

AS-TRUST: A Trust Quantification Scheme for Autonomous Systems in BGP*

Jian Chang, Krishna K. Venkatasubramanian, Andrew G. West, Sampath Kannan, Boon Thau Loo, Oleg Sokolsky, and Insup Lee

Department of Computer and Information Science,
University of Pennsylvania, Philadelphia, PA, 19104
{jianchan, vkris, westand, kannan, boonloo, sokolsky, lee}@cis.upenn.edu

Abstract. The Border Gateway Protocol (BGP) works by frequently exchanging updates that disseminate *reachability information* about IP prefixes (*i.e.*, IP address blocks) between Autonomous Systems (ASes) on the Internet. The ideal operation of BGP relies on three major *behavioral assumptions* (BAs): (1) information contained in the update is legal and correct, (2) a route to a prefix is stable, and (3) the route adheres to the valley free routing policy. The current operation of BGP implicitly trusts all ASes to adhere to these assumptions. However, several documented violations of these assumptions attest to the fact that such an assumption of trust is perilous. This paper presents *AS-TRUST*, a scheme that comprehensively characterizes the trustworthiness of ASes with respect to their adherence of the behavioral assumptions. AS-TRUST quantifies trust using the notion of *AS reputation*. To compute reputation, AS-TRUST analyzes updates received in the past. It then classifies the resulting observations into multiple types of *feedback*. The feedback is used by a *reputation function* that uses Bayesian statistics to compute a probabilistic view of AS trustworthiness. This information can then be used for improving quotidian BGP operation by enabling improved route preference and dampening decision making at the ASes. Our implementation of AS-TRUST scheme using publicly available BGP traces demonstrates: (1) the number of ASes involved in violating the BGP behavioral assumptions is significant, and (2) the proposed reputation mechanism provides multi-fold improvement in the ability of ASes to operate in the presence of BA violations.

1 Introduction

Large IP domains, called *Autonomous Systems* (ASes) use the Border Gateway Protocol (BGP) as the standard communication protocol. BGP enables ASes to exchange IP prefix (*i.e.*, IP address blocks) reachability information with each other through periodic propagation of BGP *update* messages. The *reachability information* within a BGP update consists of IP prefixes, and an ordered list of

* This research was supported in part by ONR MURI N00014-07-1-0907. POC: Insup Lee, lee@cis.upenn.edu

ASes, called *AS_PATH*, through which the prefix is reachable. Additionally, BGP relies on three major *behavioral assumptions* (BAs) to operate: (1) information contained in the update is legal and correct, (2) a route to a prefix is stable, and (3) the route adheres to the valley free routing policy. Violation to any of these behavioral assumptions can have severe consequences for the inter-domain routing. The past decade has seen numerous incidences of BA violations. For instance *prefix hijacking*, when an AS claims to directly reach (*i.e.*, own) a prefix contrary to its actual capability [5] and [10]; *valley route*, which might prevent BGP convergence [2]; and *unstable or potentially spoofed link insertion* in the *AS_PATH* to make the route more attractive [25].

In this paper, we present *AS-TRUST*, a novel scheme for quantifying the level of trust¹ one can have on the ASes based on their adherence to the BAs. This trust quantification has many benefits: (1) obtain a succinct but global view of the current state of inter-domain routing and the extent to which it is plagued by the aforementioned hijacking, stability and policy violation issues, (2) potentially minimize the ASes that violate BAs by making the AS trustworthiness information available to the entire BGP community, and (3) use the information to make informed policy decisions about any new updates received from the ASes.

In AS-TRUST, trust is quantified using the notion of *AS reputation*. This is based on the observation that AS behavior is repetitive. To compute the reputation of an AS, AS-TRUST evaluates past BGP updates received for exhibition of specific behaviors, based on well-defined properties. The behavior evaluation provides feedback to a Bayesian reputation function to generate a probabilistic view of the trustworthiness of all the observable ASes in the Internet. Note that, a low reputation for an AS does not mean it is necessarily malicious. It simply means it has violated one or more of the behavioral assumptions. It could have done it for a variety of reasons including misconfiguration, traffic engineering purposes, as well as malice. We argue that, as it has been shown that violating these individual assumption has consequences for the inter-domain routing space [2] [5] and [25], one needs to be aware of their occurrence. The AS-TRUST reputation values allow us to achieve this sense of awareness. Our implementation of AS-TRUST demonstrates the following: (1) incidents of BA violation are consistently present, (2) a considerable percentage of ASes (5-6%) are involved in some form of BA violation with a handful exhibiting poor behavior exclusively, and (3) the proposed reputation mechanism significantly improves the ability of ASes to operate in the presence of BA violations. To the best of our knowledge, this is the first attempt to quantify the behavioral assumptions of BGP in a systematic manner.

The paper is organized as follows. Section 2 presents background on BGP and the problem statement. Section 3 presents details of AS-TRUST including the notion of BGP service, feedback mechanism, and the reputation function employed. Section 4 presents the properties for evaluating the BGP services. Section

¹ Trust is defined as the competence of an entity to exhibit a specific behavior(s) [16].

5 presents the AS reputation computation and analysis. Section 6 presents the related work, followed by Section 7, which concludes the paper.

2 Preliminaries

2.1 The Border Gateway Protocol

The Border Gateway Protocol is a path-vector routing protocol for exchanging information about reaching IP prefixes. Using BGP, each AS informs its neighbors about the best available route to prefixes it can reach. In this regard, AS sends out a BGP update message *announcing* the prefix. Similarly, an AS can *withdraw* a prefix that it has previously announced. Each AS through which the update passes adds its AS number to the message. This ordered list of ASes called the *AS_PATH* informs an AS receiving the update, the path through which the prefix can be reached. When an update is received by an AS containing a prefix announcement, it has to determine whether it should be accepted or not. Acceptance means that the AS is willing to add the route to its routing information base. Each AS has its own policies that determine whether it accepts a BGP update and whether it forwards it (the update) to its neighbors. Routing policies serve an important purpose in BGP and provide an AS with not only the capability to prefer routes over others to reach a prefix, but also to filter and/or tag an update to change the route’s relative preference downstream.

2.2 Problem Statement and Approach

The current version of BGP [1] was designed with only effectiveness in mind. It implicitly assumes ASes can be trusted to adhere to certain *behavioral assumptions* (BAs). We say that an AS is violating (*i.e.*, not adhering to) the BAs, if it displays any of the following five behaviors:

- *Illegality*: The values of the AS number and the prefixes in the update are from a restricted range, that is the AS numbers are private or the prefixes are bogons.
- *Hijack*: An AS falsely claims to own a prefix in the update. Such false claims on someone else’s prefixes can have adverse consequences including loss of service [10], or can be used for spamming purposes [28].
- *Vacillation*: An update is deemed vacillating if it is part of a quick succession of announcements and withdrawals involving a specific prefix perpetuated by an AS that owns it. Vacillation can cause frequent route-flapping at the upstream ASes, which is detrimental to BGP stability and router performance [?].
- *Valley Route*: The *AS_PATH* of an update has one or more ASes that form a valley. An AS in the *AS_PATH* is said to form a valley if: (1) it forwards a route received from its provider² to another provider, or (2) it forwards a

² ASes and their neighbors usually have one of the four relationships: provider-to-customer (Pv2C), customer-to-provider (C2Pv), sibling-to-sibling (S2S), and (P2P) peer-to-peer[15].

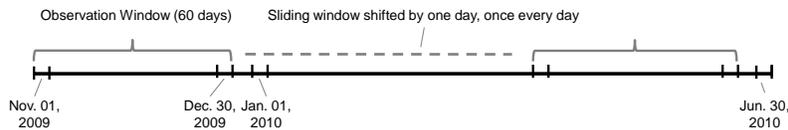


Fig. 1. Data Source Time Windows

route with an existing peer-to-peer link to one of its own peers. Most ASes try to follow a *valley-free routing* (VFR) guideline in their export policy settings [26] as VFR have been shown to be a sufficient condition to ensure that convergence of BGP [15].

- *Unstable AS-Links*: An AS propagates updates through a short-lived AS-link (*i.e.*, a hop between individual ASes in the *AS_PATH*). Detecting such unstable AS-link bindings is important, since ASes which chose a path with one or more unstable AS-links may increase the latency of data delivery, increase the number of BGP updates exchanged within the inter-domain routing space, and may be indicative of link spoofing [23]

The principal question this paper tries to address is “*what is the probability with which an AS adheres and violates the behavioral assumptions of BGP?*” In this regard, we use the notion of reputation. *Reputation* is a quantitative measure of an entity’s likelihood to perform a specific task based on its past behavior [20]. The idea is to compute the reputation for all the ASes in the Internet based on the updates received in the past and analyze them for adherence to the BAs. This is done in four steps: (1) collecting BGP updates in a database; (2) evaluating the data in the database, over a well-defined duration called the *observation window*, for the exhibition of the aforementioned five behaviors; (3) recording the results of the analysis as feedback; and (4) using feedback to compute reputation for the ASes. Reputation is a dynamic value which changes as the AS behavior changes, over time. This is accomplished by repeating the evaluation process over a sliding observation window and generating updated feedback.

2.3 Experiment Setup

We implemented the proposed scheme and conducted a six month long experiment measuring the evolving trustworthiness of ASes, on an Internet-scale. To receive the latest BGP updates, we use the RouteViews BGP trace collector, maintained by University of Oregon [11]. The RouteViews trace collector is a group of BGP routers which peer with a large number of ISPs via BGP sessions. At the time of writing, the RouteViews received BGP updates from 46 ASes. It has been shown in [30] that RouteViews receives updates from almost all the ASes currently active within the Internet and is therefore a good source for computing reputation of ASes. Just as many of the past works in BGP security [27], we assume the RouteViews repository to be trustworthy and provides us with accurate information.

In this work, we use BGP update data from Nov. 1, 2009 - Jun. 28, 2010 (see Figure 1). We take BGP updates received over a 60 day period called the *observation window*, evaluate the AS behavior, and compute reputation for the ASes on the 61st day. For example, data from Nov. 1, 2009 to Dec. 30, 2009 is analyzed to compute AS reputation on Jan. 1, 2010. The observation window is then slid forward by one day and the process is repeated. In order to be fair to the ASes, we did not consider updates announced within 24 hours of the end of the observation window in computing the reputation of the ASes as they have not had enough time to prove themselves. There are over 180 observation windows between Nov 1, 2009 and Jun. 28, 2010. The 60 day observation window was chosen as it was long enough to prevent the behavior evaluation from being biased by transient AS behavior.

3 AS-TRUST Reputation Computation

This section provides an overview of the principal aspects of computing trustworthiness of ASes by AS-TRUST. We begin by formalizing the notion of *BGP service* which forms the basis of the whole process. We then present the mechanism for obtaining feedbacks. In the subsequent sections we describe the evaluation process and reputation computation, respectively.

3.1 BGP Service

The principal task of BGP is to facilitate the dissemination of reachability information through updates. We model this dissemination using a novel notion of *BGP service*. A BGP service is a formal way of viewing reachability information provided collectively by the ASes in the *AS_PATH*, called *providers*³, to an *observer* AS receiving a BGP update. It is defined as: $S_i = \{p_i, AS_PATH = [AS_0, \dots, AS_N]\}$. Here S_i is the service identifier indexed by i , p_i is the prefix being announced by AS_0 as a part of the service S_i , and $AS_0, \dots, AS_N \in AS_PATH$ are the provider ASes which forward the reachability information as a part of the service S_i . Figure 2 (a) illustrates the principal concepts and entities of a BGP service. A service said to have *started* when a provider AS announces a particular prefix and *ended* when the prefix is withdrawn. A service can therefore be in two modes: *active* and *inactive*. A service is said to be active if it has started but not ended; and inactive otherwise. An inactive service has to have been active, at least once, in the past. Each time a service is active, it is called an *instance* of that service. The bottom half of the Figure 2 (a) illustrates some of these concepts over a *time-line* of a service.

A BGP service can be decomposed into three orthogonal *service elements*, each of which are provided by a subset of providers: (1) *AS-prefix binding*: a

³ These are different from notion used in the context of VFR. Here, the term provider ASes mean the provider of a service. In the rest of the paper, unless otherwise specified, the term provider refers to provider ASes.

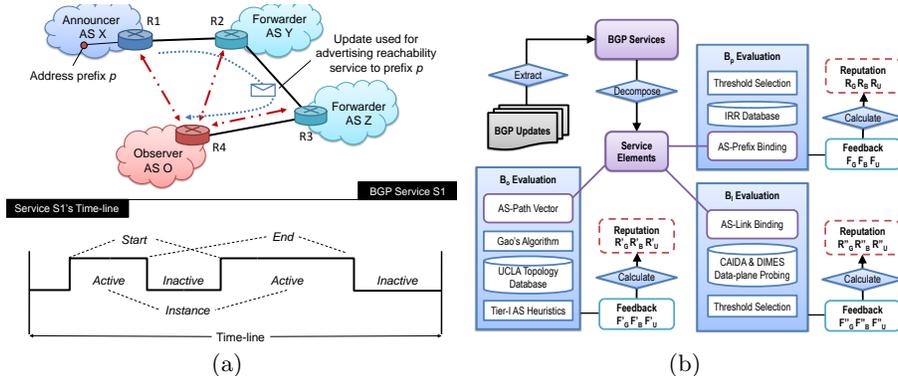


Fig. 2. (a) BGP Service and Timeline, (b) AS-TRUST Service Analysis

tuple of the form (AS_0, p) , which is established when an AS_0 announces a prefix p and is broken when the prefix is withdrawn. Each BGP service has one AS-prefix binding in it. This service element is provided by AS_0 . (2) *AS-path vector*: is synonymous with the AS_PATH in the service. It is said to be provided *collectively* by all the providers; and (3) *AS-link binding*: a tuple of the form (AS_i, AS_j) , which is established when AS_i forwards an update to AS_j . The AS-link binding is broken when no service uses it. A service has $N - 1$ AS-link bindings; one between each of the N ASes in the AS_PATH . This service element is said to be provided *individually* by all the providers in the AS_PATH to the observer. In the rest of the paper, we use the term AS-link bindings and AS-links, interchangeably.

Upon observing a BGP service, the observer decomposes it into its constituent service elements, each of which is then evaluated on its validity. The results of the evaluation act as a feedback on the providers of the service element. The next sub-section details how the behaviors described in Section 2.2 can be evaluated for the service elements, followed by the feedback mechanism used. *Note that, as the feedback is generated locally, we do not have to consider the case of potentially dishonest external feedback affecting our reputation computation outcome.*

3.2 Behavior Evaluation

We propose three behavior sets, one corresponding to each service element, for behavior evaluation. The three *behavior sets* comprehensively cover the behavioral assumptions on which BGP operates. They are: (1) *Prefix Behavior Set* (B_p): Requires that AS in the AS-prefix binding service element *does not* exhibit prefix value illegality, hijacking or vacillating behavior; (2) *Path Behavior Set* (B_o): Requires that *none* of the ASes in the AS-path vector service element form a valley or exhibit AS number illegality⁴; (3) *Link Behavior Set* (B_l): Re-

⁴ A 16-bit AS number is *illegal* if its value is in the range of 64496-64511, which is reserved for use in documentation and sample code, 64512-65534, which is designated for private use, or 65535, which is reserved [6].

quires that *none* of the ASes create an AS-link binding service element that is short-lived.

It can be seen that there is a one-to-one mapping between the service elements and the behavior sets. Therefore, evaluating a service involves evaluating whether AS-prefix binding, AS-path vector, and AS-link binding service elements satisfy B_p , B_o , and B_l , respectively. However, before we delve into the details of evaluation, we provide an overview of our feedback mechanism, which is essential for reputation computation, and forms an integral part of the evaluation process.

3.3 Feedback Mechanism

Evaluation of a BGP service element provides one of three mutually exclusive feedbacks. The feedback can have one of three values: (1) *Feedback G*: this feedback is given on the providers that satisfies the requirements of the appropriate behavior set; (2) *Feedback B*: this feedback is given on the providers that do *not* satisfy the requirements of the behavior sets, however, they do not disrupt BGP operation; and (3) *Feedback U*: this feedback is given on the providers that not only violate the requirements of the behavior sets, but also disrupt BGP operation.

In the rest of the paper, we use the term *GBU feedbacks* to refer to our feedback types. When the service element implemented by the provider(s) receives *Feedback G*, it is referred to as *good behavior*. Conversely, a *Feedback B* or a *Feedback U* for an AS is referred to as the demonstration of *poor behavior*. Essentially, a good behavior adheres to the behavioral assumptions, while poor behaviors violates them. In general, there exists a 3×3 feedback matrix for every provider AS_a , at the observer, of the form:

$$F_a = \begin{pmatrix} F_G & F_B & F_U \\ F'_G & F'_B & F'_U \\ F''_G & F''_B & F''_U \end{pmatrix}$$

where the element $F_a(1, j)$, $F_a(2, j)$, and $F_a(3, j)$ stores the details of the BGP service, which AS a provided when evaluated with respect to B_p , B_o , and B_l , respectively. Finally, as the feedback are generated locally at the observer AS, we do not have to consider the case of potentially dishonest feedback affecting our reputation computation outcome.

4 BGP Service Evaluation and Feedback

In this section, we describe the metrics used in the behavior evaluation of service elements. These metrics allow the feedback matrix to be populated, which will subsequently be used to compute reputation. As mentioned earlier, the behavior evaluation considers BGP services received during a 60 day observation window and produces feedback on the providers. Figure 2 (b) illustrates the work-flow of the evaluation process discussed in this section. The boxes with dashed outlines illustrate the output produced at the end of analyzing a service based on each of three behavior sets.

Table 1. Feedback for Behavior Evaluation based on B_p

Prevalence	Persistence	Classification	Feedback
high	high	Good	F_G
high	low	Vacillation	F_B
low	high	Good	F_G
low	low	Hijack	F_U

4.1 Evaluation of Service using B_p

Determining whether an AS-prefix binding (AS_0, p) exhibits B_p builds on [13]. Therefore, in the rest of the section we briefly summarize the metrics used and evaluation described therein. The evaluation is a three step process:

- *Stability Analysis:* For each (AS_0, p) observed during the observation window, we compute two temporal metrics: persistence and prevalence⁵. *Prevalence* (Ps), is the total percentage of time an AS-prefix binding is active within the observation window. *Persistence* (Pr), on the other hand, is the average duration of time an AS-prefix binding is active, within the observation window.
- *Providing Feedback:* The value of the Ps and Pr are compared against a set of thresholds T_{pr} (1% of the observation window) and T_{Ps} (10 hours)⁶ and feedback provided. Table 1 shows the feedback matrix element updated for different Pr and Ps values.
- *Detecting Bogons:* (AS_0, p) is also statically checked for the presence of bogons, and their discovery results in F_U element being updated in the feedback matrix F_a associated with AS_0 .

The case of Pr being high and Ps being low demonstrates a *vacillating* nature of an AS-prefix binding. Detailed analysis of such bindings demonstrate that they are usually legitimate [13]. However, the AS-prefix binding service element itself vacillates between being active and inactive at a rate which is not conducive for data communication. Further, it causes significant increase in the number of updates exchanged to manage the changes causing frequent route flapping [13]. Consequently, we give such vacillating behavior *Feedback B* because the ASes execute BGP’s functionality correctly but fail to meet the requirement of the behavior set. As for bogons, we believe their announcement subverts the operation of BGP and we therefore give them *Feedback U*.

The results of the evaluation, based on B_p , are summarized in Figure 3. An average of 421704.1 AS-prefix bindings were observed every observation window, out of which an average of 4.0% were found to be hijacked⁷ involving 1.7% of all the ASes. Similarly, about 6.9% of AS-prefix bindings were classified as vacillating, involving 3.1% of all the ASes. The number of ASes displaying exclusively

⁵ The principle idea of evaluating temporal characteristics comes from the observation that legitimate AS-prefix pairs last long periods of time [21].

⁶ Both the thresholds have been established empirically, based on lowest false positive and false negative rates when compared with Internet Route Registries (IRR) [13].

⁷ This number is unusually high due to the Internet-scale prefix hijacking attempt on April 8th, 2010 by AS_{23724} .

Behavior Evaluation Results from Jan. 1, 2010 and Jun. 30, 2010

(For each day, the analysis considers: I. BGP updates from the past 60 days for B_p and B_o , and II. AS relationship annotated topologies of the past 24 hours for B_o)

Analysis of B_p^*		Analysis of B_o	
Property	Value	Property	Value
Avg. # of ASPB** Observed	421704.1	Avg. # of Paths	661395.7
Avg. # of ASPB Provide Feedback U	6955.61***	Avg. # of Valley Routes	3447.8
Avg. # of ASPB Provide Feedback B	29256.6	Avg. # of AS Creating Valley Routes	89.2
		Avg. # of BGP Services Containing Illegal AS Number	44.1

Analysis of B_i		Analysis of B_o	
Property	Value	Property	Value
Avg. # of AS-links Observed	94754.2	Avg. # of Unique AS Observed	35667.2
Avg. # of Stable AS-links	91143.6	Avg. # of AS Using Unstable AS-links	1945.7
Avg. # of Unstable AS-links	3610.6	Avg. # of AS Only Using Unstable AS-links	67.4

* No bogons were observed during the experiment periods
 ** ASPB: AS-prefix bindings
 *** The actual value was higher, due to the Internet scale prefix hijacking mounted by AS23724 on Apr. 8, 2010.

Fig. 3. AS Behavior Evaluation Statistics

poor behaviors is lower still. Finally, we observed zero occurrence of AS-prefix bindings with bogon prefixes during any of the observation windows. We believe this is because bogons are invariably filtered out by ASes that detect them. *The results demonstrate that a relatively large number of ASes (3-5%) are involved in announcing vacillating and hijacked prefixes.*

4.2 Evaluation of Service using B_o

To evaluate an AS-path vector based on B_o is a four step process: (1) *Generating AS Relationship Map*: We download that day’s annotated topology from the UCLA’s Internet topology site [30] and merge it with a topology inferred by applying Gao’s algorithm [15] to the previous day’s RouteViews data; (2) *Introduce Peers*: We obtain the list of all tier-1 ASes from [2]. All links between tier-1 ASes are re-labeled peer-to-peer (P2P), and links between tier-1 AS and lower-tier AS are re-labeled Pv2C where the tier-1 AS is the provider (Pv); (3) *Providing Feedback*: Once the merged annotated topology has been created, the AS-path service element of all the services announced that day is evaluated for the existence of ASes which might violate VFR. If such an AS is found, then its F'_B entry in its feedback matrix is updated; and (4) *Identifying Illegal ASes*: The AS_PATH is finally examined for illegal AS numbers. This is done based on a static check. The first legal AS, after the set of illegal ones is blamed and its F'_B updated.

The use of two well-known AS topology relationship inference techniques increases the confidence of our own relationship labeling. The violation of VFR is given *Feedback B* because, though not good in the long run, it does not necessarily affect the operation of BGP in providing knowledge about routes to prefixes. In the case of illegal ASes, the first legal AS after a set of private ASes is blamed because such leaking of private numbers usually happens when an AS forgets to filter out local AS numbers before forwarding the update [23].

The results of the evaluation, based on B_o , are summarized in Figure 3. We found that, an average of 661395.7 paths were observed per day. Out of these, 0.5% paths were found to violate VFR per day. Finally, an average of 89.2 providers out of over 35K were seen violating VFR per day. On average, we found only about 44.1 ASes involved in allowing illegal AS numbers in the AS_PATH , per day, during the six months of behavior analysis with respect to

B_o . In summary, the violation to B_o , especially valley routes are prevalent and a recurring event in the day to day operation of BGP.

Table 2. Feedback for Behavior Evaluation based on B_l

Prevalence	Persistence	Classification	Feedback
high	high	Good	F''_G
high	low	Good	F''_G
low	high	Good	F''_G
low	low	Unstable	F''_U

4.3 Evaluation of Service using B_l

Computing the stability of an AS-link binding (AS_i, AS_j) in the AS_PATH follows a similar approach to AS-prefix binding stability evaluation and uses the *prevalence* and *persistence* metrics. The evaluation and feedback with respect to B_l is done in three steps: (1) *Identifying AS-link Bindings*: This generates a set L of all the AS-link bindings, decomposed from the services observed during the observation window; (2) *Computing Stability Metrics*: This step computes the prevalence (Pr) and persistence (Ps) for each of the AS-link in set L ; (3) *Providing Feedback*: The computed Pr and Ps values are then compared with a threshold Tl_{Pr} and Tl_{Ps} and a feedback is provided. Table 2 shows the feedback matrix element updated for different Pr and Ps values.

The reason we give unstable AS-links the *Feedback U* because it is possible that poor AS-link stability is due to an attempted link spoofing which could subvert the intended BGP operation. It should be noted that it is difficult to get conclusive proof for the spoofing given a lack of ground truth, though we find an interesting result which strengthens the case for their occurrence (see Section 5.2). We therefore argue that the potential of spoofing merits a punitive feedback for ASes involved in unstable AS-links. The value of thresholds Tl_{Pr} and Tl_{Ps} are set to 1% of the observation window and one hour, respectively. These values are established empirically based on comparison with a set of AS-links D . The set D is obtained from data-plane probing database provided by the CAIDA [3] and DIMES [4] projects. The thresholds are the values below which, all the AS-links in the set L , with the particular Pr and Ps, have the smallest intersection with the set D . Data-plane probing is used because if a AS-link is ephemeral, it has a low probability of being found in data-plane probing. Further, it is the only form of ground-truth available that can reliably identify AS-links stable enough to allow data traffic to pass through them [24].

Figure 3 shows the results of analysis, based on B_l . An average of 95640.4 AS-links were observed during the each of the observation windows. Out of these over 96.1% AS-links received *Feedback G*. From the perspective of the ASes, on average of 35667 ASes were seen every observation window, out of which 5.4% ASes announced unstable AS-links at least once. Only about an average of 0.18% of ASes announced purely unstable AS-links. Figure 4(a) visualizes 4625 unstable AS-links seen each month over the course of the experiment involving 2305 ASes. The dots are the ASes and the lines between them are unstable AS-links. The

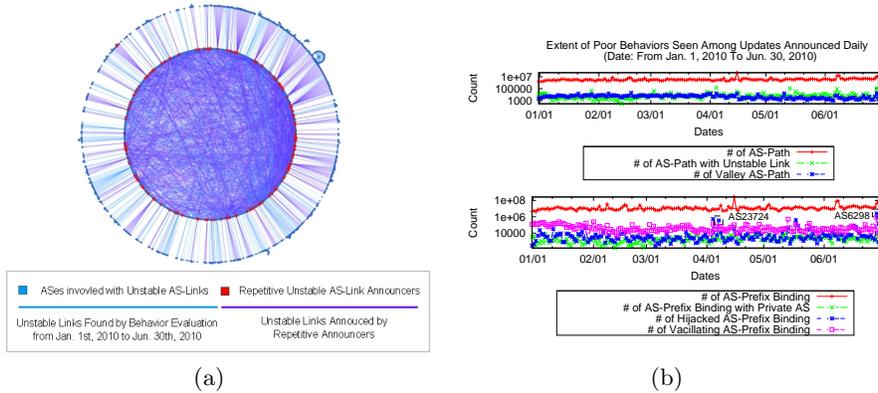


Fig. 4. (a) Visualization of Unstable AS-Links and the Provider ASes Involved, (b) Extent of Poor Behaviors in the Internet

red dots represent the 149 ASes which have established unstable AS-links at least once every month, during the course of our experiment. *In summary, these results demonstrate that unstable AS-links are a repetitive phenomena, affecting a substantial number (5-6%) of ASes.*

Figure 4(b) presents the extent of poor behavior seen every day between Jan. 1, 2010 and Jun. 30, 2010. The trend graphs provide an overview of the extent the poor behaviors of ASes afflicting inter-domain routing and how they have evolved over time. The numbers are raw-values and include all AS-prefix bindings, *AS_PATHS* and AS-links, decomposed from the observed BGP services. Overall, the problem of poor behavior is consistently present over the course of the six months and is largely stable in its intensity, with occasional spikes. These results do indicate the importance of monitoring AS-prefix binding vacillation, AS-link stability and presence of valley route with the same diligence as prefix hijacking.

5 Reputation Computation

At the end of behavior evaluation, we have a feedback matrix for each AS. This will now be used to compute reputation. The reputation will allow the observer AS to know *the probability of a service element in a BGP service provided by an AS, being given Feedback G (or Feedback B or Feedback U)*. Given the service elements are orthogonal to each other, the reputation for an AS a is computed as a 3×3 matrix (just like F_a).

$$R_a = \begin{pmatrix} R_G & R_B & R_U \\ R'_G & R'_B & R'_U \\ R''_G & R''_B & R''_U \end{pmatrix}$$

Here, the rows correspond to the reputation of the AS with respect to an AS-prefix binding, AS-path, and AS-link binding service elements, respectively. We do not arrive at a single number for reputation here, as it would not be able to describe the behavior of an AS with the same level of detail.

5.1 Reputation Computation

To calculate the reputation we use Bayesian statistics. Intuitively, if we have random events with k possible outcomes, we can compute the posteriori probabilities, *i.e.*, the probabilities of observing these behaviors, using the Dirichlet distribution if we assume the prior to be a Dirichlet distribution as well [8]. Then, the reputation is the expected value of a posteriori probability distribution. The reputation model presented here is a generalization of [20].

An important property of the Internet is that poor behaviors usually have very short duration [13] [25]. Reputation value is designed to determine *the expected probability of a service element in active service provided by an AS being given Feedback G*. Consequently, the reputation is calculated by weighing the entries in the feedback set to a value proportional to the time the service remained active within an observation window. In other words, after M trials, let $|F_X|$ be the count of BGP services contained in F_X , where $X \in \{G, B, U\}$, and $|F_G| + |F_B| + |F_U| = M$. Then, we define $|F_X| = \sum_{i=1}^k t(s_i)$, where k is the total number of BGP services in F_X and $t(s_i)$ is the percentage of time a BGP service s_i in F_X is active within the observation window. As good behaviors last a long duration compared to poor ones, at any given time within our observation window, the probability of an active service having good service elements will be much higher than probability of an active service with poor service elements.

Example: Let an AS0 provide four services $S1, S2, S3, S4$ where $t(S1) = 0.95$, $t(S2) = 0.70$, $t(S3) = 0.05$ and $t(S4) = 0.40$, respectively. After evaluating the services based on B_p , let $S1$ and $S2$ be give the *Feedback G*, $S3$ is given *Feedback B* and $S4$ is given *Feedback U*. As each service element in a service is independent of the others, $t(S1) + t(S2)$ may be greater than 1. The reputation of ASes is given by: $|F_G| = 0.95 + 0.70 = 1.65$, $|F_B| = 0.05$, $|F_U| = 0.4$, resulting in $R_G = (1.65 + 1)/(1.65 + 0.05 + 0.4 + 3) = 0.51$, $R_B = (0.05 + 1)/(1.65 + 0.05 + 0.4 + 3) = 0.21$, and $R_U = (0.4 + 1)/(1.65 + 0.05 + 0.4 + 3) = 0.28$. The value one added encodes the prior observation. We can compute reputation values corresponding to other two behavior sets in a similar manner. It can be seen that the reputation function redistributes the probability of poor behavior in a manner proportional to the duration for which the service was active.

5.2 Reputation Analysis

In this section, we analyze the reputation of ASes, generated over a period of six months from Jan. 1, 2010 - Jun. 30, 2010. We focus on presenting only the results of the reputation due to poor behaviors. The reputation due to good behavior is a complement of the results and can be easily extracted from these. Figure 5 (a) shows the CDF of reputations of ASes which have at least one *Feedback B* or *Feedback U* for B_p, B_o and B_l . As reputation of ASes is computed every day, we illustrate the CDFs for a sampling of six days during the six month period. These graphs demonstrate three important points: (1) among the ASes that do demonstrate poor behaviors, over 85% of them do so infrequently (in

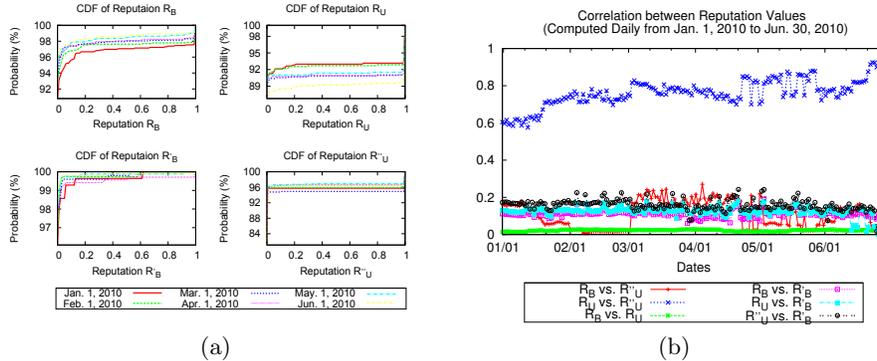


Fig. 5. (a) CDF of ASes with Poor Reputation Values, (b) Correlation between Various Reputation Values

the case of B_o this number is over 99%); (2) about 2-8% of the ASes which demonstrate poor behavior do so exclusively (*i.e.*, the spike near the very end of the distribution) for B_p and B_i ; and (3) over 99% of the ASes have a reputation close to zero for B_o , which means they are rarely involved in valley routes. This result demonstrates the *sensitivity of the reputation metric* as it is able to capture even those ASes which seldom violate BAs.

The availability of a quantitative value for different aspects of AS trustworthiness in the form of reputation allows us to mine emergent AS behavior trends that were heretofore difficult to identify. Figure 5 (b) illustrates the results of the correlation between different elements of the reputation matrix. Interestingly, we find R_U (*i.e.*, prefix hijacking) and R'_U (*i.e.*, unstable AS-links) have a very high correlation, over the six months of our experiments. This is very intriguing as it increases the potential for low stability AS-links to be malicious (spoofed links). None of the other reputation values are strongly correlated.

One of the ways of using the AS reputation information is to improve the routing policies at ASes in order to minimize the effects of BA violations. In this regard, we built a BGP simulator that uses AS reputation information to make policy decisions about route-preference and route-dampening (*i.e.*, a process that prevents routers from thrashing while trying to incorporate a large number of route updates, thus producing a more stable routing table). The idea is to demonstrate that reputation information improves the number of quality routes added to the routing tables of an AS compared to reputation-less scheme. We find that, the reputation-based policy is particularly effective for reducing hijacked entries in the routing table, with an average 13 fold reduction compared to non-reputation case. Additionally, the reputation-based policy achieves two and four fold reduction in adding valley routes and routes with unstable-links, respectively. For more details on the simulation please refer to Section 6 in the technical report [?].

6 Related Work

Little work has been done with respect to characterizing AS behaviors. Most of the work has focused on detecting prefix hijacking using control-plane [21] [25] or data-plane probing [18] [32]. In [13], we present a preliminary version of this system called AS-CRED. AS-CRED, computes reputations for ASes based on their tendencies to hijack or announce short-lived prefixes as in this work. AS-TRUST, on the other hand, considers many more aspects in reputation computation including valley free routing and AS-link stability. However, the principal difference between the two is in the semantics of the reputation value. For example, in AS-TRUST R_U indicates how many prefix hijacking an AS performed. This value is indifferent to the number of stable AS-prefix bindings the AS had. For AS-CRED, reputation is simply a statement of how many prefix hijacking an AS has mounted. AS-CRED reputations can therefore be compared with the reputations of other ASes to see how they fare in comparison. With AS-TRUST, each row in the reputation matrix is a normalized value and has a probabilistic meaning. Therefore, poor behaviors of an AS cannot be seen independently of its good behaviors. We believe, reputation values of AS-TRUST thus provide a complimentary view of the AS behavior compared to the one provided by AS-CRED.

7 Conclusions

In this paper we presented AS-TRUST, a reputation-based trust quantification scheme with respect to the adherence of an AS to the three behavioral assumptions of BGP. Reputation is computed by evaluating past updates announced by each observable AS in the Internet for the exhibition of specific behaviors. The evaluation utilizes well-defined properties for this purpose including the presence of stable AS-prefix binding, stable AS-links, and valley free routes. It then classifies the resulting observations into multiple types of feedback. The feedback values are input into a reputation function that provides a probabilistic view of trust. Analysis of AS-TRUST shows that the incidents of assumption violation is consistently present, and that the proposed reputation mechanism can significantly improves the ability of ASes to function in the presence of violations of behavioral assumptions. In the future, we plan to study the effectiveness of other possible ways of using AS reputation to improve AS' policies.

References

1. A Border Gateway Protocol 4 (BGP-4) RFC. <http://www.rfc-editor.org/rfc/rfc4271.txt>.
2. BGP Routing Leak Detection System Routing Leak Detection System. <http://puck.nether.net/bgp/leakinfo.cgi>.
3. Macroscopic Topology Measurements . <http://www.caida.org/projects/macroscopic/>.
4. The DIMES project. <http://www.netdimes.org/new/>.

5. 7007 Explanation and Apology. <http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html/>.
6. Autonomous System (AS) Numbers. <http://www.iana.org/assignments/as-numbers/>.
7. Chinese ISP hijacks the Internet. <http://bgpmon.net/blog/?p=282>.
8. Dirichlet distribution. <http://www.cis.hut.fi/ahonkela/dippa/node95.html>.
9. ListWare: BGP Update Report. <http://www.listware.net/201007/nanog/6379-bgp-update-report.html>.
10. Pakistan hijacks YouTube. http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml/.
11. RouteViews. <http://www.routeviews.org/>.
12. P. Boothe, J. Hiebert, and R. Bush. Short-lived prefix hijacking on the Internet. In *In Proc. of the NANOG 36*, February 2006.
13. J. Chang, K. Venkatasubramanian, A. G. West, S. Kannan, I. Lee, B. Loo, and O. Sokolsky. AS-CRED: Reputation service for trustworthy inter-domain routing. In *University of Pennsylvania Technical Report, MS-CIS-10-17*, April 2010.
14. N. Feamster, H. Balakrishnan, and J. Rexford. Some foundational problems in interdomain routing. In *In HotNets, 2004*, pages 41–46, 2004.
15. L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Trans. Netw.*, 9(6):733–745, 2001.
16. T. Grandison and M. Sloman. A survey of trust in Internet applications. *IEEE Communications Surveys and Tutorials*, 3(4), August 2000.
17. N. Hu, P. Zhu, and P. Zou. Reputation mechanism for inter-domain routing security management. In *In Proc. of the 9th International Conference on Computer and Information Technology*, pages 98–103, October 2009.
18. X. Hu and Z. M. Mao. Accurate real-time identification of IP prefix hijacking. In *SP '07: Proceedings of the 2007 IEEE Symposium on Security and Privacy*, pages 3–17, Washington, DC, USA, 2007. IEEE Computer Society.
19. G. Huston. *Interconnection, Peering and Settlements*, 2(1):2–16, March 1999.
20. A. Josang and R. Ismail. The beta reputation system. In *In Proceedings of the 15th Bled Electronic Commerce Conference*, 2002.
21. J. Karlin, S. Forrest, and J. Rexford. Autonomous security for autonomous systems. *Comput. Netw.*, 52(15):2908–2923, 2008.
22. M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang. PHAS: a prefix hijack alert system. In *In Proc. of the 15th conference on USENIX Security Symposium*, Berkeley, CA, USA, 2006. USENIX Association.
23. R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP misconfiguration. In *In Proc. of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 3–16, 2002.
24. M. Nicholes and B. Mukherjee. A survey of security techniques for the Border Gateway Protocol (BGP). *IEEE Communications Surveys and Tutorials*, 11(1), First Quarter 2009.
25. J. Qiu, L. Gao, S. Ranjan, and A. Nucci. Detecting bogus BGP route information: Going beyond prefix hijacking. In *Security and Privacy in Communications Networks and the Workshops, 2007. SecureComm 2007. Third International Conference on*, pages 381–390, Sept. 2007.
26. S. Qiu, P. McDaniel, and F. Monrose. Toward valley-free inter-domain routing. In *Communications, 2007. ICC '07. IEEE International Conference on*, pages 2009–2016, 2007.
27. T. Qiu, L. Ji, D. Pei, J. Wang, J. Xu, and H. Ballani. Locating prefix hijackers using LOCK. In *18th USENIX Security Symposium*, Aug. 2009.

28. A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. *SIGCOMM Computation and Communication Review*, 36(4):291–302, 2006.
29. H. Yu, J. Rexford, and E. Felten. A distributed reputation approach to cooperative Internet routing protection. In *Secure Network Protocols, 2005. (NPSec). 1st IEEE ICNP Workshop on*, pages 73–78, Nov. 2005.
30. B. Zhang, R. Liu, D. Massey, and L. Zhang. Collecting the Internet AS-level topology. *SIGCOMM Comput. Commun. Rev.*, 35(1):53–61, 2005.
31. Z. Zhang, Y. Zhang, Y. C. Hu, and Z. M. Mao. Practical defenses against BGP prefix hijacking. In *CoNEXT '07: Proceedings of the 2007 ACM CoNEXT conference*, pages 1–12, 2007.
32. Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush. iSPY: detecting IP prefix hijacking on my own. *SIGCOMM Comput. Commun. Rev.*, 38(4):327–338, 2008.
33. C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis. A light-weight distributed scheme for detecting IP prefix hijacks in real-time. *SIGCOMM Comput. Commun. Rev.*, 37(4):277–288, 2007.