considered an attack.

As the de facto inter-domain routing protocol employed by edge routers, BGP is responsible for spreading the routing information throughout the Internet. BGP is a path vector protocol. BGP routing information carries a sequence of AS number, indicating the path a route has traversed. Routing information is exchanged between the two BGP routers through UPDATE messages. Incremental updates are sent as network information changes.

BGP routing behavior has received a lot of examination in the research literature. Labovitz et al. [5] showed that unstable and pathological routing behaviors dominated the Internet around 1996. Later, they presented potential explanations for these anomalies [6]. Other BGP routing problems, such as slow convergence [1], persistent MED oscillation [7], [8], have also been well examined. Rexford et al. defined the metric "BGP event" to measure the BGP routing instability and concluded that routes to the popular destinations were generally stable [9]. In order to generate the realistic BGP traffic for testing, Maennel et al. extensively studied BGP traffic and characterized the statistics properties of BGP traffic [10].

Also concerning about the abnormal BGP route changes, Wang et al. [11] proposed a path-filtering approach to validate the correct route changes for DNS prefixes. Teoh et al. develop an interactive visualization process to explore BGP data [12]. These works are complementary to our approaches described in this paper.

III. DATASET

The dataset we examine in the paper are BGP UPDATE messages collected by the collector RRC00 of the Routing Information Service of RIPE [13]. RRC00 has multi-hop BGP sessions with 9 peer ASes that are located at different countries. Because BGP updates reflect route changes of peering AS, we use BGP updates as our target. In addition, because the anomalies that we concerned in this paper are anomalous routing behaviors for a single prefix, we only select BGP updates for the selected prefixes.

Two major steps are needed to transform BGP raw updates into the input of our BGP anomaly detector. First, these messages are converted from the binary format to ASCII format through `route_btoa`, a software

provided by Multi-threaded Routing Toolkit (MRT). Second, the ASCII messages are parsed by a script one by one. This script filters out irrelevant messages based on three metrics. A message is relevant if it satisfies all three of the following conditions:

- The message must announce or withdraw the prefix currently being monitored.
- The message must come from the source AS which is our chosen observation point AS.
- The message must come from one particular router in the source AS. Messages with the same AS path information might come from different routers within the same source AS redundantly. Therefore, this metric filters out the duplicated effects.

In addition, due to the known implementation problem in some vendor's BGP router [6], duplicated updates may be announced for the same prefixes consequently. Duplicate update itself is a type of anomaly since it does not reflect route changes, However, in our experiments, we remove them from dataset because the origin of this anomaly is well-known.

Different prefixes may show different behaviors, and even for a simple prefix, the routing behaviors may be different from different observation point. In our experiments, we select 20 prefixes (4 prefixes for root DNS servers, 4 prefixes for gTLD servers, 4 prefixes for popular destinations, 4 prefixes from Department of Defense, 2 from Korean, 2 from China) from 9 peer ASes.

IV. SIGNATURE DETECTION

In this part, we describe a set of signatures that depict the anomalous routing dynamics.

A. *Patterns of Anomalous BGP Dynamics*

A route announced by a BGP router is generally the best route at that moment. Comparing the consecutive 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. We first briefly describe the route selection process [14].

Given a set of different BGP route announcements,

- 1) Accept the routes with the highest LOCAL_PREFER.
- 2) If the LOCAL_PREFER is the same, prefer the route with the shortest AS_PATH.
- 3) If the AS_PATH length is the same, prefer the route with lowest origin type, where a route originally learned from internal protocol (IGP) is preferable to a route learned from external protocol (EGP), which is preferable to a route injected into BGP via redistribution statically or dynamically. (INCOMPLETE).

- 4) If the origin type is the same, prefer the route with the lowest MED.
- 5) If the routes have the same MED, prefer the route learned from an external BGP (EBGP) advertisement over an internal BGP (IBGP) advertisements.
- 6) If all the preceding scenarios are identical, break ties by accepting the advertisement with the smallest intra-domain (IGP) cost to the NEXT_HOP router.
- 7) Break any remaining ties by accepting the routes announced by the router with lowest router ID.

Note that, since the BGP updates are collected through the EBGP sessions, we cannot acquire some 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 by comparing the AS_PATH length, origin type and MED value of each consecutive route announcement.

Based on the relative preference value of two consecutive routes, we define four terms.

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, we label the second route as FLAT.

WD: if the second announcement is 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 and the time interval between update messages is less than T and the average update rate $> \alpha$. In the experiments, we set $K = 4$, $T = 240s$, and $\alpha = 1/90$.

Then, given a BGP update burst, we map the updates into a {UP, DOWN, FLAT, WD} sequence. Following rules are used to define the signatures.

Type A: If the update burst ends with WD, then we believe it should be an incidence of slow convergence.

Type B: If the update burst has WD in the middle, it indicates transient failure followed by 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 transient failure (or congestion) followed by fast fail-over, or normal route changes.

TABLE I
SIGNATURES OF BGP UPDATE BURST

TYPE	Pattern	Examples	Indication
A	A sequence of updates ends with WD	<D,D,F,D,W> <U,F,F,D,W> <U,D,U,D,W> <D,W,D,U,W>	Slow convergence due to link/router failure
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
D	A sequence of updates with monotonic increasing or decreasing route preference	<U,U,U,U,U> <D,D,D,D,D> <U,F,F,U,F> <D,F,D,F,F>	Normal, but relative slow convergence process
E	A sequence of updates with two more route preference fluctuation	<U,D,U,D,U> <D,U,D,U,D> <D,U,U,D,U>	Routes flap
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

Type D: If the update burst does not include WD and route preference is monotonic increasing or decreasing, it indicates a normal, but relatively slow convergence process.

Type E: If the update burst consists of two more <UP, DOWN> or <DOWN, UP> sequence, the update burst should be anomalous. Whenever a higher preferable route is replaced by a lower preferable route, either the higher preferable route is withdrawn or some attributes change, such as LOCAL_PREFER or MED. In the short time window, two or more preference changes indicates that the route to destination is oscillating.

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. The only difference might be 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 and other information, we do not know whether or not the same length AS paths have different preference. If the difference lies in the content of AS_PATH, we

can only label the sequence as a speculative incidence. However, in the special case, if two or more same length AS paths oscillate in the BGP burst, then it should be anomalous, because that kind of burst indicates the route to a prefix is highly unstable. If the difference lies in community attributes, the burst might be anomalous due to the policy of the ISP. Our observation reveals that a major ISP (AS3549) sent five same routes with different community attributes in a very short period of time (6 seconds). Although the route is stable, if the downstream AS performs BGP route flap damping, this stable route will be suppressed. For the same reason, if the stable route frequently changes the ATOMIC_AGGREGAT or AGGREGATOR, it can also be suppressed.

B. Experiment Results for Signature Detection

We have performed the signature detection on 20 prefixes over the period from Feb. 2002 to Jan. 2003. Here we only show the detected anomalies for four representative prefixes. Other prefixes have similar results. We observe these prefixes from 9 peer ASes, however, we only show the results from 5 ASes where the behaviors are representative.

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. The five ASes are AS2914, AS3333, AS13129, AS3549, AS3257.

Figure 1, 2 plot the number of the each type of incidences at five ASes. From the figure, we can notice that AS2914 and AS13129 are more stable than AS3549 and AS3257 for these particular prefixes. For example, no anomaly for Yahoo prefix is observed from AS2914 while 28 anomalies have been noted from AS3257; for prefix 166.111/16, 8 anomalies are seen from AS13129 while 195 anomalies are observed from AS3549. In addition, Yahoo prefix is the most stable prefix, confirming the conclusion that popular prefixes are usually stable in the paper [9]. Although the DNS root server-A prefix is well engineered and maintained, it still shows some unstable incidences. It is probably due to the fact that bgp flap damping [15] is unused on the root server prefixes, some of frequent route changes are not suppressed.

We notice that the DoD prefixes has total 69 type F incidences 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 rate of AGGREGATOR substitution is very high, once per minute on average, which deserves more attention from the operators to figure out what happened in the DoD networks. In addition, as paper [16] pointed out, the local AGGREGATOR changes in theory should be restricted in the local area, but not be propagated

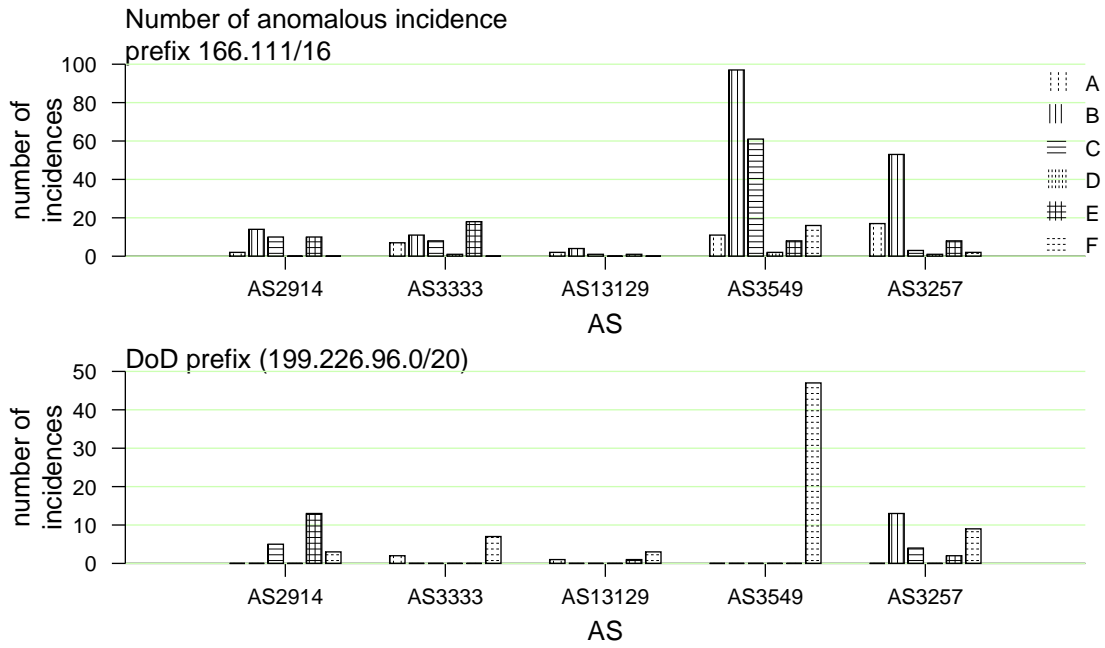


Fig. 1. anomalies detected by Signature Detector




Fig. 2. anomalies detected by Signature Detector

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 of 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 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 [17]) 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 incidences, 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 potential 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 this kind of incidence should be anomalous because the frequent route changes can degrade the packet forwarding.

V. STATISTICS-BASED ANOMALY DETECTION

We apply a statistics-based anomaly detection method, NIDES/STAT [18], which is introduced first by SRI. 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. This is achieved through a χ^2 -like test for comparing the similarity between the testing data and training data.

In this section, we will first introduce the algorithm and then elaborate on the measures specially designed for BGP anomaly detection.

A. Algorithm

Here are some notations for the NIDES/STAT algorithm: let the current system behavior be a random variable under sample space S . Event, E_1, E_2, \dots, E_k , represent a partition of S , where these k events are mutually exclusive and exhaustive. Let p_1, p_2, \dots, p_k be the expected probabilities of the occurrence corresponding to events E_1, E_2, \dots, E_k . To verify if the random variable actually has the distribution depicted by p_1, p_2, \dots, p_k , experiments are repeated N times independently, where N is a large number.

Let Y_i represents the real number of occurrence for event E_i , we have $\sum_{i=1}^k (Y_i) = N$. Let p'_i represents the empirical probability for event, i.e. $p'_i = Y_i/N$. We then test the hypothesis

$$H_0 : p'_i = p_i, i = 1, 2, \dots, k$$

and

$$H_1 : H_0 \text{ is not true}$$

Let

$$Q = \sum_{i=1}^k \frac{(Y_i - N \times P_i)^2}{N \times p_i}$$

If the independence is assumed between events $E_i, 1 \leq i \leq k$, it has been proven that, for a large N' , Q has an approximate χ^2 distribution with $(k - 1)$ degrees of freedom.

Let q be an instance of Q . If $\Pr(Q > q) < \alpha$ (or $q < \chi^2_{\alpha}(k - 1)$) where α is the desired significance level of the test, the hypothesis is rejected. In the context of our application, it means that the short-term profile is statistically different from its long-term profile which allows us to draw a conclusion that an anomalous behavior has occurred.

In practice, the assumption of independence may not be true. Furthermore, there may be insufficient observations for some bins in the data stream on which Q is based. Therefore, Q may not have a χ^2 distribution. SRI's NIDES/STAT proposed a way (exponentially weighted sums) to track the values of Q in order to establish an empirical probability distribution for Q . This distribution, along with distribution of the system's expected behavior, is saved in a long-term profile, which is updated with a predefined frequency.

B. Measures

NIDES/STAT defines four classes of measures: (1) activity intensity measure, which can detect bursts of activity. (2) categorical measure, whose values are by nature categorical. (3) counting measures, whose values are numerical. (4) audit record distribution. In our statistics-based BGP anomaly detector, five measures are defined.

1) Introduction of five measures:

TABLE II
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

a) 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 the huge number of update messages can occupy the overall resource of a BGP router and freeze a 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 fomula.

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

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

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

c) BGP Update Type (M3): Similar to [19], we classify BGP update messages into 7 types. At the top of the class hierarchy are two major classes: announcements and withdraws. An announcement contains the sender's BGP route to an address prefix, while a withdrawal indicates that the sender wants to remove a previously announced route.

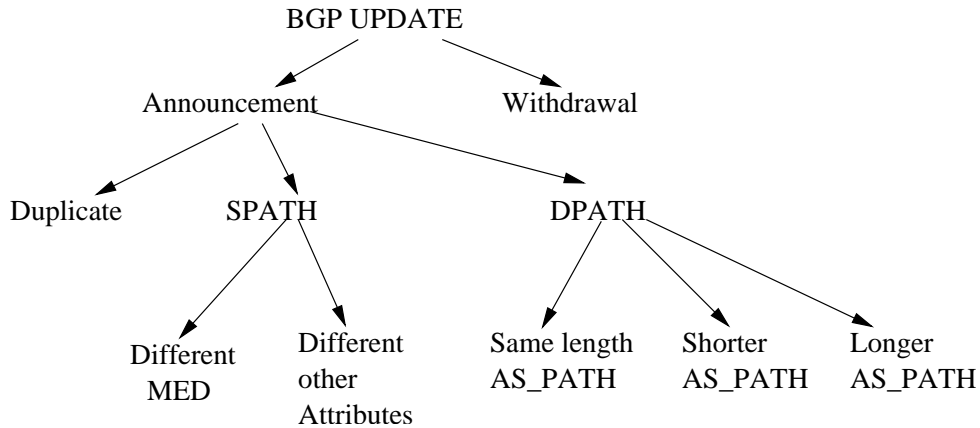


Fig. 3. BGP Update Class Hierarchy

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-classes into 3 more specific areas: 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.

d) 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. Network operators may predefine the probability of this category. Because BGP has no mechanism to ensure the authenticity of AS_PATH announced by an AS, the attacker can potentially announce arbitrary AS_PATH to disturb the whole Internet. This measure is designed to detect this kind of attack.

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 the detailed computation of these variables, please refer [18].

e) AS path difference (M5): In order to compare the current AS path with the common used 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. Edit distance is the smallest number of insertions, deletions, and substitutions required to change one string to another. 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 more the difference.

2) 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:

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

Large values of T^2 are indicative of abnormal behavior, and values close to zero are indicative of normal behavior (i.e. the behavior is consistent with expected behavior). Thus, T^2 itself is a summary judgment of the abnormality of BGP update messages. The network operators should pay more attention to the larger T^2 values.

Because Q computation is different for intensity measures and non-intensity measures, the transformation of Q to S is slightly different according to the type of measures.

For intensity measures, the S derivation from the tail probability of Q distribution is calculated through the following ways:

- 1) Let P_m denotes the relative frequency with which Q belongs to the m^{th} interval. In our experiment, there are 32 values for P_m , with $0 \leq m \leq 31$.
- 2) For the m^{th} interval, let $TPROB_m$ denote the sum of P_m and all other P values that are smaller than or equal to P_m in magnitude.
- 3) For the m interval, let s_m be the value such that the probability that a normally distributed variable

with mean 0 and variance 1 is larger than s in absolute value equals $TPROB_m$.

On the other hand, transformation of Q to S for the non-intensity measures is as follows:

- 1) We let $TPROB_m = P_m + P_{m-1} + \dots + P_{31}$.
- 2) S is derived by the following formula from $TPROB_m$

$$s_m = \phi^{-1}(1 - (TPROB_m/2))$$

where ϕ is the cumulative normal distribution function of an $N(0, 1)$ variable.

Because 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, because the chance that T^2 has a greater value is very small based on our past experience.

C. Experiments for Statistics-based Anomaly Detection

1) *Experiments overview:* Our experiments consist of two major parts, historical profile training and detecting process.

- 1) Long term historical profile training is the process by which the anomaly detector learns the past behaviors for a subject. (Please note here we only declare they are past behavior not normal behavior)
- 2) 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.

In the rest of this subsection, we will discuss the critical parameters for our anomaly detector and show the experiment results.

2) *Experiments parameters:* Certain parameters have significant impacts on the performance of the statistics-based BGP anomaly detector. As the preliminary stage of our experiment, we choose the parameters based on our past experience. In practice, choice of these parameters needs comprehensive consideration.

First, the decay factor has a significant impact on our detector. According to [4], most convergence time is about 3 minutes. Thus, in the case of inter-arrival time measure, the decay factor is set to be 0.00333333, which corresponds to the half-life of 300 seconds or 5 minutes. 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.

TABLE III
SQL WORM ATTACK TESTING

Prefixes		Observation point		
		AS3333	AS2914	AS7018
Popular	Yahoo	Normal	Normal	Normal
Prefixes	Realnetworks	Normal	Normal	Normal
DNS Root server A 198.41.0.0/24		Normal	Normal	Normal
DoD prefix 199.226.96.0/20		Warning	Warning	Warning
Korean's prefix 203.250.84.0/24		Warning	Warning	Warning
China's prefix 166.111.0.0/16		Warning	Warning	Warning

TABLE IV
 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

We 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 $r = 0.05$.

3) *Experiment results:* Anomaly detection is capable of finding the known and unknown anomalous behaviors. 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 to analyze the BGP routing dynamics. We conduct our test in two directions.

- 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. The results show that there are indeed anomalous routing behaviors during that period, and the information provided by the detector gives confirmatory supports to the speculation that the major anomaly is slow convergence correlated to SQL worm.
- In order to compare the results with those from signature detection, we perform the test on the same prefixes. We notice that statistics-based anomaly detector identified fewer incidences than signature-

TABLE V
EXAMPLE OF AN ANOMALY

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

TABLE VI
EXAMPLE OF AN ANOMALY

Time	AS_PATH
16:39:13	2914 701 10913 19836
16:39:39	2914 701 11840 19836 19836 19836 19836 19836 19836 19836 19836 19836 19836
16:40:34	2914 3549 10913 10913 10913 19836
16:41:01	2914 701 10913 19836

based detector. Furthermore, most of the incidences identified by statistics-based anomaly detector belong to the subset of anomalies identified by signature. However, those incidences appear to be more deviant indicating that statistical anomaly detector can select the most anomalous incidences.

In the following part, we will show the experiment results for these tests and conduct an analysis.

a) 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 log data to see if the detector can raise alarm. Part of the results is listed in the table III. From the table, we can notice 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.

For a more careful analysis, figure 4 plots the T^2 value for the prefix 166.111.0.0/16 that is the address block for a university in China. The X-axis denotes the index of each update. All updates were recorded from 1st Sept. 2002 to 31st Jan. 2003. The highlighted updates (between the arrow “begin” and “end” in the figure) were recorded on 25th Jan 2003 when SQL worm attacked the whole Internet. Observing many UPDATES with large T^2 value (statistically significant deviation) at that day, we can easily infer that

BGP routing has produced many highly abnormal behaviours. One of the incidences was shown in table V. The last BGP withdrawal message generates the biggest T^2 . The corresponding S values are listed in the first row of table IV. S_{M3} is large because the probability for the router to withdraw the route to the prefix is small in long-term training data. When the withdrawal arrived, the detector thought it was not expected then raised alert. The large S_{M2} and S_{M5} value were due to the previous two updates, because both updates gave the new paths that AS3333 never propagated before. Large S_{M2} value indicates the arrival of a new path, and large S_{M5} value indicates that the path is significantly different from the dominant path.

In fact, we find that this is a case of classical slow convergence incidence [1]. We speculate that since SQL worms generated huge amount of traffic and congested some link (we believe in this example, the link was between AS1239 and AS9407), the BGP session between that two AS was down, so AS1239 announced a withdrawal. Due to propagation delay, AS3333 had to exploit some backup routes that actually had already become invalid but AS3333 did not know at that moment. AS3333 suffered the slow convergence and at last it withdrew the route to that prefix. 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 the anomalies.

In addition, through the comparison of the prefixes from DoD, Korean and China, we can infer that their abnormal behaviors are similar in essential, while their BGP update sequences are different. Table IV lists the S and T^2 values of the most abnormal update for each prefix. From the table, we observe some similarity in S distribution among these prefixes. For the anomalous behaviors in the three prefixes, three assigned S values (S_{M2} , S_{M3} , and S_{M5}) all are abnormal large (greater than 1.96, indicating that the probability for the update message occurrence is less than 5%). This similarity leads us to further manually examine the BGP traffic of the three prefixes. The results confirm that all of them suffered slow convergence during worm attack.

b) 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 incidences identified by statistics-based detection is much smaller. That is because the long term historical profile used to train the detector includes many incidences identified by the signatures as “abnormal”. In the training phase, the anomaly detector learns these incidences as the normal ones. Consequently, in the testing phase, many similar incidences are neglected. We will show the an example to explain this case later in the paper.

First, we show the results of detection on the DNS root server A prefix (198.41.0.0/24). The dataset

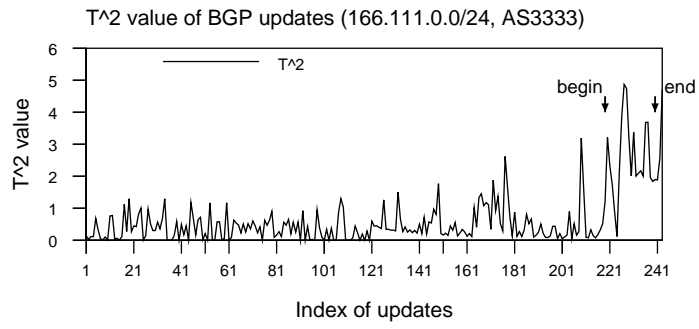


Fig. 4.




Fig. 5.

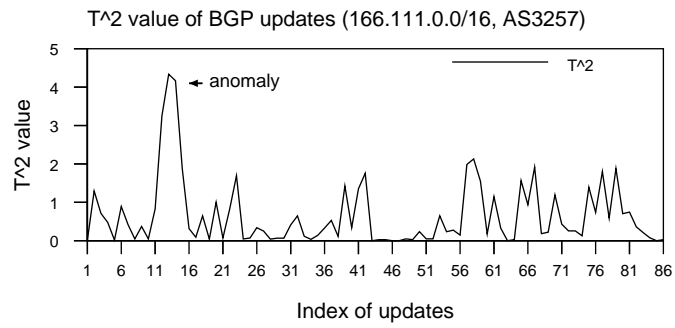


Fig. 6.

is all BGP updates from AS2914 for that prefixes in 2002. Although prefixes for DNS root servers are well maintained, we still find out one “relative anomaly”. For this case, the result of statistics anomaly detection is the same as that of signature detection. Signature detector only detects this incidence as Type B anomaly. We will show this anomaly and explain why the statistics based detector flags it as anomalous routing later in the paper.

Figure 5 plots the T^2 value of each update in the year 2002. Although the dataset appear normal (only 66 updates in a whole year), we are still able to find one anomalous incidence that consists of a sequence of updates with relatively large T^2 value. This anomaly is showed in table VI. In this incidence, four close BGP updates and three AS paths are announced. The second path is a brand new path, and the third path is a transient path. This event indicates that BGP route exploits some very strange route that router usually discards and then changes back to preferred previous route. One probable explanation for this is that AS 11840 provides a backup route for AS19836, since AS11840 is under same administrative domain as AS19836. AS2914 cannot observe this backup link in normal situation, unless some special events happened, such as link failure or severe congestion. Although the root cause may be the normal BGP operation, this special event should still deserve attention due to its rarity.

We also examined BGP updates for the prefix 166.111/16 from AS3257 in the dataset of Jan. 2003 (figure 6). It is interesting to notice that the detector did not flag warning during the SQL worm attack. However, the signature detector identified two Type A and two Type B incidences. Why statistics anomaly detector failed to detect the slow convergences and transient failures? It is because in the training dataset, the updates in 2002, there are 16 Type A and 47 Type B incidences. The statistical detector learns these incidences as normal incidences, and records the statistics properties into long-term profile. In the testing phase, when the detector scans these similar incidences, it cannot find out the significant difference so that no warning is flagged.

Although the statistics-based detector fails to detect anomalies caused by the worm, it did flag a warning for one type B incidence that is worth of more investigation than the incidences caused by the worm. (highlighted in figure 4). The abnormal in that type B incidence is the brand new fail-over path, {3257, 3356, 12013, 3681, 20080, 11537, 9405, 4538}. This path is never seen before and only remains for 500 seconds. Moreover, the new path is very different from the primary path {3257,1239,9405,4538}. We believe this incidence is more interesting than other type B incidences because this incidence involves an extremely strange path which might be a symbol of potential attack.

VI. DISCUSSION

In this section, we discuss the advantages and limitations of both signature-based and statistics-based BGP anomaly detection techniques.

Signature-based BGP anomaly detection can be easily performed and the processing overhead is trivial. If the signatures of anomalies are well defined and persistently updated, this method should have a lower false rate compared to other detection techniques (since, in theory, it knows what it is looking for). Matching examined data with predefined signatures, people can roughly know what might have happened in the networks.

However, in the BGP scenario, it is very difficult to accurately define signatures. In this paper, We set the parameters to define some simple signatures based on our experience. Admittedly, these parameters may not be optimal. For example, we set $K = 4$ and $T = 240s$ in the definition of BGP update burst to space out the consecutive updates and to locate the anomalous events. Since K and T are hard-coded, we may miss some anomalous incidences with only three close updates, or may incorrectly treat two consecutive updates as two different events. In order to design precise and complete signatures, a lot of inputs from network operators and researchers are indispensable.

Statistics-based BGP anomaly detection does not need to be equipped with knowledge about patterns of anomalies in advance. It simply compares historical BGP routing behaviors with current behaviors and identifies significant differences. Further, the statistics-based detector can assign BGP updates with different deviation scores, providing an objective measure that can tell which incidence 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. With this additional information, operators might be able to speculate what might have accounted for the statistically significant deviation. Table VII presents comparison of observed and expected distributions of 19 AS paths occurrence frequency(captured by measure M4). Among the 19 paths, only AS path 5 occurs much more frequently than expected, indicating that this path probably has triggered the warning.

Limitation also exists 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

TABLE VII
EXPECTED AND OBSERVED DISTRIBUTION

ID	AS_PATH	Expected Dist.	Observed Dist.
1	3333 5378 6660 3561 6245	0.00431348	0.00548961
2	3333 1103 6453 1 6245	0.00680711	0.00866316
3	3333 1103 6453 701 6245	0.00261389	0.00332661
4	3333 5378 6660 4544 6245	0.000874996	0.00111358
5	3333 9057 3356 701 6245	0.37579	0.478254
6	3333 9057 702 701 6245	0.00201142	0.00255986
7	3333 286 209 6082 6245	0.00795883	0.0101289
...			
19	3333 9057 1755 3561 6245	0.0261248	0.0332481

anomalies to some extent, the list of detected anomalies is not exhaustive. Signature-based detection can detect typical anomalous incidences but miss the incidences that do not match the predefined fixed patterns. These missed incidences could be detected by the statistics-based detection as long as their observed behaviours are statistically different from their expected behaviors. Conversely, while statistics-based detection tends to miss some possible anomalies which are statistically consistent with training data, signature-based detection can effectively identify them. Experiments demonstrate that combination of the two approaches generates the results with lower false rate.

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 the real operational network. Further, root cause analysis for BGP anomalies may need cooperations among ASes, because under some circumstances, identification of certain causes is almost a mission impossible for an individual AS. However, given the current selfish routing environment, this cooperation is hard to archive. 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. We have proposed a future project targeting at this problem, in which we plan to build a large-scale BGP testbed [20]. Whithin this testbed, we are capable to fully control the networks, and even simulate the network failures and attacks. At that stage, since we can manipulate the root causes, we will be in a better place to evaluate the detector's performance .

VII. CONCLUSION AND FUTURE WORK

As the de facto inter-domain routing protocol, BGP operation deserves extensive monitoring and examination. However, in the run-time operational network, operators are not able to analyze the whole BGP dataset due to its large quantity. The most reasonable way is to mainly focus on a small set of valuable data. With this goal in mind, we propose two approaches, signature-based detection and statistics-based detection, to search for anomalous BGP routing dynamics. We apply these approaches to examine the real BGP UPDATE data collected by RIPE-NCC. The results show that these approaches can effectively and efficiently identify the anomalous BGP routing behaviors.

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 [20]. 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.

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, R. Wattenhofer, S. Venkatachary, and A. Ahuja. The Impact of Internet Policy and Topology on Delayed Routing Convergence. In *Proceedings of the IEEE INFOCOM*, April 2001.
- [3] T. Griffin and B. Premore. An Experimental Analysis of BGP Convergence Time. In *Proceedings of ICNP*, November 2001.
- [4] Snort Project: the Open Source Intrusion Detection System. <http://www.snort.org>.
- [5] C. Labovitz, G. Malan, and F. Jahanian. Internet Routing Instability. In *Proceedings of ACM Sigcomm*, September 1997.
- [6] C. Labovitz, F. Jahanian, and G.R.Manlan. Origin of Internet Routing Stability. In *Proceedings of the IEEE INFOCOM*, June 1999.

- [7] 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.
- [8] T. Griffin and G. Wilfong. Analysis of the MED Oscillation Problem in BGP. In *Proceedings of ICNP*, November 2002.
- [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] Olaf Maennel and Anja Feldmann. Realistic BGP Traffic for Test Labs. In *Proceedings of the ACM SIGCOMM '02*, August 2002.
- [11] 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.
- [12] 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.
- [13] The RIPE Routing Information Services. <http://www.ris.ripe.net>.
- [14] B. Halabi. *Internet Routing Architectures*. Cisco Press, second edition, 2001.
- [15] R. Chandra C. Villamizar and R. Govindan. BGP Route Damping. RFC 2439, SRI Network Information Center, May 1998.
- [16] 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.
- [17] Z. Mao, R. Govindan, G. Varghese, and R. Katz. Route Flap Damping Exacerbates Internet Routing Convergence. In *Proceedings of ACM Sigcomm*, August 2002.
- [18] H.S. Javitz and A. Valdes. The NIDES Statistical Components: Description and Justification. Technical report, SRI Network Information Center, March 1993.
- [19] 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.
- [20] R. Bazjcsy, T. Benzel, 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.