On Reverse Engineering the Management Actions from Observed BGP Data

Shih Ming Tseng, S. Felix Wu
University of California, Davis
Email: {smtseng,sfwu}@ucdavis.edu

Xiaoliang Zhao
Verizon Communications Inc.
Email: xiaoliang.zhao@verizonbusiness.com

Ke Zhang
Cisco Systems Inc.
Email: kezhan@cisco.com

***Abstract*—While most of the research work on BGP has focused on detecting and characterizing large-scale routing anomalies from the perspective of network operations and management, it is important to monitor the management actions taken by the network operators in response to global BGP network failures. A fundamental question to answer is the following: by utilizing only public BGP observation data under today's Internet environment, can we reverse engineer the management actions taken by specific autonomous systems? In this paper, we propose a formal framework to describe and analyze MOAS [1] events and possible management actions. We use BGP data and a two-step learning approach to evaluate each possible action then determine the most likely one. Through this process, we discovered that early actions were taken by multiple ASes before the faulty originator corrected its mistake. Furthermore, the results show that only a handful of ASes took such early corrective action, but the effect is disproportional: a significant portion, more than 90%, of affected prefixes were routed back to their correct routing path.**

I. Introduction

As of February 2008, our Internet consists of 27500 autonomous systems (ASes), together, announcing more than 250,000 IP prefixes [2] via the standard inter-domain routing protocol, BGP (Border Gateway Protocol) [3]. Mainly due to misconfiguration, global BGP failures such as AS7007's false deaggregation of 65000+ network prefixes in 1997 severely disabled the operations of the Internet [4]. Similarly, in both 2001 and 2004, we observed that thousands of false Multiple Origin AS (MOAS) announcements directed a significant portion of Internet traffic toward incorrect destinations [5]. On Feb 24, 2008, one of world's largest online video distributor, YouTube, experienced hours-long service disruption because AS17557 (Pakistan Telecom) hijacked YouTube's prefix [6].

When a prefix P appears to originate from more than one AS, i.e., two or more ASes claim ownership of the same prefix, an instance of Multiple Origin AS (MOAS) occurs [1]. More precisely, for any two AS paths to reach prefix P. e.g. "AS_{x1} AS_{x2} ... AS_i" and "AS_{y1} AS_{y2} ... AS_j". A MOAS occurs if $AS_i \neq AS_j$.

Most of the research work on BGP has focused on detecting and characterizing these large-scale MOAS events, without identifying the exact root causes of the anomalies. From the perspective of network operations and management, it is also very interesting to monitor the management actions taken by the network operators in response to global BGP failures. Sometimes, as we will show later in the paper, even though the operators claim to fix the problems during a specific time frame, the public Internet data, if carefully analyzed, supports a different conclusion. A fundamental question to answer is this: under a particular event, has ISP_X (or AS_X) performed the "right" action to correct the problem? However, without knowing what action was taken, we can not verify the correctness of the action.

In this paper, given a set of MOAS events, we examine publicly available BGP traces to determine the management actions taken by different network operators. Under the context of automated network management, especially in an inter-domain environment, understanding which management actions being launched by other ASes is critical to possibly avoid unexpected interferences among different administrative domains. The other possible application of our approach is, for an AS, to monitor and verify the expected effectiveness of a particular management action against the target network problem.

According to [1], [7], one plausible hypothesis regarding a large-scale BGP MOAS problem is that the MOAS conflicts were withdrawn eventually by the faulty AS itself. However, is that hypothesis consistent with the network-level BGP forensics we have? While there can be an infinite number of hypotheses regarding what happened, surprisingly many plausible explanations or hypothesis about BGP global failures can be shown to be inconsistent with a significant percentage of BGP update messages. For example, on April 6th of 2001 at 5:27pm (UTC), AS3549 (GBLX Global Crossing Ltd.) received a false MOAS for prefix 140.113.0.0/16 (National Chao-Tung University, Taiwan) from AS15412 (Flag Telecom Global Internet AS). By examining the daily routing table changes, in [8], it is shown that AS15412 did not fix the MOAS until 2 days later. However, the BGP per-update data [5] observed at AS3549 clearly indicates that AS3549 removed the MOAS 15 minutes (i.e., at 5:42 pm) after the event itself. Similarly, data from other BGP peers such as AS2914, AS3257, and AS13129 shows that some MOAS conflicts did not last longer than 8 hours.

If the original hypothesis (i.e., AS15412 itself withdrew the conflicting MOAS event) is true, then most if not all of the Internet ASes should observe the correction for the prefix at roughly the same time. Then, the question becomes "how were the MOAS conflicts caused by AS15412 fixed?" One better (and more consistent) hypothesis is called **"routing policy changes,"** where a number of individual ASes intentionally

dropped certain BGP updates from AS15412 and MOAS conflicts were resolved. Then, the problem becomes how many and exactly which ASes were involved in **"routing policy changes."**

We propose a language to describe these changes formally. Once we have a formal way to describe MOAS events, analyzing and learning from BGP traces can be accomplished using the principle of computational learning theory. Because BGP traces are finite, learning theory behaves within probabilistic bounds. The result returned is correct with high probability. We develop a two-step approach to validate each possible action then determine the most likely one. First, we learn by analyzing BGP traces from a single data collector. Then, we correlate the results by resolving conflicts among all data collectors.

We discovered that some ASes took early actions to block "bad"[1] MOAS routes and fix MOAS prior to the faulty originator correcting its mistake. [6] also shows that some service providers took early actions to fix YouTube services. Furthermore, our results show only a modicum of ASes were able to block more than 90% of the "bad" MOAS routes.

The rest of the paper is organized as follows. Section II presents our scheme for discovering routing policy changes and we define a formal notation to describe it. Section III describes our two-step approach in great detail. In Section IV we report our findings derived from real BGP traffic traces obtained during three MOAS storms. Section V concludes the paper.

II. Detection of Routing Policy Changes

In this section, we explain how we discover possible routing policy changes and propose a language to describe the changes formally.

A. The discovery of routing policy changes

Our objective is to analyze public BGP data such as RIPE [9] or Oregon Route Views [10] and to determine the possible management actions taken by some ASes during particular MOAS events. For example, in Table I, we listed two BGP update messages observed by AS3549 for the prefix 140.113.0.0/16. The first message indicates that at 17:27 the origin AS for this prefix was changed from AS9916 to AS15412, while the second message states the opposite – the origin AS changed back to AS9916. Given just these two updates, we can have at least the following possible explanations:

AS15412 withdrew (Case A) the faulty originator detected and corrected itself and the data collector saw the route was changed back. This was the explanation given by paper [8].

AS3561 blocked(Case B) the immediate neighbor of the faulty AS blocked the faulty route. AS3561 detected and blocked the announcement from AS15412. Therefore, the data collector AS3549 lost the route to AS15412 and switched back to the previous route. Please note that this hypothesis was consistent with a posted article [11] on NANOG from James A. Farrar of AS3561 around 17:37 pm. on April 6th.

[1]Illegitimate BGP route

Assuming we did not see Farrar's message, can we tell, by only examining the BGP updates, which hypothesis or explanation is more likely to reflect the truth? Surprisingly, it turns out that neither Case A nor Case B matches the data in Table II, which contains additional data from another data collector - AS9177.

With two extra BGP updates observed by AS9177 for the same prefix, we can clearly tell that AS15412 did not withdraw the ownership of prefix 140.113.0.0/16 on April 6, at 17:42 pm. Moreover, until April 10, AS9177 (as well as AS8210) still believed that AS15412 provided the best route to reach prefix 140.113.0.0/16. Hence, "Case A" is clearly false.

On the other hand, if AS3561 had completely withdrawn the bogus announcement from AS15412 on April 6 around 17:37 (according to the claim by Farrar) , then both AS8210 and AS9177 would have been able to use the correct route on April 6th right after the fix claimed by Farrar. However, there is strong evidence in Table II line 5 which clearly points out that AS3561 used the bad route originated by AS15412 until April 10. Therefore, "Case B" is also not true.

As a result, the big question is: what actually happened during the MOAS of April 6, 2001? With merely four BGP update messages, we quickly disqualified two previously well accepted explanations of operations(management actions) for the MOAS instance of April 2001.

All we know from BGP traces is that, from some data collector's point of view, the announcement from AS15412 was gone. This could be an implicit withdraw(AS paths and/or other attributes are changed by new announcement messages) or explicit withdraw(prefixes are withdrawn by withdrawal messages). Therefore, we discuss all of the possible scenarios as follows:

One possible explanation is link failure. On April 6th, the link between AS3561 and AS3549 failed such that AS3549 changed the best route (Table II line 3). However, data show that there was still other BGP traffic along the link "AS3561-AS3549." Thus, link failure is clearly not true.

Another possibility is that some ASes setup an in-bound and out-bound route filtering policy such that no "bad" route announced by AS15412 passed through. Please note that each AS determines what is "bad." Based on Table II, a reasonable assumption is that AS3549 employed such route filters which made the best route change as shown in line 3.

In this paper, we focus our study on examining **Routing Policy Changes**, *i.e.,* a small number of ASes taking filtering actions on all BGP update messages originated from a particular AS. For example, to be more specific, during the April 2001 MOAS case, we are interested in finding the small set of ASes filtering all BGP update messages from AS15412, and of course, how well a particular filtering hypothesis matches the relevant BGP data.

Line	Time	AS path changes	Remarks
1	04/06 17:27	3549 7018 1659 9916 ⇒ 3549 3561 15412	Change to AS15412
2	04/06 17:42	3549 3561 15412 ⇒ 3549 7018 1659 9916	Change back to AS9916

TABLE I
TWO BGP UPDATES FOR PREFIX 140.113.0.0/16 ON APRIL 6TH OF 2001

Line	Time	OP	AS Path Changes	Remarks
1	04/06 17:27	AS3549	3549 7018 1659 9916 ⇒ 3549 3561 15412	MOAS first seen (CSM)
2	04/06 17:28	AS9177	9177 8210 1833 7018 1659 9916 ⇒ 9177 8210 1755 3561 3561 15412	still in MOAS
3	04/06 17:42	AS3549	3549 3561 15412 ⇒ 3549 7018 1659 9916	Fixed (but still in MOAS)
4	...	...	...	Nothing happened
5	04/10 13:11	AS9177	9177 8210 1755 3561 3561[2]15412 ⇒ 9177 8210 1833 7018 1659 9916	MOAS fixed (CMS [3])

TABLE II
FOUR BGP UPDATES FOR PREFIX 140.113.0.0/16 ON APRIL 2001

```
1.  L ::= <MOAS>*
2.  <MOAS>::= prefix @ <UTCTIME>,<PATHCHANGES>
3.  <UTCTIME>::= second | <STARTING> to <ENDING>
4.  <STARTING>::= second
5.  <ENDING>::= second
6.  <PATHCHANGES>::= <ONECHANGE><PATHCHANGES>*
7.  <ONECHANGE>::= <OLDPATH> → <NEWPATH> <EXPLANATIONS>
8.  <OLDPATH>::= <ASPATH>
9.  <NEWPATH>::= <ASPATH>
10. <ASPATH>::= nil | asnumber | <ASPATH>*
11. <EXPLANATIONS>::=<ONE_EXPLANATION><EXPLANATIONS>*
12. < ONE_EXPLANATION >::=<WHO><ACTION><TARGET>
13. <WHO>::= asnumber
14. <ACTION>::= announce | withdraw
15. <TARGET>::= asnumber
```

Fig. 1. Grammar of Language L

B. Formal Notation of Routing Policy Changes

A formal notation is essential since we need to describe raw data as well as produce corresponding results. In this section we propose language L to describe MOAS events and corresponding BGP operations. Figure 1 shows L in BNF like grammar. L can be empty if no MOAS happens. No matter what human operators do, all policies result in only two types of BGP updates (A and W). Hence, without the knowledge of actual operations, BGP traces clearly indicate the outcomes – prefixes were announced or withdrawn. Therefore, "ACTION" is an abstraction of either an announce or withdraw.

III. METHODOLOGY

In this section, we describe how to derive possible policy changes for a certain MOAS. We use the divide-and-conquer approach to solve one prefix at a time. We start by looking for a single AS path change and solve the problem by using the T-Shape analysis in Section III-A. Since there were many different data collectors, i.e. observation points (OP), and each OP may have a different point of view, we incorporate all the data from different OPs and correlate the information by resolving conflict between each data collector in Section III-B.

A. T-Shape Analysis for Single AS Path Change

When one data collector OP records an AS path change from old path $\{OP\ AS_{x1}\ AS_{x2}\ ...\ AS_{xm}\ \}$ to new path $\{OP\ AS_{y1}\ AS_{y2}\ ...\ AS_{yn}\ \}$, it is clear to see these two paths must have one or more common nodes. When having only one common node, the node is the data collector OP. When having two or more common nodes, the node closest to AS_{xm} is labeled as AS_{xk} where AS_{xk}=AS_{yj}(common node).

In the case where the old path was withdrawn, from data collector OP's point of view, any node between AS_{xk} and AS_{xm} (including AS_{xk} and AS_{xm}) is a possible candidate who withdrew the prefix. Since every AS within this segment is able to withdraw the BGP announcement from AS_{xm}, we need to evaluate and examine each of them. In other words, we consider "AS_Z withdraw AS_{xm}" where Z = {xk....xm}.

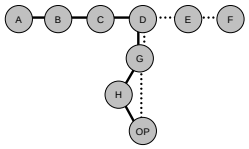

Fig. 2. T-Shape Analysis

When a new and better AS path emerges, it is possible that AS_{yn} announces prefix ownership or the nodes between AS_{yj} and AS_{yn} (including AS_{yj} and AS_{yn}) changed the current best route due to local preference/policy changes or reuse of previously damped routes. Regardless of the cause, we consider " AS_Z announce AS_{yn}" where Z = {yj....yn}.

To illustrate the process, we show an AS path change "OP G D E F ->OP H G D C B A" in Figure 2 which represents an AS path change recorded by data collector OP. The old path is shown with a dotted line and the new path is shown with a solid line.

The common nodes on both paths are D, G and OP. Among these three common nodes, D is the closest to F. Thus, in "withdraw" cases, nodes D, E and F are possible nodes to withdraw/block BGP updates from F. Similarly, in "announce" cases, nodes D, C, B and A are the candidates who announce, reuse/unblock the route from A.

B. Explanation Correlation

We describe how to find all of the possibilities for any MOAS observed from single point of view. However, different observation points may have different opinions. Sometimes, these opinions conflict with each other.

To resolve the conflicts from different observation points, we propose the following algorithm shown in Figure 3. There are two steps : (1) learn the facts, and (2) use the facts to resolve conflicts.

The main idea is to treat observations as facts and use these facts to remove any conflicting possibilities derived from section III-A. Since other observation points provide the current best path, we are able to clearly derive which AS "announce" which originator and we place that data point in "facts." Also, P denotes all possible explanations derived from section III-A. OP denotes other observation points and the

[2]Prepend, a common BGP technique
[3]Change from Multi-Home to Single-Home (CMS)

```
1.  while ( OP=pop( @All_OPs) )
2.          $aspath = readpath (OP)
3.          $originator = tail ($aspath)
4.          @ASes = split (' ', $aspath)
5.          foreach $AS  (@ASes) {
6.                  push ( @facts, "$AS announce $originator") }
7.  endwhile

8.  while ( explanation=pop(@P) )
9.          if    ( check conflict ( explanation,@facts ) == 0 )
10.                 output @explanation
11. endwhile
```

Fig. 3. Explanation Correlation

function $readpath(OP)$ returns the current best path of each OP. The function $checkconflict$ is used to compare P and the "$facts$" for conflicts.

Using the same example in Table II, our objective is to provide a possible explanation for Line 3. First, there are two OPs (AS3549 and AS9177) and we know both AS9916 and AS15412 claimed the ownership of prefix 140.113.0.0/16 (MOAS event). The first MOAS was observed by AS3549 at 5:27pm (Line 1). The announcement from AS15412 continued to propagate to AS9177 at 5:28pm(Line 2). Suddenly, as shown in Line 3, from AS3549's point of view, the prefix ownership changed back to AS9916, i.e., "MOAS was fixed."

To analyze the changes in "withdraw" scenario, we have the following possibilities: "AS15412 withdraw AS15412", "AS3561 withdraw AS15412" and "AS3549 withdraw AS15412."

By using $readpath$(AS9177), we have some conflicting facts from AS9177: "AS15412 announce AS15412" and "AS3561 announce AS15412". Please note that Table II Line 5 also implies the same facts. Line 5 indicates the AS path did not change until April 10, which implies "AS15412 announce AS15412","AS3561 announce AS15412" on April 6 since the old path from AS15412 was still valid.

By removing the conflicts, we have only one possible policy change, which is "AS3549 withdraw AS15412." Therefore, AS3549 is the possible AS who blocked the bad route from AS15412 at 5:42pm on April 6, 2001 (Line 3).

The next question is "why did AS3549 withdraw the route from AS15412?" While the previous AS path was still valid (path segment "3561-15412" was still in use until April 10), why would AS3549 use a longer route? The most plausible explanation is that AS3549 found the error (of AS15412) and took an early action to fix it by blocking the bad announcement from AS15412. Then, AS3549 selected another route from its routing table, which was the path originated by AS9916.

IV. Evaluation

In this section we discuss how we validate our algorithm. A challenge of validation is to define a suitable evaluation metric. We propose our metric and show the results quantitatively.

A. MOAS Dataset

This section presents results of our study based on the dataset collected from three major MOAS storms – two in April 2001 and one in December 2006.

The data we used was collected by rrc00.ripe.net - a RIPE BGP data collector [9]. First, we convert BGP raw data into formal notation as describe in the language L. Please note that some prefixes not affected by MOAS storm are shown as "null" in L's notation. With L's notation, the prefixes affected by MOAS storms are shown as a series of AS path changes.

We focus on studying BGP MOAS events. We want to answer the following:

- Was the MOAS fixed by the faulty originator withdrawing the false announcement?
- If not, what happened? Which ASes changed their policies to block bad routes?

For legitimate MOAS instances, it is a reasonable assumption that not all ASes can afford to wait for days until the faulty originator fixes the problem. Instead, some ASes may try to fix or reduce MOAS themselves by withdrawing the bad route[4]. Such actions will be identified as "withdraw."

B. Quality Measurement

We derive a number of possible explanations (series of actions) for all AS path changes. We are interested in finding the explanations that best fit real BGP MOAS traffic.

During MOAS, for each prefix, there are only two AS path changes that actually matter, one is considered "damage", i.e., to change to bad routes, the other is considered "fix", i.e., to change back to good routes. We have already determined which ASes caused the damage – AS15412 in 2001 and AS9121 in 2004. Thus, the task is to find the set of actions that fit into "AS_X withdraw AS15412".

We consider an explanation a "match" if and only if a series of BGP actions can be used to explain a particular set of AS path changes. For example "AS_X withdraw AS15412," is a reasonable explanation of an AS path change, and of course, causes a change from the incorrect path to the correct path. Thus, we are able to identify that AS_X "fixed" the damage.

Please note that a series of BGP actions may "fix" more than one prefix. i.e., the same series of BGP actions may lead to multiple "matches." Therefore, we can quantify the quality of a particular series of BGP actions based on the number of AS paths which have a match. Using the same example in Table II Line 3, because "AS3549 withdraw AS15412" can explain the AS path change "3549 3561 15412 $\Rightarrow$ 3549 7018 1659 9916", there is a match.

In a nut shell, a better explanation (series of BGP actions) implies more matches and "fix" more prefixes. The measure of quality Q is defined as the percentage of the total affected prefixes fixed by a particular series of BGP actions.

After determining quality Q for each explanation (series of actions), we are able to evaluate which set of the combination of explanations best fit real BGP traffic. To solve the problem, we use a greedy algorithm. First, we obtain the quality of each explanation (each series of actions). Then select the candidate (one explanation) with the best quality and add the candidate to our final solution. We keep adding the next best candidate until our final solution matches our goal G. The goal is the user defined parameter – percentage of total traffic matches.

[4]In this paper, we do not focus on how ASes consider a route "bad".

AS	(A)	(B)	(C)
3561	58.6%	41.40%	0.00%
7018	57.8%	0.20%	42.02%
209	41.7%	7.29%	51.06%
2914	41.4%	1.24%	57.40%
9057	25.6%	3.28%	71.09%
3356	24.4%	2.48%	73.07%
2497	24.4%	8.70%	66.91%
701	21.8%	3.47%	74.75%
3549	21.7%	0.01%	78.33%
1	20.4%	0.96%	78.65%
1103	18.0%	0.10%	81.88%
8210	17.9%	27.41%	54.66%
1755	17.0%	0.81%	82.19%
1221	15.3%	0.58%	84.12%
4777	8.9%	17.81%	73.32%

(a) April 6, 2001

AS	(A)	(B)	(C)
3561	41.4%	58.6%	0.0%
8210	34.2%	57.4%	8.4%
9177	28.0%	57.4%	14.5%
209	5.8%	0.0%	94.2%
3333	5.5%	0.5%	93.9%
286	5.5%	0.0%	94.5%
6730	2.4%	0.0%	97.6%
701	2.0%	0.8%	97.2%

(b) April 12, 2001

AS	(A)	(B)	(C)
6762	54.32%	45.61%	0.07%
3333	48.08%	45.57%	6.35%
174	15.30%	0.04%	84.66%
701	13.60%	0.00%	86.40%
4637	9.55%	0.06%	90.39%
3741	7.56%	0.05%	92.39%
3549	7.13%	0.06%	92.81%
3320	6.62%	0.07%	93.32%
2497	6.23%	0.07%	93.70%
2914	5.81%	0.06%	94.14%
6453	5.21%	0.04%	94.74%

(c) December 24, 2004

Fig. 4. Prefixes Fixed by AS Actions

A G set to 100% is literally setting the greedy algorithm to find all BGP actions which explain 100% of the BGP traffic.

C. Analysis of Real BGP Traffic

In this section, our goal is to find which set of BGP actions "fix" most of the MOAS. We set our goal G equal to 97%.

We use the data of the first 8 hours of MOAS storm on April 6, 2001(AS15412 falsely announced 30088 prefixes), April 12, 2001 (AS15412 falsely announced 920 prefixes) and December 24, 2004 (AS9121 falsely announced 106793 prefixes).

In Figure 4, we show the percentage of prefixes that have been fixed versus actions taken by different ASes. Since we focus on the "fixing" process, the BGP actions are "AS_X withdraw AS15412" for Figure 5(a), 5(b) where X is shown in the AS column and "AS_y withdraw AS9121" for Figure 5(c) where y is shown in the AS column.

Column (A) represents the quality Q – percentage of prefixes which a particular policy change (action) matches. For instance, in Figure 4(a), "AS7018 withdraw AS15412" matches 17385 prefixes (57.8%= 17385 / 30088). Hence, AS7018's action rescued 57.8% of affected prefixes. Column (B) represents the percentage of prefixes which a particular policy change does "NOT" match. (i.e. a bad route was in use and not fixed). For instance, "AS3549 withdraw AS15412" did not match 0.01% of all affected prefixes. AS3549 still used bad routes from AS15412 for 0.01% of all affected prefixes. Column (C) represents $100\% - (A) - (B)$ which refers to prefixes for which we cannot make a determination based on the given BGP data.

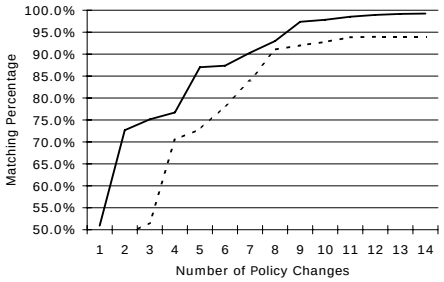

(a) April 6, 2001

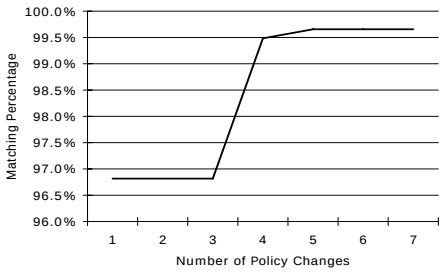

(b) April 12, 2001

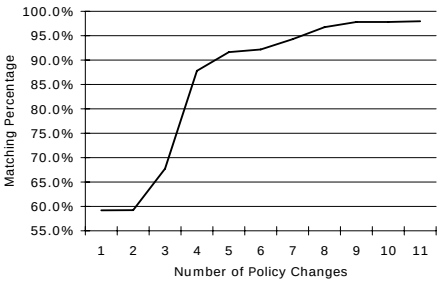

(c) December 24, 2004

Fig. 5. Real Traffic Matches versus Number of Actions(Policy Changes)

As a reminder, ASes may have different policies for different prefixes. For instance, AS3549 might decide to rescue prefix 140.113.0.0/16 but not prefix 169.237.0.0/16. Also, more than one AS might rescue exactly the same prefix by blocking the bad traffic. For instance, our results indicate both AS7018 and AS701 blocked the faulty route from AS15412 to rescue prefix 140.113.0.0/16. Hence, the sum of Column (A) may exceed 100%. An interesting question arises: "did ASes collaborate to rescue prefixes?"

Then, we apply the greedy based algorithm to find out which set of BGP actions reaches our goal. Figure 5 represents the results of three different MOAS storms corresponding to Figure 4. Figure 5(a), Figure 5(b) and Figure 5(c) show the results corresponding to Figure 4(a) ,Figure 4(b) and Figure 4(c) respectively. The Y-axis shows the percentage of matches. The X-axis represents the number of "series of BGP actions." The results are easy to interpret, for example, in Figure 5(c), 5 actions (rescuing actions took by 5 ASes) explain 91.6% of real BGP MOAS traffic.

In April 2001, AS15412 did not fix the problem until several days later while some ASes took actions to fix MOAS prior to AS15412. There were a total of 124090 fix-

related "PATHCHANGES." As shown with the solid line in Figure 5(a), 3 actions fixed 75% and 10 actions fixed 98% of the damaged Internet. In other words, 10 actions explain 98% of those AS path changes. It is surprising to see only a small number of actions are sufficient to fix most of the MOAS problems. Besides, data shows ASes respond to MOAS quickly, for example, "AS3549 withdraw AS15412" within 15 minutes (Table II).

We have another interesting finding when computing Figure 4(a), column (B). There are exactly the same 12455 prefixes which both AS3561 and AS15412 did not rescue (i.e. still used bad routes). Hence, there is a good chance that in the peering relationship "AS3561-AS15412", the BGP in-bound filter of AS3561 was "accept all" such that AS3561 believed and propagated all BGP announcements from AS15412. Since AS3561 most likely had an in-bound filter of "accept all" from AS15412, it is very interesting to see the effect if AS3561 took no action to fix any prefix within the first 8 hours (even AS3561 claimed they fixed the problem) [11]. We remove one particular action – "AS3561 withdraw AS15412" from the candidates in Figure 4(a) and have the results plotted on a dotted line in Figure 5(a). Surprisingly, the results are similar to what we have shown on the solid line — it takes only a small set of ASes to rescue the majority of damaged Internet routes.

Figure 5(b) shows a small scale MOAS (1124 "rescuing" actions) caused by AS15412 and Figure 5(c) shows a large scale MOAS (115162 "rescuing" actions) caused by AS9121. Comparing the MOAS on April 6 and April 12, we find the number of affected prefixes was much smaller on April 12 (around 1K) than on April 6 (around 30K). The scale of MOAS on April 12 2001 was indeed smaller than in 2004. It is clear to see that fewer actions are needed to fix the Internet. For example, in Figure 4(b), it took only three ASes to rescue 99% of the Internet.

All three sets of experiments indicate that only a few transit ASes are able to mitigate MOAS damage by blocking/withdrawing invalid routes. This makes perfect sense. Due to business relationships, most traffic is usually routed through certain AS paths(primary paths). Thus, it is fairly easy to block inappropriate BGP traffic at certain points where traffic converges, causing the effects of the MOAS to be dramatically reduced.

V. Conclusion and Future Work

In this paper, our goal is to understand and learn how network operators react to BGP events, such as MOAS instances which disrupt network traffic.

Ideally, as part of the critical Internet infrastructure, BGP is expected to behave robustly against human errors and intentional attacks. However, by using a naive trust model without adequate validation checks, BGP does not achieve that expectation. The BGP MOAS problem is one of many examples that illustrate such vulnerabilities. These BGP vulnerabilities have accounted for several Internet-wide incidents since the late 90's. Most of the BGP research work has focused on how to prevent BGP from being misused. In this paper, we examine MOAS from a different angle – we take a closer look at past MOAS events and focus on how people fixed MOAS in a timely manner.

First, we show how to identify policy changes which "fix" the Internet and provide a validation process. This framework can be used by other researchers to learn about BGP local policies which have been kept as top secret within ASes. Second, we show that it is not always the case that the originating AS withdrew the bad announcements to fix the problem; instead, other ASes took early actions to stop the bogus information. We also determine that there were ASes who used non-suitable configurations(e.g. accept-all), which magnified the impact of MOAS by propagation. Another very interesting observation is that a few ASes took early actions and it benefited a large portion of the Internet. Still there are some open questions: what factors did those few ASes consider before creating their policy? What are the relationships between those few ASes?

We also propose a formal notation to describe and analyze MOAS events. The language can be extended to describe more attributes and provide more detail and complex explanations. Using the language, we are capable of learning automatically from BGP traces by utilizing data mining techniques.

One of the requirements from network operation is to automate the detection of and reaction to network events. In this regard, we would like to build an auxiliary system for Inter-domain routing which is capable of detecting and correcting invalid entries automatically in real time. The techniques and algorithms used by this paper can be further utilized in building such a system. The observations and findings revealed by this paper may provide some heuristics for its deployment. We see this paper as one step further in the direction of building an automated network management system.

References

[1] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, "An Analysis BGP Multiple Origin AS(MOAS) Conflicts," in *IMW '01: Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*. New York, NY, USA: ACM Press, 2001, pp. 31–35.

[2] T. Bates, P. Smith, and G. Huston, "Cidr report," http://www.cidr-report.org/as2.0/.

[3] Y. Rekhter and T. Li, "A border gateway protocol 4 (bgp-4)," RFC1771, 1995.

[4] S. Misel, "Wow, as7007!" http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html, April 1997, North American Network Operators Group.

[5] S. M. Tseng, K. Zhang, S. F. Wu, K.-L. Ma, S. T. Teoh, and X. Zhao, "Analysis of BGP origin as changes among brazil-related autonomous systems," in *IP Operations and Management (IPOM 2007)*, November 2007.

[6] M. A. Brown, "Pakistan hijacks YouTube," http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml, Feb. 2008.

[7] S. T. Teoh, K.-L. Ma, and F. S. Wu, "A visual exploration process for the analysis of internet routing data," IEEE Visualization 2003 Conference. IEEE, 2003.

[8] S. T. Teoh, K.-L. Ma, F. S. Wu, D. Massey, X. Zhao, D. Pei, L. Wang, L. Zhang, and R. Bush, "Visual-based Anomaly Detection for BGP Origin AS Change Events." 14th IFIP/IEEE Workshop on Distributed Systems: Operations and Management, 2003.

[9] "RIPE NCC," http://www.ripe.net/.

[10] "The Route Views Project," http://www.routeviews.org/.

[11] J. A. Farrar, "Re: C&W routing instability," http://www.merit.edu/mail.archives/nanog/2001-04/msg00209.html, April 2001, Nanog Mailing list.