New Approach For Path Vector Protocol

Dr.E.Gajendran[1], Dr.S.Britto Raj[2], Mohd Ayaz Uddin[3],Dr.S.Sharavanan[4],Dr.Leonard Gibson Moses[5], Dr.K.Ravikumar[6] , Dr.E.Mohan[7]

[1,2]Associate Professor, [7] Professor, Department of Computer Science Engineering, Mohamed Sathak A.J College of engineering, Chennai, India
[3]Assistant Professor, Department of AIML, Sphoorthy Engineering College, Nadegul, Hyderabad, India
[4]Professor, Department of Computer Science Engineering, Karpagam Institute of Technology, Coimbatore, India
[5]Associate Professor, Department of ECE, Kings Engineering College, Chennai, India
[6]Associate Professor, Department of CSE, RRASE College of Engineering, Chennai, India

[1]gajendrane@gmail.com, [2]brittorajs@gmail.com, [3]mohdayazuddin@sphoorthyengg.ac.in, [4]sharavanan.cse@karpagamtech.ac.in, [5]gibs.ml@gmail.com, [6]ravikumarcsephd@gmail.com, [7]emohan1971@gmail.com

ABSTRACT

Previous measurement studies have shown the existence of path exploration and slow convergence in the global Internet routing system, and a number of protocol enhancements have been proposed to remedy the problem. However, existing measurements were conducted only over a small number of testing prefixes. There has been no systematic study to quantify the pervasiveness of Border Gateway Protocol (BGP) slow convergence inthe operational Internet, nor any known effort to deploy any of the proposed solutions. In thispaper, to present the measurement results that identify BGP slow convergence events across the entire global routing table. The data shows that the severity of path exploration and slow convergence varies depending on where prefixes are originated and where the observations are made in the Internet routing hierarchy. In general, routers in tier-1 Internet service providers (ISPs) observe less path exploration, hence they experience shorter convergence delays than routers in edge ASs; prefixes originated from tier-1 ISPs also experience less path exploration than our data show that the convergence time of route fail- over events is similar to that of new route announcements and is significantly shorter than that of route failures. This observation is contrary to the widely held view from previous experiments but confirms the earlier analytical results. The effort also led to the development of a path-preference inference method based on the path usage time, which can be used by future studies of BGP dynamics.

Keywords: BGP, ISP, Protocol, Autonomous Systems ,router

1. INTRODUCTION

The first quantitative assessment on path explorations for the entire Internet destination prefixes. There results confirmed the wide existence of path exploration andslow convergence in the Internet, but also revealed that the extent of the problem depends on where a prefix is originated and where the observation is made in the Internet routing hierarchy. In other words, the existing widely different opinions on the extent of path exploration and slow convergence may be a reflection of where one takes measurement and which prefixes are being examined. Finally, route failure events, denoted as T-$_{down}$, have a substantially longer delay than all the above events. In short, we have T-$_{short}$ < T-$_{up}$ < T-$_{long}$<< T-$_{down.}$ A major challenge in our data analysis is how to differentiate T-$_{long}$ and T-$_{short}$ events, which requires knowing routers' path preferences. We have developed a new path ranking algorithm to infer relative preference of each path among all the alternative paths tothe same destination prefix. We believe that our path ranking algorithm can be of useful in many other BGP data analysis studies. The rest of the paper is organized as follows. Section 2 describes our general methodology and data set where we develop a path ranking algorithm to classify events into different types.

1. PROBLEM STATEMENT

In this section, we propose a simple and novel mechanism forward edge sequence numbers to annotate routing updates with path dependency information, so as to effectively address the path exploration problem. Using this mechanism, we develop an enhanced path vector routing protocol, EPIC, which limits path exploration and thereby leads to faster protocol convergence after network failure andrepair events. Our solution has the following properties. In contrast to previous solutions which assume a simplified setting, our solutionis based on a more general and realistic modelof BGP operation and AS topology: ASes may contain internal routers and share multiple edges with neighboring ASes .It does not require ASes to expose detailed connectivity information. It can be implemented with fairlymodest communication and memory overheadthe remainder of this paper is structured reviews BGP operation and illustrates the path exploration problem. We introduce the proposed novel mechanism for embedding path dependency, i.e., forward edge sequence numbers and use examples to show how they are used. A detailed description of EPIC is presented along with correctness results and lists some analytical results for EPIC and simulation results are presented. Finally, we review some related work and conclude BGP is used between ASes to exchange network reach ability information. Each AS has one or more border routers that connect to routers in neighboring ASes, and possibly a number of internal BGP routers. BGP sessions between routers in neighboring ASes are called eBGP (external BGP) sessions, while those between routers in the same AS are called iBGP (internal BGP) sessions. Note that adjacentASes may have more than one eBGP session. They now briefly describes the relevant operation at a BGP router.

BGP routers distribute reach ability information about destinations by sending route updates, containing announcements or withdrawals, to their neighbors. In the rest of this paper, we implicitly assume a destination, say. A route announcement contains a destination and a setof route attributes, including the AS path attribute, which is a sequence of AS numbers that enumerates all the ASes traversed by the route. We denote an AS path as in the origin AS to which belongs. In contrast, route withdrawals only contain the destination and implicitly tell the receiver to invalidate (orremove) the route previously announced by the sender. When a router receives a routeannouncement, it applies an littering process (using some import policies). If accepted, the route is stored in the local routing table. After this, the selected best route is subjected to some export policies and thenannounced to all the router's neighbors. Importantly, prior to being announced to an external neighbor, but not to an internal neighbor in the same AS, the AS path carried in the announcement is pretended with the ASN of the local AS. Vectoring protocols are inherently associated with path dependencies: the path selected by a router depends on paths learned by its neighbors which, in turn is influenced by the paths selected at the neighbors' peers, and so on. This natural property leads to the so-called path explorationphenomenon that prolongs protocol convergence. Note that in path vector protocols, the path vectors are used to prevent routing loops, but they cannot avoid path exploration. As a path vector protocol, BGP exhibits path exploration. More significantly, it introduces additional complexity that makesit particularly difficult to address this problem.In the rest of this section, we illustrate the path exploration phenomenon by an example, then describe why, in general, it is impossible to avoid it by solely relying on the AS paths associated with BGP routes.

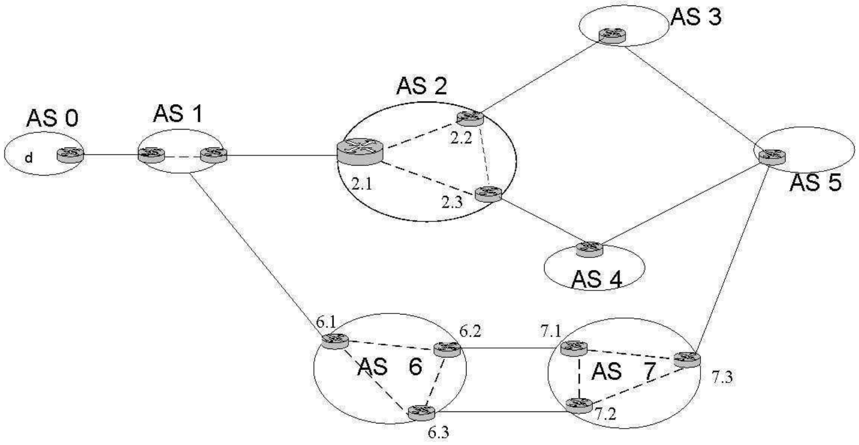

Fig. 1. BGP and Path Exploration

Consider the topology in Fig. 1. Now suppose it announces a path to destination this announcement is received at its neighbors and propagated hop by hop. Finally, when the network converges knows three paths to reach .Now consider what happens when the link between and it fails, making unreachable This failure triggers the following sequence of events In turn, each of them sends withdrawals to their own neighbors. Eventually the "best path" and sends it to its (other) neighbors. However, if the withdrawal arrives next, then this "best route" is invalidated Finally, after receives the withdrawal from invalidates the path announced earlier and sends a withdrawal. This cycle of selecting and propagating (invalid) paths is termed path exploration. Clearly, the cycle stops after all the obsolete routes have been explored and invalidated.

2. RELATED WORK

2.1 Limiting Path Exploration in BGP

The Internet is a collection of independently administered Autonomous Systems (ASes) gluedtogether by the Border Gateway Protocol (BGP) [1], the *de facto* inter-domain routing protocol in the Internet. BGP is a path vector routing protocol where the list of ASes along the path to a destination (AS path) iscarried in the BGP routing messages. Usingthese .path vectors., BGP can avoid the looping problems associated with traditional distance vector protocols. However, BGP may still take relatively long time to converge following a network failure. Experimental studies show that, in practice, BGP can take up to 15 minutes to converge after a failure The root cause of this slow convergence is the dependency among paths announced through the network, leading to path exploration:when a previously announced path is withdrawn, other paths that *depend* on the withdrawn path (now *invalid*) may still be chosen and announced, only to be removedlater one by one. During path exploration, thenetwork as a whole may explore a large number of (valid and invalid) routes beforearriving at a stable state. Theoretically, in the worst case, a path vector routing protocol canexplore as many as O(n!) alternative routes before converging. Addressing path exploration within the framework of BGP is particularly challenging: AS paths carried inBGP route advertisements are highly summarized, making it difficult to capturedependencies between different paths and tocorrectly distinguish between valid and invalidpaths.

2.2 The Temporal and TopologicalCharacteristics of BGP Path Changes :

BGP is a policy-based path-vectorrouting protocol deployed in Internet for inter- domain routing. The Internet is divided intotens of thousands of autonomous routingdomains, of which over 15 thousand are currently associated with Autonomous System Numbers (ASNs) for the purpose of inter- domain routing. BGP routers in each AS transmit routing messages to other BGProuters in the same AS and other ASes throughinternal and external BGP connections, respectively. Routing messages containingreach ability information are called BGPupdates. To facilitate the study on theoperational use of BGP, there are public BGP routing message collection sites such as RIPE’s RRCs (Remote Route Collectors) and Oregon University’s Route- Views that collectBGP updates and routing tables from tens ofBGP routers located in various ASes. These data sets provide researchers and operators alocal perspective on the visible Internet BGProuting status. Researchers have been usingthe collected routing tables and routingmessages to study the Internet topology at AS- level , monitor the Internet growth , examine the inter-domain routing stability, investigate BGP router miss-configuration , and derive the model for BGP traffic . In this paper we present a systematic approach to decompose the stream of BGP updates into small sequences of path advertisements with the purpose of distinguishing the routing eventsthat cause the BGP routing changes.

2.2.1 Experimental Study of Internet Stability and Wide-Area Backbone Failure :

In this paper, we describe an experimental study of Internet stability and the origins of failure in Internet protocol backbones. Unlike telephony networks, the stability of end-to-endInternet paths is dependent both on theunderlying telecommunication switching system, as well as the higher level software and hardware components specific to the Internet's packet-switched forwarding, name resolution and routing architecture. Although a number of vendors provide mean-time to failure statistics for specific hardware components used in the construction of wide- area networks (e.g. power supplies, switches, etc.), estimations of the failure rates for IP backbones at a systemic level remain problematic. As we describe below, the interactions between the underlying components of the Internet are poorly understood . Typical analysis of faults in telephony networks has focused on the number of customers affected by an outage . The US Federal Communication Commission requires service providers to report all outages lasting 30 minutes or more and affecting 30,000 customers or more . No such reporting requirements yet exist for Internet providers. And, if such requirements did exist, the same estimations of the impact of failures would beproblematic for Internet providers. Both thedefinition of failure and even \end-user" aresomewhat ambiguous on the Internet. Incontrast to the fixed bandwidth used bytelephony, Internet applications and end-users have widely disparate bandwidth, latency and loss requirements. For example, the failure of an Internet T3 link (45 MB) may impact one large weather simulation at a supercomputercenter, or several thousand web-surfing dial-up users. In our analysis, we make no effort to quantify the significance of Internet outages based on the number of users affected. Instead, we focus on the number of individual link or interface failures,and the number of unreachable network destinations. In general, the Internet exhibits a number of engineering and operational challenges distinct from those associated with telephony networks andapplications. Most significantly, unlike switched telephony networks, the Internet is a conglomeration of thousands of heterogeneous dynamically packet switched IP backbones. No resources are explicitly reserved for each datagram or IP data flow. Instead, the end-to-end quality of Internet performance dependson the impact of loss, queuing 2 delay and network congestion on each of the flow's individual datagram packets. So, for example, although the initial \call setup" of an Internet telephony application may succeed, all subsequent voice data grams in the connection may be lost due to network congestion. The relationship between loss, latency and end-to-end performance remains an area of activeresearch. In addition, the explosive growth in demand for Internet facilities and features has resulted in a significantly more rapid Internet software and hardware evolutionary testing and development cycle than traditional amongst PSTN equipment suppliers. For example, telephony switches typically undergo development cycles on the order of several years or even decades. In contrast, some Internet backbone routers and switches have development cycles lasting six months or less. Internet vendors regularly market backbone equipment featuring new software algorithms even before these protocols have advanced into official standards .The technological demands associated with the Internet's growth are so severe that Internet providers often depend on these newly released products or software features to sustain their network's continued expansion.

As a result ,the reliability of the Internetinfrastructure has arguably suffered. The rapid growth of IP backbones has also led to adecline in the relative level of experience and degree of coordination amongst Internet backbone operators. A number of significant recent Inter- net outages have stemmed from human error. Other outages have originated, or been exacerbated by lack of coordination between the backbone engineering staff of different Internet providers. In the PSTN network, a comparatively small number of telecommunication companies interact via well-defined, standardized channels using

uniform management, measurement and operational procedures. The significantly more diverse and less uniform Internet does not enjoy the same degree of coordination. Specifically, the Internet lacks central administration and coordination. Unlike traditional PSTN standards bodies whose formal membership requirements are defined by inter- national treaty, the only requirement for participation in the three yearly Internet standards meetings is showing up . We briefly describe some recent Internet outages which directly, or indirectly, impacted a majority of Internet backbone paths. Although several major incidents stemmed from underlying PSTN failures, we focus below on faults specific to the Internet.

3. PROJECT DESCRIPTION

3.1 Data Set and Pre-processing

To develop and calibrate our update grouping and path ranking heuristics, we used eight BGP beacons, one from PSG, the other seven from RIPE. All the eight beacon prefixes are announced and withdrawn alternately every 2 hours. First, we removed from the update stream all the duplicate updates, as well as the updates that differ only in Community Or Med attribute values, because these updates are usually caused by internal dynamics inside the last-hop AS. Second, we used the anchor prefix of each beacon to detect routing changes other than those generated by the beacon origins.

3.2 Clustering Updates

Based on the observation that BGPupdates come in bursts, two adjacent updates for the same Prefix is assumed to be due to the same routing event if they are separated by a time interval less than a threshold T. A critical step in taking this approach is to find an appropriate value for T. A value that is too high can incorrectly group multiple events intoone. On the other hand, a value that is too low may divide a single event into multiple ones. Since the root causes of beacon routing events are known, and the beacon update streams contain little noise after the preprocessing, we use beacon prefixes to find an appropriate value for T.

3.3 Classifying Routing Events

After the routing updates are grouped into events, we classify the events intodifferent types based on the effect that each event has on the routing path. Let us consider two consecutive events n and n+1 for the sameprefix observed by the same monitor. We define the path in the last update of event n as the ending path of event n, which is also the starting path for event n + 1.

3.4 Comparing AS Paths

If a routing event has non-empty p- start and p-end, then the relative preference between p-start and p-end determines whetherthe event is a T-long or T-short. This would bean easy task for controlled experiments using beacon prefixes, since one simply create such events by manipulating AS paths. We compare this new Usage Time based approach with three other existing methods for inferring pathpreference: Length, Policy, and Policy + Length. One challenge in conducting this comparison is how to verify the path ranking results without knowing the router’s routing policy configurations.

.

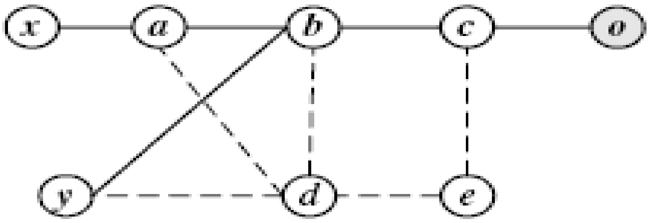

Fig 2:.Example of AS Topology

3.5 Characterizing Events

After applying the classification algorithm to BGP data, we count the number of T-down events observed by each monitor as a sanity check. A T-down event means that a previously reachable prefix becomes unreachable, suggesting that the root cause of the failure is very likely at the AS that originates the prefix, and should be observed by all the monitors. Therefore, we expect every monitor to observe roughly the same number of T-down events.

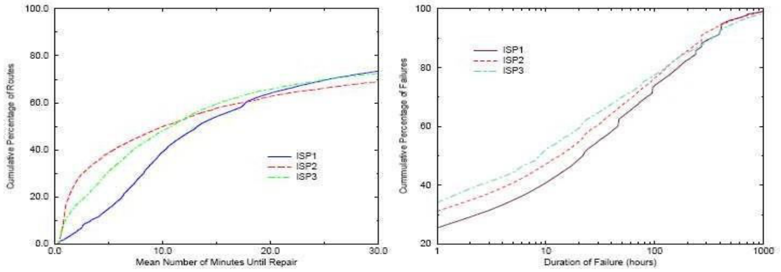

Fig 3.Cumulative distribution of MTTR

4. CONCLUSION

We conducted the first systematic measurement study to quantify the existence of path exploration and slow convergence in the global Internet routing system. We first developed a new path ranking method based on the usage time of each path and validated its effectiveness using data from controlled experiments with beacon prefixes. We then applied our path ranking method to BGP updates of all the prefixes in the global routing table and classified each observed routing event into three classes: Path Change, Path Disturbance, and Same Path. For Path Change events, we further classified them into 4 sub- categories: T-down, T-up, T-long, and T-short. We measured the path exploration, convergence duration, and update count for each type of events.

5. REFERENCES

1. PSG Beacon List. Available from:http://www.psg.com/zmao/BGPBeacon..RIPE Beacon List. Available from:ttp://www.ripe.net/ris/docs/beaconlist.html [Cited 05/11/2006].
2. Shailedra Kumar Shrivastava, Sanjay S. Gharde. Support Vector Machine for Handwritten Devanagari Numeral Recognition. International Journal of Computer Applications (0975 – 8887) Volume 7– No.11, October 2010.

3. A.Rajesh, Dr.E.Mohan, "Classification of microcalcification based on wave atom transform", Journal of computer science, 10 (9), 1543-1547, 2014.
4. Venkatachalam, K., Reddy, V. P., Amudhan, M., Raguraman, A., & Mohan, E. (2021, June). An implementation of K-means clustering for efficient image segmentation. In 2021 10th IEEE international conference on Communication Systems and Network Technologies (CSNT) (pp. 224-229). IEEE.
5. Gladstan, T., & Mohan, E. (2017). A Novel Approach Object Recognition Using Efficient Support Vector Machine Classifier. International Journal of Electronics and Communication Engineering and Technology, 8(2), 81-90.
6. Rajesh, A., & Mohan, E. (2016). Classification of Mammogram Using Wave Atom Transform and Support Vector Machine Classifier. International Journal of Computer Science and, 467-470
7. Mohan, E., Sugumar, R., & Venkatachalam, K. (2014). Automatic brain and tumor segmentation in MRI using fuzzy classification with integrated Bayesian. *Int. J. Appl. Eng. Res*, *9*(24), 25859-25870.

8. Thambu Gladstan, Dr.E.Mohan. Object Recognition Based on Wave Atom Transform**.** RJAET. 8(13): 1613-1617, 2014.