

A Distributed Reputation Approach to Cooperative Internet Routing Protection

Harlan Yu
Princeton University
harlanyu@cs.princeton.edu

Jennifer Rexford
Princeton University
jrex@cs.princeton.edu

Edward W. Felten
Princeton University
felten@cs.princeton.edu

Abstract

The security of the Internet’s interdomain routing system hinges on whether autonomous systems (ASes) can trust the information they receive from each other via the Border Gateway Protocol (BGP). Frequently, this trust has been misguided, resulting in wide-spread outages and significant concerns about future attacks. Despite the seriousness of these problems, proposals for a more secure version of BGP have been stymied by serious impediments to practical deployment. Instead, we argue that the existing trust relationships between network operators (and the institutions they represent) are a powerful force for improving the security of BGP, without changing the underlying routing protocol. Our approach leverages ideas from online reputation systems to allow ASes to form a peer-to-peer overlay that integrates results from local network-management tools for detecting attacks and configuration errors. The proposed architecture is incrementally deployable, protects against shilling attacks, and deters malicious operator behavior.

1. Introduction

The Border Gateway Protocol [21] is the Internet’s de facto interdomain routing protocol. The veracity of BGP routing information, passed as messages among ASes, is vital to the proper functioning of the Internet. BGP provides no intrinsic facility for an AS to validate the truthfulness of a received path update message— an AS must blindly trust the legitimacy of each route it receives. Illegitimate routing information can be propagated both by unintentional routing misconfigurations [17] and attacks by a malicious adversary [20]. Both types often cause traffic instability and reachability problems for at least a small subset of ASes on the Internet. BGP failures can even cause severe widespread outage as witnessed by the infamous AS7007 incident [10], when a flood of er-

roneous route advertisements caused global connectivity problems. BGP attacks can also force traffic redirection, interception, or modification in the case of a malevolent eavesdropper.

In order to defend against fallacious BGP update messages, the research community has focused on two broad classes of possible solutions. On one hand, researchers have proposed modifications to the actual BGP protocol and the addition of a centralized public key infrastructure (PKI) or a routing registry [14] [15] [19]. Adoption of these schemes have thus far failed for a variety of reasons: the potential for out-of-date central databases, the inability for incremental deployment or the prohibitive overhead of expensive cryptographic operations. On the other hand, researchers have suggested a multitude of analytic tools to mitigate the effects of erroneous updates locally at each AS [12] [22] [24]. While many have proven useful, many detect only specific types of misconfigurations and attacks. These locally-deployed tools also do not take advantage of AS coordination to resolve interdomain routing problems from multiple vantage points. Others have recognized the need for real-time inter-AS collaboration [11] [4] [18] but have either restricted adoption using a centralized “members only” paradigm or have focused only on tactics to mitigate a particular class of attacks.

In this paper, we contend that network operators can capitalize on existing trust relationships to better address any BGP routing fault by sharing existing results from local debugging tools. We propose a novel distributed reputation protocol that is particularly well-suited for inter-AS cooperation. The key idea is to mimic real-world trust relationships in a peer-to-peer (P2P) social network and weigh the information gathered from more trusted colleagues. Our approach requires no modifications to the BGP protocol and does not call for a centralized PKI or registry as in previous proposals. We demonstrate the system’s potential to complement existing network management techniques and amplify the detection of both misconfigurations and attacks.

2. Distributed Reputation

In this section, we first describe the challenges posed by online reputation systems in general and propose a novel solution of building real-world trust relationships into distributed online environments. We then discuss specifically the protocol's advantages and limitations in a BGP-specific context.

2.1. Essential Properties

When people look for information about unknown entities, they generally seek the opinions and recommendations of those who they regard to have good reputations. This works quite well in the real world— authentication is no harder than recognizing one's voice or face, and reputations are persistent and permanent. Major problems arise when we apply the concept of reputation to the online world. The Internet enables people to have ephemeral identities, pseudo-anonymous communications, and little accountability for their actions. Those who have built up bad reputations in an online community can simply shed this identity and return with a clean slate using a new pseudonym.

Trust in the interdomain routing context is based on the personal relationships between human network operators who manage ASes. Especially among backbone and other large transit ASes, many network operators have already developed trust relationships through personal and repeated contact at meetings of the North American Network Operators' Group (NANOG) [7] or similar groups in other geographic regions (*i.e.*, AfNOG [1] and SANOG [5]). Additionally, trust between ASes can be formulated through successful business relationships or known adherence to best common security practices [6]. An example is the Ivy Plus Administrative Computing group, a community of network operators from university campuses with an adopted set of group practices [2]. Network operator communities are suitable for online reputation systems since the barrier to entry is high— it is difficult for a notorious operator to shed a bad reputation and start afresh.

In general, any useful online reputation scheme requires at least the following four properties:

1. **A foundation of trust:** Any reliable reputation system must build upon repeated interactions between any two principals in the system. A principal must be able to assess its past interactions with other individuals in the system and be able to act accordingly to those judgments.
2. **Carrots and sticks:** To be useful, a system must provide strong disincentives to deter malicious be-

havior and levy effective punishments appropriately. Analogously, a system can also encourage and reward trustworthy behavior.

3. **Robustness against shilling:** The system's architecture must be robust against shilling attacks by dishonest principals. False ratings injected by an adversary must either have minimal affect on the system or be easily detectable. It must also be difficult for an existing principal to drop a bad reputation and start anew. This limits the number of repeated "hit and run" attacks by the same principal.
4. **An accurate scoring system:** Ratings distributed to and calculated by principals in the system must be both understandable and verifiable. Users must be able to accurately interpret the ratings in order to act knowledgeably on the recommendation. The ratings might also be weighted towards the most recent indicators of behavior.

2.2. A Trust-based Overlay Network

The essence of the distributed reputation architecture is to construct a P2P overlay topology that mirrors existing real-world trust relationships. This overlay network will then be used to implement distributed voting where peers vouch for the truthfulness of boolean propositions. We make the assumption that each AS is represented by a single node in the network.

A logical link in the overlay network is constructed between two nodes if one AS has an offline trust relationship with another AS. As such, nodes must swap authentication information through some form of out-of-band communication to establish the link. Authentication can come in many forms depending on the specific implementation of the protocol. For example, network operators might swap identifying information in person at NANOG key signing parties [3], over the phone or using e-mail. Note that the logical links are not limited to the set of immediate BGP neighbors, but can reach far across the physical network. The links are also unidirectional since the trust might not be reciprocated. The result will effectively be an online social network, where neighbors in the overlay are trusted acquaintances in real life: a node will trust its direct (one-hop) neighbors the most, the neighbors of those nodes (two hops away) less than its direct neighbors but more than others, and so forth.

Formalizing this idea, we define a variable weight factor α , where $0 < \alpha < 1$, that geometrically decreases the level of trust as the relationship between two nodes grows distant. We will show in the next section how to

use this weight factor to share trustworthy information through distributed voting.

2.3. Distributed Voting

Having established the trust-based overlay, nodes will initiate queries in the form of a boolean bit-string, also called a proposition. For example, a node can broadcast the simple proposition “Is $AS3 \rightarrow AS7$ a spoofed edge?” and solicit responses from other peers in the network. A response is simply a vote cast in ternary fashion $(-1, 0, +1)$: -1 indicates that the node believes the proposition to be false; $+1$ for true; and 0 means that the node is either neutral to the proposition or doesn’t have enough information to make a judgment.

As votes are cast, each node in the system will iteratively recompute its overall vote for each proposition based on both its own judgment and the votes of all other nodes in the system. To demonstrate this calculation, we consider a single node N in the network, with weight factor α , for a proposition P . Node N will cast its own vote, a ternary choice V_N , based on its local analysis of edge $AS3 \rightarrow AS7$. Node N will also receive overall votes from its direct one-hop neighbors dynamically as they are cast and pushed out toward N . It will then compute the average of its neighbors’ scores to get V_{avg} . To mimic taking friends’ recommendations in real-life decision-making, node N ’s overall vote for P will be the weighted value $\alpha V_N + (1 - \alpha)V_{avg}$. By assuming that each of N ’s neighbors use the same weight value α , the average of all votes from node N ’s L^{th} -degree friends contribute $\alpha(1 - \alpha)^L$ to its overall score for P . Since $0 < \alpha < 1$, N ’s overall score for P will converge despite the iterative recomputation and mutual dependency of weighted votes between itself and its neighbors.

At a particular AS, a separate server will run the distributed reputation service that verifies characteristics about underlying BGP paths with the opinions of trusted peers in the overlay. The server will monitor received paths through internal BGP sessions with the AS’s border routers. When a BGP route is flagged as suspicious (*i.e.*, a query returns a negative response, below a given threshold $-T$), the server will automatically reconfigure the the router’s policies to filter out the suspect route. To implement a reputation node, a network operator need only modify the configuration of the router, not the underlying hardware or software of the existing routing mechanism.

2.4. Protocol Advantages

The distributed reputation protocol does not require a centralized PKI or registry. By constructing an overlay

according to the protocol, trust relationships are automatically built into the network. We show below that our protocol adheres to the three properties of a useful reputation scheme and assert three additional systemic advantages pertinent to BGP security.

1. **Strong deterrent for malicious behavior:** The primary architectural benefit is the protocol’s suppression and deterrence of misbehavior. Acting maliciously hurts an AS’s direct neighbors the most since the weight factor α skews the impact of lies to hurt those who trust him. Repeated lies will not only jeopardize his online reputation, but may also lead to other real-world repercussions through the loss of real-world trust. Even if a node wrongly trusts a misbehaving or compromised AS, the node can simply disconnect it as a direct neighbor in the overlay network to prevent future harm.
2. **Difficult to shill the entire system:** The weight factor α also forces a node’s vote to have geometrically smaller impact the further away it propagates. The effect of a single malicious vote, even coming from a highly regarded AS, will have only a small impact on direct peers and a negligible effect on ASes multiple hops away in the overlay. Even more importantly, colluding adversaries who form a clique in the overlay network can only do as much damage as a single isolated adversary. The best attack, where colluding adversaries attach to many locations in the network, is made difficult by the need of out-of-band authentication and real-life trust relationships.
3. **Intuitive scoring system:** With weight factor α , an AS can look at its overall vote for a proposition and infer both its own and its neighbors’ beliefs on that query. A overall vote $> \alpha$ or $< -\alpha$ signifies a personal “true” or “false” belief, respectively. The vote’s variation from $\pm\alpha$ signifies the confidence level of its peers. A vote between $-\alpha$ and α denotes a lack of information to make an informed decision. The AS either has a conflicting opinion from its neighbors, or has no opinion and is relying solely on his peers’ opinion of the proposition.
4. **Incrementally deployable:** Our proposed solution makes no changes to the de facto BGP protocol or the packet format of current BGP path announcements. The reputation system is completely disjoint from BGP’s control plane and can be adopted voluntarily at any time by network operators. The system would also function regardless of the choice of local detection tools run at each AS.

5. **Automated voting and decision-making:** Since the protocol depends on local analysis tools, it is possible to build in hooks that allow the existing mechanisms to automatically cast votes in the system. In addition, if an AS computes an overall score for a proposition that is above or below a set threshold T , it can automatically advise its routers to filter out suspicious paths and choose a more trustworthy one.
6. **Confidentiality of AS relationships:** ASes closely guard confidential information about their business relationships with other ASes. An AS can simply abstain from voting on a proposition that might divulge its business relationships with its neighbors. Overall votes are also passed only to one's direct neighbors since votes will be aggregated and averaged before being propagated further. Cryptographic mechanisms will encrypt and sign the votes cast in the overlay and provide repudiability such that the votes will be unverifiable to a third party.

2.5. Protocol Limitations

As a general framework, the reputation system is *not* inherently capable of detecting BGP misconfigurations and attacks itself. It is only useful when deployed in conjunction with the collection of available tools for debugging local networks. The positive tradeoff, though, is that the architecture itself is agnostic to a specific problem in question, allowing *any* boolean proposition to be raised and voted on. For a given proposition, the protocol works well if a minority of participants lie or believe a particular lie. Otherwise, the system could provide a slippery slope as principals unknowingly trust and propagate a faulty belief.

An honest AS might propagate false information in the reputation system if its local detection tools are not sufficient to uncover the fault. Given that *some* nodes in the overlay are able to diagnose the fault, the node would ideally discover the fault based on information collected from these trusted overlay peers. A trustworthy AS might also advertise bogus votes if it is compromised by a masquerading attacker. To reduce this threat, one basis of trust must be an understanding of a peer's security policies (*i.e.*, it follows best common security practices) before making a direct overlay connection. It matters not why an AS is receiving false information from another AS as long as it can manually drop the neighbor as a peer in the overlay when necessary.

As such, the reputation system is initially a reactive solution that may yield a small time lag between the injection of a new fault and its global detection. This time lag stems from the duration necessary for ASes to

run local detection mechanisms and for these results to propagate through the reputation network. Once enough ASes discover the fault locally, the scheme begins to behave proactively for ASes who have not yet detected this fault. For example, upon receiving a BGP update message, a node can immediately reject a route based on a precomputed negative rating $< -T$ about this route gleaned from the trust system. An AS could alternatively accept a positively-rated route $> T$ despite its appearance on bogon filter lists [8].

By building the overlay network on top of routes in the potentially faulty underlying network, an adversary could also attack the overlay itself. End-to-end encryption and signing of votes would protect the injection of bogus votes, but an adversary could still divert legitimate votes from its intended destination. A plausible yet expensive solution would be to create a separate out-of-band network that directly connects trusted peers.

3. Practical Applications

Currently in BGP, we can already verify *who* is speaking by using standard cryptographic techniques pairwise between communicating ASes. The preeminent problem in interdomain routing is verifying the truth of *what* ASes actually say. We hypothesize that our distributed reputation protocol can facilitate the detection of both misconfigurations and attacks. Our approach would make existing tools more effective by allowing trusted ASes to share data about widespread routing errors. Below, we speculate about the reputation system's BGP security potential using a preliminary examination of two types of common faults: prefix hijacks and invalid AS paths.

Our sketch relies on the common notion that it is impossible to detect all BGP faults affecting a single AS locally at that AS. While ASes are well-equipped to make truthful claims about routing connectivity within very close proximity (*i.e.*, prefixes originated by an immediate customer), they cannot be confident about connectivity claims many hops away. The vast majority of conflicts only affect a subset of ASes that are oblivious to the routing fault. Thus, it is imperative that ASes possess a means to verify with trusted peers the routes propagated from across the network.

3.1. Prefix Hijacks

One type of invalid route announcement is an IP prefix hijack. This occurs when an AS announces direct reachability to an IP prefix it does not actually own, either inadvertently or with ill-intent. This scenario often yields two announcements, a legitimate and a spoof,

from different ASes claiming to originate the same IP prefix. Zhao *et al.* studies extensively these Multiple Origin AS (MOAS) cases and suggests a method for detecting when these types of failures occur [24]. Their solution seems effective at raising alarms upon detection but offers no real capability to determine which announcement is actually legitimate. Their suggestion of a DNS database with valid origin information will inevitably become outdated and unreliable for practical use.

As an alternative, operators can amplify the above method by utilizing a distributed reputation system to identify the legitimate announcement upon a raised alarm. A hijacked prefix announcement only reaches a subset of the network since ASes closer to the real origin will continue to use the legitimate route. Assuming that an AS has overlay peers at diverse locations in the physical network, it can check upon detecting a new MOAS alarm whether the new route is a hijacked prefix. The conflicted AS could propose a boolean inquiry such as “Is AS88 entitled to originate prefix 128.112.0.0/16?” to assess the legitimacy of the route at hand. Unlike connections in the physical AS topology, trusted peers in the overlay would presumably have better information since well-configured ASes are more likely to filter bogus advertisements and less prone to adversarial reconfigurations. Another example is if AS_1 in the United States trusts AS_2 in Asia to originate prefix P . AS_1 could then easily disregard a future prefix hijack of P even though it is many hops away in the physical AS connectivity. Even if the AS has no overlay peers outside the affected area, it may still be helpful to share local traceroute and data plane verification [22] information with other affected peers.

3.2. Invalid AS Paths

Another type of BGP fault is the announcement of an AS path that does not actually exist. Kruegel *et al.* have formulated a heuristic to detect invalid path anomalies based on insights about the Internet’s general topology [16]. They divide the topology into a group of core backbone ASes and clusters of periphery ASes. First, they claim that any valid AS path can never traverse the Internet backbone more than once and thus may contain only a single subsequence of core ASes. Their second constraint stipulates that all consecutive pairs of periphery ASes in the path must either be in the same cluster or two geographically proximate clusters. With the Internet’s topology constantly evolving, it can be practically difficult to obtain a precise and up-to-date topology using a service like RouteViews [9]. Moreover, an AS might not trust the BGP routing information provided by

these data publishing services.

To add precision to the above heuristic, an AS could issue propositions in the reputation system for each AS-AS edge in question to determine topological connectivity. The distributed reputation system can even be used to determine the cause of the invalid path. For example, a network operator concerned about a potential spoofing attack can poll her trusted peers with the proposition “Is the edge $AS3 \rightarrow AS7$ a spoofed edge?”. Alternatively, she could also pose the proposition “Is the edge $AS3 \rightarrow AS7$ currently down?” if she suspects a temporary failure. Augmenting the ideas proposed by Kruegel with reputation could yield a more accurate network model with built-in confidence. It will also allow ASes to probe globally for fine-grained information to pinpoint local conflicts.

3.3. Other Applications

Our solution is obviously not limited to these two above possibilities, but can be used generally any time information from collaborating vantage points might be useful. Another potential application would be to apply root-cause analysis in real time using the techniques described by Feldmann *et al.* [13]. Using RouteViews data from multiple vantage points, they demonstrate a method to locate a set of suspect edges that might have initiated an AS path change. A network operator might not trust information provided by the vantage points or might not have a complete view of other ASes. A distributed reputation system would introduce confidence in the accuracy of collected data and would integrate information from all border routers at an AS into individual votes.

Wu *et al.* developed a tool to pinpoint anomalies such as flapping prefixes and large traffic shifts near an AS by monitoring BGP updates at its border routers [23]. An AS implementing this tool could make assertions about traffic anomalies within close proximity and merge these results with analysis from other trusted vantage points using the reputation system. One might also utilize reputation systems to generate better AS topologies, to detect policy conflicts, to defend against denial of service attacks or to resolve other interdomain security issues.

4. Conclusion and Future Work

This paper demonstrates the promise of using our novel distributed reputation protocol to protect the Internet’s interdomain routing infrastructure. Our solution diverges from recent work by focusing on how ASes can collaborate in a trustworthy manner to improve resilience against attack and misconfiguration. We de-

scribe a lightweight overlay protocol that can be adopted incrementally by the AS network operator community. The reputation architecture significantly raises the difficulty of performing large shilling attacks and provides a strong deterrent against malicious behavior. Adopters will be able to receive and share valuable routing information with well-reputed peers at many vantage points and still keep sensitive data confidential.

As this paper hypothesizes about our reputation system's potential benefits for BGP security, future work will focus on implementation specifics and validation. We plan on devising efficient methods to dynamically recalculate and manage the large number of received votes. We also aim to quantify how resilient the system is against shilling attacks by isolated and colluding adversaries in a realistic Internet topology. We will also look to determine what specific queries work best to debug different types of BGP faults in practice and attempt to integrate existing local tools to optimize accurate and instantaneous voting.

Acknowledgment

We thank Nick Feamster and the anonymous reviewers for their many detailed and insightful comments.

References

- [1] Africa Network Operators' Group. <http://www.afnog.org>.
- [2] IvyPlus Admin Computing Practices. <http://ivyplus.stanford.edu/practices.html>.
- [3] NANOG PGP Key Signing. <http://www.nanog.org/pgp.abley.html>.
- [4] NSP-Security (NSP-SEC) Info Page. <https://puck.nether.net/mailman/listinfo/nsp-security>.
- [5] South Asian Network Operators Group. <http://www.sanog.org>.
- [6] The Network Reliability and Interoperability Council: Best Practices. <http://www.bell-labs.com/user/krauscher/nric>.
- [7] The North American Network Operators' Group. <http://www.nanog.org>.
- [8] The Team Cymru Bogon Reference Page. <http://www.cymru.com/Bogons>.
- [9] University of Oregon Route Views Project. <http://www.routeviews.org/>.
- [10] V. Bono. 7007 Explanation and Apology. <http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html>.
- [11] K. Cooper. An Information Sharing and Analysis Center for the Internet. <http://www.nanog.org/mtg-0105/cooper.html>, 2001.
- [12] N. Feamster and H. Balakrishnan. Detecting BGP Configuration Faults with Static Analysis. In *Proc. USENIX NSDI*, 2005.
- [13] A. Feldmann, O. Maennel, Z. Mao, A. Berger, and B. Maggs. Locating Internet Routing Instabilities. In *Proc. ACM SIGCOMM*, 2004.
- [14] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin. Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing. In *Proc. ISOC Network and Distributed Systems Security*, 2003.
- [15] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (S-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, 2000.
- [16] C. Kruegel, D. Mutz, W. Robertson, and F. Valeur. Topology-based Detection of Anomalous BGP Messages. In *Proc. RAID*, 2003.
- [17] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP Misconfiguration. In *Proc. ACM SIGCOMM*, 2002.
- [18] K. Moriarty. Distributed Denial of Service Incident Handling: Real-Time Inter-Network Defense. <http://ietfreport.isoc.org/all-ids/draft-moriarty-ddos-rid-06.txt>.
- [19] J. Ng. Extensions to BGP to Support Secure Origin BGP (soBGP). IETF Internet Draft draft-ng-sobgp-bgp-extensions-02.txt, April 2004.
- [20] O. Nordstrom and C. Dovrolis. Beware of BGP Attacks. *ACM SIGCOMM Computer Communications Review*, 34:1–8, 2004.
- [21] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). RFC 1771, March 1995.
- [22] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. Katz. Listen and Whisper: Security Mechanisms for BGP. In *Proc. USENIX NSDI*, 2004.
- [23] J. Wu, Z. Mao, J. Rexford, and J. Wang. Finding a Needle in a Haystack: Pinpointing Significant BGP Routing Changes in an IP Network. In *Proc. USENIX NSDI*, 2005.
- [24] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang. Detection of Invalid Routing Announcement in the Internet. In *Proc. IEEE International Conference on Dependable Systems*, 2002.