Validation of the MOAS conflicts through Assertions *

Xiaoliang Zhao, Dan Pei, Lan Wang, Dan Massey, Allison Mankin, S. Felix Wu,Lixia Zhang†

December 7, 2001

Abstract

In BGP routing updates a Multiple Origin AS (MOAS) conflict occurs when a particular IP address prefix appears to originate from more than one autonomous system (AS). In certain cases MOAS conflicts are a result of current operational practice; in other cases, however, MOAS conflicts are caused by mis-configurations, protocol implementation errors, or even intentional attacks. The current BGP specification and implementation provide no distinction between the former and the latter cases. In order to detect network faults and attacks, in this paper we present a simple spoof-resilient solution to the MOAS conflict problem. Our solution is backwards compatible and incrementally deployable. It allows BGP routers to distinguish valid MOAS conflicts from invalid ones, thus it substantially reduces, if not eliminate, the risk of traffic hijacking through false routing announcement due to either malicious attacks or protocol faults.

1 Introduction

This paper presents a backwards compatible and incrementally deployable BGP extension for detecting invalid Multiple Origin AS (MOAS) conflicts. The Internet is made of thousands of Autonomous Systems, loosely defined as a connected group of one or more IP prefixes which have a single and

*This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No DABT63-00-C-1027. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the DARPA.

†xzhao@unity.ncsu.edu, peidan@cs.ucla.edu, lanw@cs.ucla.edu, masseyd@isi.edu, mankin@isi.edu, wu@cs.ucdavis.edu, lixia@cs.ucla..edu

clearly defined routing policy[1]. BGP[2] is the standard inter-AS routing protocol. A BGP route lists a particular prefix (destination) and the path of ASes used to reach that prefix. The last AS in an AS path should be the origin of the BGP routes. A Multiple Origin Autonomous System (MOAS) conflict occurs if a prefix appears to originate from more than one ASes. More precisely, suppose prefix d is associated with AS paths $asp_1 = (p_1, p_2, \ldots p_n)$ and $asp_2 = (q_1, q_2, \ldots q_m)$. We say a MOAS conflict occurs if $p_n \neq q_m$.

In an effort to improve the fault-tolerance and security properties of the BGP routing protocol, we have been measuring the behavior of BGP. The MOAS conflicts are interesting to us for a number of reasons. First, RFC 1930[1] recommends that a prefix should originate from a single AS, but MOAS conflict may occur for a limited number of valid reasons. Second, MOAS conflicts could be the result of a fault or an attack, where a BGP router falsely originates routes to some other organization's prefixes. There is no BGP mechanism for distinguishing between the valid MOAS that occur due to legitimate operational policies and the MOAS conflicts caused by faults or attacks. We propose a new use of the existing BGP community attribute and a new DNS resource record that allow us to distinguish the valid conflicts from the invalid ones.

The remainder of the paper is organized as follows. Section 2 reviews the related work. Section 3 describes the MOAS conflict data and the methodology we used to collect and process the data. Section 4 provide detailed analysis of the results and our explanations of the results. Section 5 proposes to use assertion to validate the MOAS conflicts. Section 7 summarizes the paper.

2 Related work

MOAS conflicts have been observed by a number of researchers, but no one has considered the problem in detail. The most relevant work comes from Geoff Huston's BGP Table Statistics website[3]. Starting on 2/18/2001, this site began tracking a daily count of MOAS conflicts[1] using data from some ISPs and from the Oregon Route Views Server. On 04/19/2001, the website switched to tracking MOAS conflicts on a bi-hourly instead of daily basis. However, the BGP Table Statistics work provides only a basic count of MOAS conflicts and no further explanations or analysis is offered.

[1] Huston uses the term "multiple-origin prefixes" in place of our term "MOAS conflicts"

The MOAS conflict issue has also been discussed within the IETF. RFC 1930[1] recommends that a prefix should belong to only one AS. If this recommendation was followed, MOAS conflicts would not occur, with the possible exception of a few unique cases discussed further in [4]. This view of MOAS conflicts does not correspond with the MOAS conflicts actually seen in the Internet.

There is also related work on determining the correct origin AS (or ASes) for a particular prefix. Bates et al [5] proposes to use DNS to store (prefix,origin AS) pairs in the originator's DNS side and then each incoming route update can be checked against the DNS record. However, this approach requires a new domain name hierarchy must be deployed in the current DNS system, which is considered to be a non-trivial administrative task. In addition, BGP routers must interact with the DNS when making routing decisions and DNS data are assumed to be correct. [6] proposes a filtering model that uses the Internet Route Registry (IRR) records to do the similar checking. However, updating the IRR record is not a mandatory operation for ISPs and some IRR records are outdated or inaccurate, reducing the effectiveness of this approach. Kent [7] proposes to use some form of public key infrastructure (PKI) to verify the origin of the route advertisement. However, this approach calls for a significant modification of the current Internet infrastructure.

3 MOAS Conflicts in the Internet

Data observed from Internet routers shows that MOAS conflicts are a common occurrence. Based on the number of conflicts and the conflict durations, we believe the nature of these conflicts differs from what one might expect based on documents such as [1]. Using off-line techniques, we identified examples of both valid caused by legitimate operational policies and and invalid MOAS conflicts caused by faults. A summary of the MOAS conflict data is provided below and a more detailed analysis is provided in [4].

3.1 Methodology

In our previous work[4], we provided detailed analysis of MOAS conflicts based on the data from the Oregon Route Views server [8]. Currently, the Oregon Route Views server peers with 54 BGP routers in 43 different ASes. Each peer exports its BGP routing table to the Route Views server.The routing table of Oregon Route Views server is archived on a daily basis from 11/08/1998 to present by NLANR[9] and PCH.net [10]. The large number

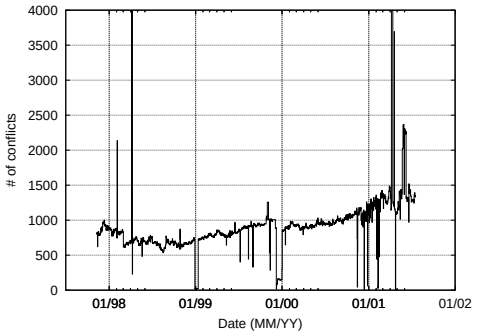

Figure 1: The number of MOAS conflicts from 11/1997 to 07/200

Year	Median of MOAS conflicts	Increasing rate
1998	683	
1999	810.5	18.7%
2000	951	17.3%
2001	1294	36.1%

Figure 2: Median of MOAS conflicts per year

of AS peers and access to archived routing data makes the Oregon Route Views data an ideal location for observing MOAS conflicts.

3.2 Results

3.2.1 Total Number of MOAS Conflicts

Figure 1 shows the total number of conflicts from 11/08/1997 to 07/18/2001 [2]. Overall 38225 conflicts were observed over 1279 days. The median number of MOAS conflicts for each year are listed in the Figure 2. There is an increase from 683 conflicts in 1998 to 1294 conflicts in 2001.

3.2.2 Duration of MOAS Conflicts

Figure 3 shows the duration of MOAS conflicts, based on the data (Figure 1). Figure 3 shows that most of the conflicts are short-lived. 13730 out of 38225 conflicts appeared only once and lasted less than one day. 11358

[2]The maximum of conflicts reached 11842 on 04/07/1998 and 10226 on 04/06/2001.

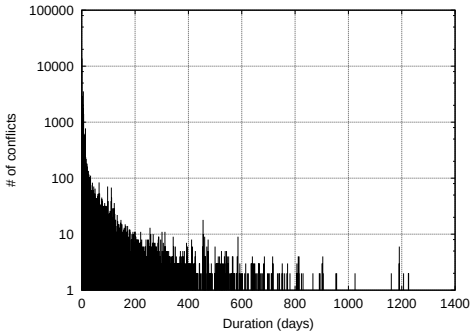

Figure 3: Duration of MOAS conflicts

Expectation (days)	Measured data set
30.9	longer than 0 day
47.7	longer than 1 days
107.5	longer than 9 days
175.3	longer than 29 days
281.8	longer than 89 days

Figure 4: Expectation of the duration of MOAS conflicts

of these one-time conflicts can be attributed to a configuration fault that occurred on April 7th, 1998. Excluding the one-time conflicts, the expectation of the duration is 30.9 days. Taking into account that many other short-lived conflicts might also be due to faults, we considered the data set which contains only conflicts whose duration is greater than 9 days (a total of 10177 conflicts). For these conflicts, the expectation of the duration is 107.5 days with 1002 conflicts lasted longer than 300 days. Figure 4 lists the expectation of the duration from the different data sets. The longest duration was 1246 days out of a possible 1279 days and 1326 conflicts were still ongoing as of the date the paper was written.

The duration of an individual conflict counts the total number days of the conflict was in existence, regardless of whether the conflict was continuous and whether the same ASes were involved.

Overall, these results show that there are number of MOAS conflicts may be caused by some other factors than multi-homing which usually implies longer MOAS conflicts durations.

4 Causes of MOAS Conflicts

There are three main causes of MOAS conflicts: some forms of multi-homing, faults or attacks, and special prefix types. [4] describes each of the potential causes in more detail and confirmed their occurrence in the Internet. In order to show our proposal can distinguish valid from invalid causes, we present the summary of causes in the following.

4.1 Special Prefix Types

Some prefixes are naturally associated with more than one AS. In particular, the prefixes associated with BGP exchange points may be advertised by each AS connected to the exchange point. Anycast address also can be advertised by several ASes. In both cases, MOAS conflicts are expected and simply reflect the fact that multiple ASes can directly reach the prefix. RFC 1930[1] also notes that aggregation could result in routes that end in AS sets and the concept of MOAS conflict is not well defined with respect to an AS set. In our study, we typically observed 12 prefixes which ended in AS sets and these AS sets were consistent with each other. Overall, the number of MOAS conflicts caused by special prefix types make up only a small percentage of the total MOAS conflicts observed in the Internet.

4.2 Multi-homing

Various forms of multi-homing policies can result in MOAS conflicts. In some cases, an organization may connect to multiple ISPs but not use BGP with all these ISPs. In this case, these ISPs may be statically configured to advertise the organization's prefixes or may import these prefixes from the IGP routing protocol. In both cases, each ISP originates the BGP route to the prefix and each ISP uses its own AS as the originating AS. From a BGP perspective, it appears as if each AS can directly reach prefixes belonging to the organization and the result is a MOAS conflict.

One would expect these conflicts to be long lasting in duration since static routes or non-BGP peering agreements are likely to have a long lifetime. Note that these MOAS conflicts are valid, but could present a problem for packet forwarding if the links necessary to support the static routes fail.

In a similar case, an organization may peer with several ISPs using a private AS number. To prevent AS number exhaustion, Haas [11] suggests that a multi-homed customer uses a private AS number which is mutually agreeable to all providers. This technique is called AS number Substitution

on Egress (ASE). If deployed, this approach could produce MOAS conflicts because the private AS number should be stripped off by the upstream providers and the real origin information will be lost.

Because the links using non-BGP routing mechanisms or private AS numbers are "hidden" to BGP, a BGP router can not tell whether or not a MOAS conflict is due to multi-homing without BGP or multi-homing with private AS number. However, by contacting individual ASes, we did confirm such occurrences.

4.3 Faulty or Malicious Configurations

MOAS conflicts can also occur when an AS incorrectly originates routes to some other organization's prefixes. This could occur due to configuration errors or even intentional attacks. Often, the faulty AS does not have a route to the incorrectly originated prefixes and packets that use the incorrectly originated route will reach the faulty AS and then be lost.

Figure 1 shows several notable examples of MOAS conflicts caused by faults. The graph shows a large spike on April 7th, 1998 and AS 8584 was involved in 11357 out of 11842 conflicts that occurred during that day. Discussions on a network operators mailing list[12] indicated that AS 8584 falsely originated routes to those conflicted prefixes. Consequently, some ASes selected the incorrectly originated route. Packets sent along this incorrectly originated route would reach AS 8584 and would then be lost.

The graph also shows a large spike on April 10th, 2001 and the sequence (AS 3561, AS 15412) was involved in 5532 out of 6627 MOAS conflicts that occurred during that day. Based on the archived data from RIPE RIS [13], AS 15412 normally originates only 5 prefixes. However, on April 6th, AS 15412 suddenly originated thousands prefixes due to a configuration error[14].

On April 25th, 1997, a severe Internet outage occurred[15] when one ISP falsely de-aggregated most of the Internet routing table and advertised the prefixes as if they originated from the faulty ISP[16]. The falsely originated prefixes resulted in MOAS conflicts. These examples show that invalid MOAS conflicts do occur and can have serious impacts on Internet routing.

Faulty aggregation could also cause MOAS conflicts. In faulty aggregation, an AS advertises an aggregated prefix, even though some of more specific prefixes are not reachable by the AS. A MOAS conflict occurs if an aggregate route is also generated by some other AS. Packets that use the faulty aggregated route will travel to the faulty AS and then may not be able to reach all the more specific prefixes.

4.4 Using Duration as Heuristic to Validate MOAS Conflicts

With the exception of faults and intentional attacks, the possible explanations should have created long duration MOAS conflicts. MOAS conflicts for exchange point prefixes should remain as long as two or more ASes choose to advertise a route to the exchange point. The data confirmed this expected pattern and exchange point MOAS conflicts persisted for most, if not all, of the study. Multi-homing without BGP and multi-homing with Private AS numbers both require router policy configurations at two or more ASes and the resulting MOAS conflicts should persist for as long as the multi-homing policy remains in place. We expected that multi-homing policies (and the resulting MOAS conflicts) would occur over months, not days. But the data in Section 3.2 shows a large number of short duration conflicts.

One possible reason for short-lived MOAS conflicts is that MOAS conflicts could occur during a transition period when a non-BGP customer switches from one provider to another. To guarantee the connectivity to the non-BGP customer, it is possible for both providers to originate the customer's prefix for a short period. Another possible and more likely reason for short-lived MOAS conflicts is router mis-configurations or other faults. These conflicts disappear when the faults are detected and corrected.

Overall, the duration can be a useful heuristic to distinguish between valid MOAS conflicts and invalid ones. However, such differentiation is not accurate enough to be a solution to validate MOAS conflicts.

5 Validation of MOAS conflict

5.1 Assertion

From the standpoint of fault tolerance and security, MOAS conflicts pose an interesting challenge. On the one hand, MOAS conflicts can occur for valid reasons, such as multi-homing without BGP and advertising exchange points addresses. Packet forwarding is not adversely affected if each AS in the MOAS conflict has a (non-BGP) route to the destination. On the other hand, router misconfigurations or intentional attacks could also produce MOAS conflicts. Packet forwarding is adversely affected if the packets rely on the invalid routes. Large scale network outages and other problems have been associated with MOAS conflicts. Thus, the validation of MOAS conflicts is necessary.

The basic idea is establish a mutual agreement between the multiple ASes who are entitled to originate a particular prefix. This agreement is at-

tached to the route advertisement so that anyone observing a MOAS conflict can verify that the conflict has been authorized by all the prefix originators. If a fault or attack creates a conflict, this faulty route will not be in agreement with the valid advertisement. This idea is captured by the following assertion:

Assertion 1 *If prefix p is originated from multiple ASes, then each originating AS must list all the potential origin ASes and the AS path must end in one of these ASes.*

Any MOAS conflicts violating the above assertion will be regarded invalid, then an alarm should be generated and further investigation should be followed.

For example, suppose multi-homing allows prefix p to be originated by both AS 1 and AS 2. Each AS will originate the route and attach the agreement indicating that both AS 1 and AS 2 can originate the route. For an exchange point prefix, any AS attached to the exchange point may choose to originate the route. An AS attached to the exchange point and originating the route should attach an agreement listing all the other ASes at the exchange. [3]

We claim the assertion based approach has the following properties:

Effectiveness: Assuming an origin AS has been configured to attach the agreement with all its route announcements, when one of its prefixes is falsely originated by either a fault or a compromised AS, other ASes can easily detect an invalid MOAS conflict. The false route announcement either carries an origin AS that is not listed in the correct agreement, or carries an agreement which is different from the one generated by the valid origin ASes.

However, because of routing policies and route propagation rules, these conflicting routes will not be observable at all points in the Internet. But at some points in the network, a BGP router will receive both of the valid route and the incorrectly originated route. This router detects a MOAS conflict. We can not make any claim about the properties of the invalid route and what agreements it might attach. But we can be certain that the valid route has not agreed upon a MOAS conflict with the faulty AS. This MOAS conflict is not agreed upon by both parties and can be definitively identified as an invalid MOAS conflict.

[3] Note that this does not imply that all ASes at the exchange actually do advertise the prefix. It only indicates that any of them may choose to do so.

Spoofing resilience: If the origin ASes has only one path to reach the rest of the Internet and the attacker blocked theirs correct route p announcements, the potential invalid MOAS conflicts cannot be easily detected. In this case the attacker has compromised the only path to reach p and can cause other arbitrary damage to p as well. However in more general cases where multiple origin ASes are making route p announcements, or an origin AS announces routes to multiple AS peers, then it is difficult for an attacker to block the correct route p announcement carrying the agreement. In this case even though an attacker can still inject invalid route with false MOAS list, other routers can easily detect the agreement inconsistency and alarm the network operator. Similarly, an attacker may try to modify the agreement on a valid announcement. As long as the attacker can not modify all routes to the origin, the change may result in an alarm.

The approach of attaching the agreement to the route advertisements might be fragile in some situations. For example, a router may decide to remove the agreements because of local policies or route aggregation. To further enhance the effectiveness of our assertion-based approach, we also suggest to put the agreements in DNS records. One could check the MOAS conflicts with DNS in the cases that the agreements information attached with the route advertisements is incomplete. In addition, supported by DNS, a router may be able to further identify the faulty AS.

5.2 Attaching Agreements to Routes

We use the existing BGP community attribute[17] to define a new BGP community, termed "MOAS List", to represent the agreement. The community attribute is an optional transitive attribute of variable length, which can be used to convey additional information to the global routing system for a group of prefixes that share some common properties. Each community attribute consists of four octets. By convention, the first two octets are used to encode an AS number and the semantics of the final two octets may be defined by the AS listed in the first two octets. We reserve one of the 2^{16} values available in the last two octets to indicate a MOAS List. This value is denoted by $MLVal$ in the remainder of this paper. Thus the community attribute $(X : MLVal)$ indicates that the route originator agrees AS X may originate a route to this prefix.

For example, if a prefix p which is originated from a set of ASes, AS_1, AS_2, ... AS_n, the route updates from $AS_i, (1 \leq i \leq n)$ will include the

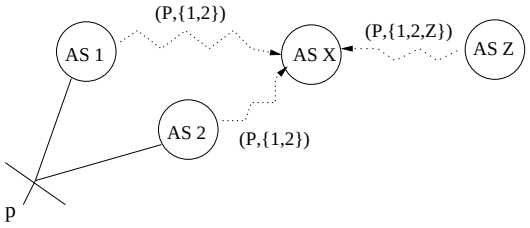

Figure 5: Example scenario

AS 1
$MLVal$
AS 2
$MLVal$

Figure 6: Example MOAS List

MOAS List $(AS_1, MLVal), ..., (AS_n, MLVal)$ In Figure 5, AS 1 and AS 2 agree that both of them may originate routes to the same prefix p and when AS 1 originates p, AS 1 will attach the MOAS List, shown as in Figure 6, and the corresponding configuration file is shown in Figure 7. So does AS 2.

When a BGP router receives route announcements for the same prefix p from multiple peers, it must check to see whether all the MOAS Lists for p are consistent. Here consistency is defined as the same set of ASes listed in the MOAS List, the order in the list may differ; if the number of ASes in the list from different peers differs, or if the AS set is different, we consider the MOAS List inconsistent.

router bgp AS1
neighbor a remote-as Y
neighbor a send-community
neighbor a route-map setcommunity out
route-map setcommunity
match ip address p
set community AS1:$MLVal$ AS2:$MLVal$

Figure 7: Example configuration file

If the BGP router notices any inconsistency in the MOAS Lists received from multiple peering neighbors, it should generate an alarm signal to inform the network operator. The operator can then further investigate the cause of the inconsistency. Also, in this case, the local policy will decide to accept the corresponding route or not.

In the example shown in Figure 5, both AS 1 and AS 2 attach a MOAS List, $(AS1 : MLVal, AS2 : MLVal)$, to their route announcement update for p. If AS Z also falsely originates p with a MOAS List $(AS1 : MLVal, AS2 : MLVal, ASZ : MLVal)$, assuming all three routes are propagated to AS X, AS X will observe an inconsistent MOAS List and should generate an alarm signal to the network operator.

Attaching MOAS List to route announcement requires only BGP configuration changes, but checking MOAS List consistency may need BGP implementation be modified accordingly. However, to quickly deploy the MOAS List checking in the operational Internet, one could run an offline monitoring process which periodically downloads the BGP routing table and check the MOAS List consistency from multiple peers. If the router is equipped to support the new BGP MIB [18], one could run a management application to access the BGP MIB to get all MOAS List, along with other information, through the MIB interface and check the MOAS List consistency.

There are some practical concerns. A MOAS List for prefix p might be dropped along some of the propagation paths due to local policy, aggregation, or improper configuration at some BGP routers. A router further down the road that receive route P announcement may see some of the announcements carrying a MOAS List, and some not. However, all of received MOAS Lists should be consistent. Otherwise, the router should generate an alarm signal. More detailed discussion can be found at [19].

5.3 Storing Agreements in DNS records

Similar to the work[5], we use DNS records to store the agreements. We introduce a new DNS Resource Record (RR) dedicated for verification of route updates involving a MOAS conflicts. The new RR, denoted by $MOASRR$, consists of two fields, 'Prefix Length' and 'AS List'. The 'Prefix Length' field contains an octet encoding the length of the address prefix associated with the node named by the RR (0 through 32). The 'AS List' field contains an list of origin ASes, each encoded as a two octet unsigned integer (0 through 65535). More detailed information about the implementation and lookup procedure could be found in [5].

In above example, name server for prefix p will add a new RR in the

zone file. The new RR will looks like:

$MOASRR\ x$ AS1 AS2

where the x denotes the prefix length.

To verify the origin AS of p, one could do a DNS lookup for prefix p by specifying the type as $MOASRR$. If the origin AS appearing in a route update doesn't match any AS number in the AS list of $MOASRR$ record, such route update should be considered as a bogus update, therefore should be dropped.

There are some practical concerns. One would be the vulnerability of the current DNS system[**?**]. If the DNS system can not guarantee the correct answer, putting agreement in the DNS record can not be effective. DNSSEC[**?**] will be a security solution for this problem. However, taking the deployment of DNSSEC into account, we do not require DNSSEC as a necessary component. Another concern would be the compatibility of current DNS software. By introducing a new RR, it implies a modification of current DNS software which may take longer time for implementation and deployment. To get around this problem, we could use TXT record as a temporary solution.

6 Simulation Results

Our proposed extension to BGP, the MOAS community attribute, enables routers to easily detect inconsistency in routing announcements, and thus it protects the routing system from potential attacks or faults. However, there is still a possibility that attackers may block valid routing announcements to propagate to parts of the Internet. Therefore, in this section, we evaluate the effectiveness of our scheme in protecting networks from false routing announcements. More specifically, we let attackers inject false routing announcements at randomly selected locations, and then we compare the damage they cause with and without our scheme. The simulation results show that our scheme substantially reduces the percentage of networks that adopt the false routes. Furthermore, we have observed that with our scheme, larger networks are even more robust against randomly selected attackers. Finally, we demonstrate that our scheme can still protect the routing system against false routes even when it is partially deployed.

In the following sections, we first describe the simulator and the topologies used in our simulation and then present the results from three experiments.

6.1 Simulation Setup

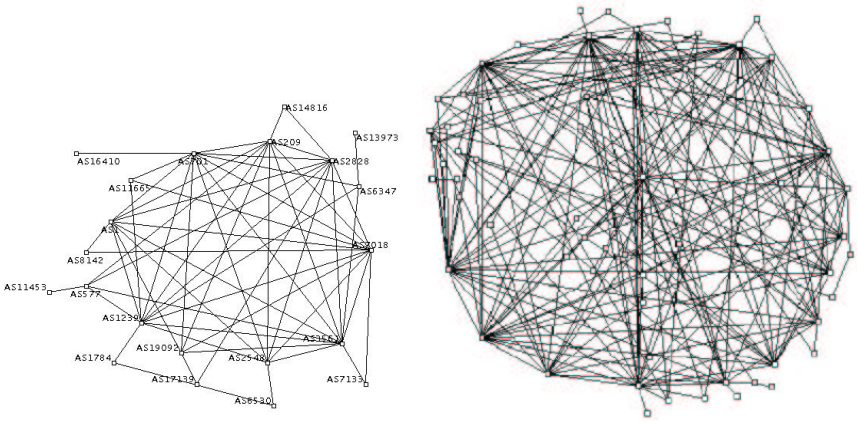

(a) 21-AS Topology (b) 63-AS Topology

Figure 8: Simulation Topologies

We use the SSFnet [20] simulator in our simulation, and we have modified the BGP protocol in SSFnet to support MOAS detection. Figure 8 shows two topologies that we mainly use, one with 21 nodes and the other with 63 nodes (see Figure 8). In these topologies, each node represents an Autonomous System (AS), and a link between two nodes represents a BGP peering relationship (i.e. the two AS's exchange routing information).

To generate these topologies, we first get the BGP routing table from the Oregon RouteViews server [8]. Then we infer BGP peering relationship based on the AS Path attribute in BGP routes. For example, if a route to a prefix p has the AS Path *1239 6453 4621*, we consider AS6453 to have two BGP peers, AS1239 and AS4621. We also mark AS6453 as a *transit AS* since packets to and from AS4621 may traverse through it (note that AS1239 is also a transit AS). If an AS does not appear to be a transit AS in any of the routes, we consider it a *stub AS*. Transit AS's are usually ISP's (e.g. AS1239 is Sprint), while stub AS's are networks at the edges of the Internet such as small organizations and universities. Next, we randomly select $x\%$ of the stub AS's and construct a topology containing these stub AS's and their peers, with the peering relationship among them completely preserved. We prune transit AS with too few peers to get the final topology.

Since most IP prefixes are originated by stub AS's, we randomly select

the origin AS's from the stub AS's. According our measurement, 96.14% of MOAS conflicts involve two ASes and 2.7% involve three ASes. Therefore, we only simulate the scenarios where a prefix is originated by one or two AS's. We choose the attackers randomly from all the AS's. Note that, although the attackers may tend to intrude small networks with lower level of security, they can cause more damage if they are located in the transit AS's.

6.2 Experiment 1: Effectiveness of MOAS Community Attribute

In this experiment, we evaluate how effectively our scheme prevents the propagation of false routing information, by comparing the number of routers adopting false routes with and without using the MOAS community attribute. We assume that all the nodes check the MOAS attributes received from their peers and, once they detect a MOAS conflict, they stop the further propagation of a false route (e.g. by checking with DNS as proposed in the paper or using some other mechanism).

We use the topology with 21 nodes (Figure 8(a)). Each simulation is performed with a random selection of N origin AS's and M attacker AS's. It is easy to see that the number of different selections can be rather large for large topologies. Therefore, rather than simulating all the possible selections, we perform 15 runs for a given number of origin AS's and attackers [4]. In other words, each data point is the average of 15 simulation runs.

In Figure 9, X is the pencentage of attackers over all the ASes, and Y is the percentage of AS's(excluding attackers) adopting attackers' routes. As you can see, when the number of attackers increases, more AS's are affected by the false routing information. However, our approach reduces the percentage of AS's adopting the attackers' routes by more than 20%. Even when 57% of the AS's are intruded by the attackers, our approach controls Y to be below 25% and 40% when the prefix is originated by 1 origin AS and 2 origin AS's respecitively.

Furthermore, we obtained similar results using another 21-AS topology with different connectivity structure (see Figure 10).

[4] To get the 15 combinations of orgin AS's and attackers, we first select 3 sets of origin AS's from the stub AS's. Then we select 5 sets of attackers for each set of orgin AS's.

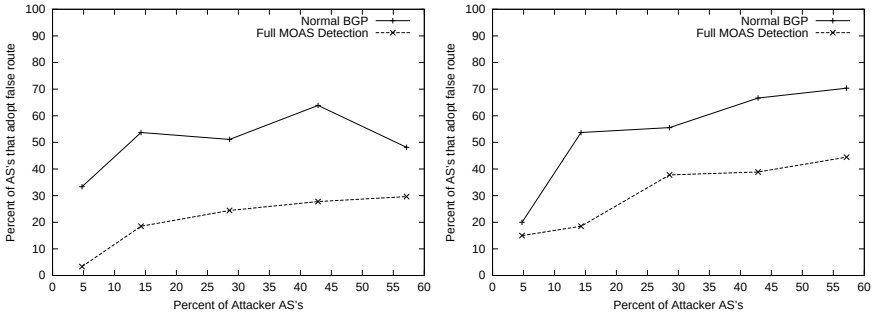

(a) 1 Origin AS (b) 2 Origin AS's

Figure 9: Spoof-Resilience of Our Scheme in the 21-AS Topology

6.3 Experiment 2: Larger Topology

In this experiment, we are interested in finding out if our scheme will still be effective in a larger topology. We also want to compare the results from this experiment with those from the previous experiment to know whether the larger topology is more or less robust against attackers with the same percentage of AS's being intruded by the attackers. The topology we use here is the network with 63 nodes (Figure 8(b)). We have run the experiment with both one origin AS and two origin AS's, but the results are similar as you can see from Figure 11.

One can make the following observations from Figure 11(a):

1. With our scheme, the larger topology is much more robust against random attackers than the smaller topology. For the larger topology, when the attackers are less than 25% of the total number of AS's, almost none of the AS's are affected by the false routing information. Even when about 50% of the AS's are attackers, Y is about 5% for the 63-AS topology, compared to about 30% for the 21-As topology.

2. Without our scheme, the effects of the attackers on the two topologies are quite similar (the gap between the top two curves is much smaller than the gap between the other two curves).

The above results suggest that our scheme is even more effective in larger topologies. We do not have a complete explanation for this phenomenon yet, and it is a topic of our future research.

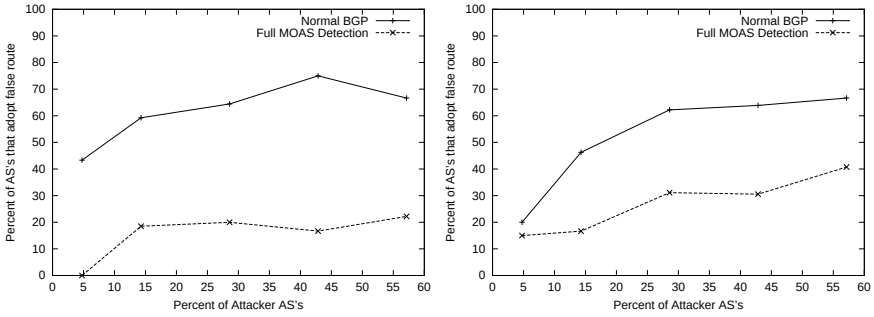

(a) 1 Origin AS

(b) 2 Origin AS's

Figure 10: Results from Another 21-AS Topology

6.4 Experiment 3: Partial Deployment of MOAS Checking Capability

The previous experiments have demonstrated that, with a simple community attribute attached to BGP updates, the Internet routing system can effectively protect itself against maliciously or inadvertently injected routing information by leveraging the distributedness of the Internet topology. Even though our scheme requires only minor modification to the existing routing software, it would be interesting to see if the scheme provides any benefits when only part of the Internet deploys it.

We randomly select 50% of the nodes to have the capability of processing MOAS community attribute. These nodes will be able to detect MOAS conflicts and eliminate false routing announcements. The other nodes will not process the attribute. This means they may install a false route in their routing table if this route is the best path to the prefix and advertise the false route to its peers.

In Figure 12, we compare the effect of partial and complete deployment of MOAS detection. As you can see, although only half of the nodes can detect MOAS conflicts, they can still provide great protection for the other nodes since they prevent the false routes from further propagation. For example, in the 63-AS topology, the effect of partial deployment is as good as full deployment when the attackers intrude less than 25% of the AS's. Again one can observe that the larger topology performs much better than the smaller topology for partial deployment except when more than 40% the AS's are attackers.

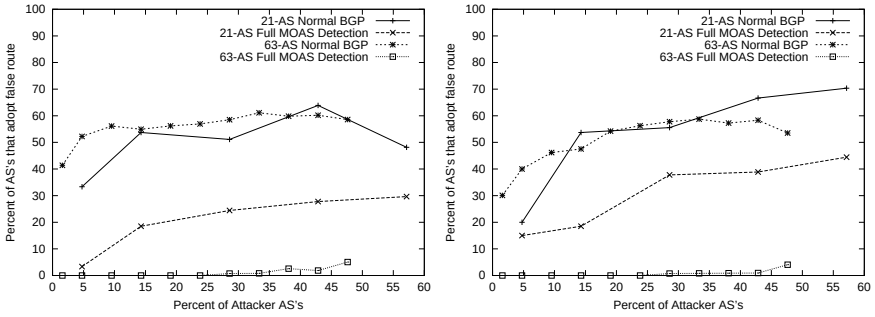

(a) 1 Origin AS

(b) 2 Origin AS's

Figure 11: Comparison between 21-AS and 63-AS Topology

7 Discussion

In this paper, we examined the MOAS conflicts from the archived BGP routing data. The basic analysis was conducted on the number of MOAS conflicts and the duration distribution. The possible reasons for MOAS conflicts are valid operation practice and faults. From the fault tolerance and security point of view, it is necessary to differentiate the valid MOAS conflicts from invalid ones. Thus an assertion based approach is proposed and both a BGP community based and a DNS based solution are suggested.

In the ideal situation, the BGP community based solution works well when all the MOAS Lists from all of origin ASes are observed. An invalid MOAS conflict will be detected and the corresponding faulty origin AS will be identified. However, there is a possibility that MOAS List might be deleted along the propagation path. In this case, it is impossible to tell whether a MOAS conflict is valid or not, even not to mention which origin AS is faulty.

The DNS based solution can guarantee to identify the invalid MOAS conflicts under the assumption that DNS provides the correct information. In addition, the agreement stored in DNS could help to identify the faulty origin AS. But this is not always the case. For example, if a prefix owned solely by a single AS was conflicted by another AS, we can see this as an invalid MOAS conflict, but we could not tell which AS incorrectly originated the prefix.

Because BGP community based solution only requires configuration level modification, as well as self-contained nature of such implementation does

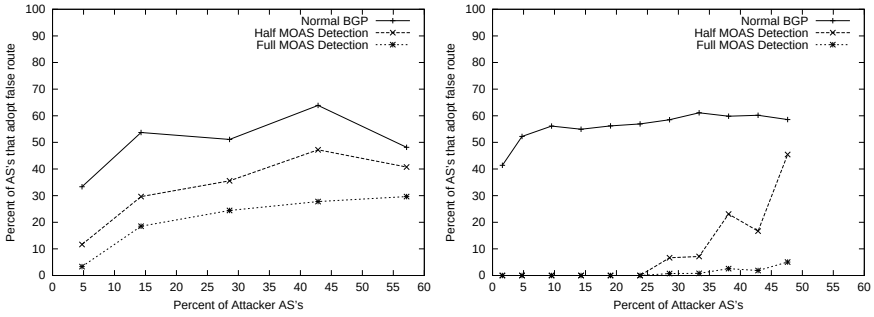

(a) 21-AS Topology (b) 63-AS Topology

Figure 12: Comparison between Partial and Complete Deployment of MOAS Detection

not interfere with other systems, such as DNS, we would like to recommend to deploy such approach at first. Although it has some limitations, we are expecting some of the invalid MOAS conflicts could be detected by this way.

To further enhance the effectiveness of our assertion-based approach, we would like to recommend the deployment of DNS based approach as soon as possible. In the cases that we couldn't identify the faulty origin AS, the further investigation approach should be developed and deployed.

SBGP is considered as a good solution for this problem. If SBGP could be used, by consulting the SBGP certificates database, we will detect the invalid MOAS conflicts as well as the faulty origin AS. However, as we mentioned before, SBGP requires a large scale modification of current Internet infrastructure, so SBGP as a solution may not be immediately available.

It should be noted that the MOAS conflicts are only one part of the general BGP authority problem, which focus on how to verify if an AS has authority to originate a particular prefix. The generalized approach to this problem is our future work.

8 Acknowledgments

We would like to thank a number of network operators and BGP routing experts for the advice. In particular, Randy Bush and Geoff Huston provided useful insights on the operation of large ISPs.

References

[1] J. Hawkinson and T. Bates, "Guidelines for creation, selection, and registration of an Autonomous System (AS)," RFC 1930, 1996.

[2] Y. Rekhter and T. Li, "Border Gateway Protocol 4," RFC 1771, SRI Network Information Center, July 1995.

[3] G. Huston, "BGP table statistics," http://www.telstra.net/ops/bgp/as6447/bgp-multi-orgas.html.

[4] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang, "An Analysis of BGP Multiple Origin AS (MOAS) Conflicts," Submitted to ACM SIGCOMM Internet Measurement Workshop 2001.

[5] T. Bates, R. Bush, T. Li, and Y. Rekhter, "DNS-based NLRI origin AS verification in BGP," Internet Draft, Work in Progress, 1998.

[6] Jessica Yu, "A Route-Filtering Model for Improving Global Internet Routing Robustness," http://www.iops.org/Documents/routing.html.

[7] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol," in *IEEE Journal of Selected Areas in Communications*, Apr. 2000, number 4.

[8] "The Route Views Project," http://www.antc.uoregon.edu/route-views/.

[9] "National Laboratory for Applied Network Research," http://moat.nlanr.net/Routing/rawdata/.

[10] "PCH.net," http://www.pch.net/documents/data/routing-tables/route-views.oregon-ix.net/.

[11] J. Haas, "Autonomous System Number Substitution on Egress," Internet Draft, Work in Progress, 2001.

[12] B. Kroenung, "AS8584 taking over the internet," NANOG Mailing List msg00047, Apr. 1998.

[13] "RIPE Routing Information Service," http://www.ripe.net/ripencc/pub-services/np/ris-index.html.

[14] J. Farrar, "C&W routing instability," NANOG Mailing List msg00209, Apr. 2001.

[15] R. Barrett et al., "Routing Snafu Causes Internet Outage," *ZDNet*, 1997.

[16] V. J. Bono, "7007 Explanation and Apology," NANOG Mailing List msg00444, Apr. 1997.

[17] R. Chandra, P. Traina, and T. Li, "BGP Communities Attribute," RFC 1997, Aug. 1996.

[18] J. Haas, S. Hares, S. Willis, and J. Chu, "Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4)," 2001.

[19] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang, "Validation of Multiple Origin ASes Conflicts through BGP Community Attribute," 2001.

[20] "The SSFNET Project," http://www.ssfnet.org.