ISSN No. 0976-5697

Volume 4, No. 6, May 2013 (Special Issue)

International Journal of Advanced Research in Computer Science

REVIEW ARTICAL

Available Online at www.ijarcs.info

"The Routing of an Autonomous System with C-Border Gateway Protocol in Network"

Miss. Snehal D. Nanhore
G.H. Raisoni college of engineering and management, Amravati
snehal.nanhore@gmail.com

Mr. Mahip M. Bartere
G.H. Raisoni college of engineering and management, Amravati
mahip_media@yahoo.com

Abstract: The Internet has quickly evolved into a vast global network owned and operated by thousands of different administrative entities. During this time, it became apparent that vanilla shortest path routing would be insufficient to handle the myriad operational, economic, and political factors involved in routing. ISPs began to modify routing configurations to support *routing policies* — goals held by the router's owner that controlled which routes were chosen and which routes were propagated to neighbors. BGP, originally a simple path vector protocol, was incrementally modified over time with a number of mechanisms to support policie, adding substantially to the complexity. Much of the mystery in BGP comes not only from the protocol complexity, but also from a lack of understanding of the underlying policies and the problems ISPs face that are addressed by these policies. Today, the complexity of ISPs' networks make it difficult to investigate the implications of internal or external changes on the distribution of traffic across their network. The complexity of building models of large ISPs' networks. We describe the various aspects important to understanding the routing inside an AS. We present an open source routing solver, C-BGP, that eases the investigation of changes in the routing or topology of large networks. We illustrate how to build a model of an ISP on a real transit network and apply the model on two "what-if" scenarios. The first scenario studies the impact of changes in the Internet connectivity of a transit network. The second investigates the impact of failures in its internal topology.

I. INTRODUCTION

When BGP was first introduced, it was a fairly simple path vector protocol. Over time, many incremental modifications to allow ISPs to control routing were proposed and added to BGP. The end result was a protocol weighed down with a huge number of mechanisms that can overlap and conflict in various unpredictable ways. These modifications can be highly mysterious since many of them, including the decision process used to select routes, are not part of the protocol specification. Moreover, their complexity gives rise to several key problems, including unforeseen security vulnerabilities spread misconfiguration, and conflicts between policies at different ISPs. Addressing BGP's problems is difficult, as changing certain aspects of BGP (e.g., changing the contents of update messages or the way they are propagated) must be coordinated and simultaneously implemented in other ISPs to support the new design. Hence, most modifications to the protocol have been made to the *decision process* BGP uses to choose routes.

The result is a protocol where most of the complexity is in the decision process and the policies used to influence decisions, while the rest of the protocol has remained fairly simple overtime. Therefore, in order to understand BGP it is necessary to understand this decision process and the policies of ISPs that gave rise to its design. Understanding policies is also key to solving BGP's problems, understanding measurement data from BGP, or determining which features to support when developing a new version of BGP. The range of policies used by operators constitutes a huge space; hence, it is impossible to list them all here. Instead, we try to list common goals of network operators and the knobs of BGP that can be used to express policies. In particular, we attempt to isolate certain *design patterns* commonly used by ISPs, the motivations behind them, and how they are implemented in an ISP's network using BGP's mechanisms. We taxonomize policies into four general categories: *business relationship* policy arising from economic or political relationships an ISP has with its neighbor, *traffic engineering* policy arising from the need to control traffic flow within an ISP and across peering links to avoid congestion and provide good service quality, policies for *scalability* to reduce control traffic and avoid overloading routers, and *security*-related policies that are often used to protect an ISP against malicious or accidental attacks.

We also discuss several avenues of research currently in progress related to BGP policies. We start by giving an overview of BGP routing in the next section. The physical topology of an AS defines feasible paths that can be used to cross the network. How traffic actually crosses the network depends on the choices made by the routing protocols.These choices depend on two major factors: the diversity of available routes and router configurations. The diversity of the available routes known by an AS depends on the routing information received from neighboring ASs. Among these available routes, the routing protocols choose which one will be used to reach each destination. This choice depends on the goals of the network operators expressed in the router configuration. understanding the routing of large ASs requires not only modeling the routing inside the AS, but also taking into account routing information received from neighbor ASs. We explain how routing in an AS works. We describe how to model the routing of an ISP's network. We explain what information is required in order to build such a model. We show how this information is processed by an open source tool we developed. Finally, we provide two applications of our tool to study the behavior of a transit AS.

A. C-BGP: A BGP Solver for Large Ass:

We are not aware of the existence of any tool that fully captures the aspects described earlier. The most closely related works from the literature are [4, 5]. The aim of [4] was to provide the networking industry with a software system to support traffic measurement and network modeling. This tool is able to model intradomain routing

CONFERENCE PAPER
"A National Level Conference on Recent Trends in Information Technology and Technical Symposium" On 09th March 2013
Organized by
Dept. of IT, Jawaharlal Darda Inst. Of Eng. & Tech., Yavatmal (MS), India

and study the implications of local traffic changes, configuration, and routing. However, It does not model the interdomain routing protocol. Reference .It proposed a BGP emulator that computes the outcome of the BGP route selection process for each router in a single AS. This tool does not model the flow of the BGP routes inside the AS, so it does not reproduce the route filtering process occurring within an AS. Neither of these tools is publicly available .In this section we describe C-BGP, an open source routing solver we developed. C-BGP can be used by ISP network operators to study routing what-if scenarios based on routing information collected in their network.

The solver takes several sources of information into account. First, it takes a description of the network topology at layer 3. Then it takes the configuration of all the routers present in the topology. This configuration describes the IGP weights of all the links, the BGP peerings of each router, and the BGP policies that must be enforced on each peering. We are able to parse Cisco and Juniper configuration files and generate configurations suitable for C-BGP. Finally, the tool takes the BGP routes learned by the ISP network on its border routers. As output, the solver computes for each router the routes selected toward all the interdomain prefixes. This output can then be used to replay how the traffic was routed by the routers of the AS. In order to accurately model the routing in an ISP's network, we need to precisely model the path selection performed by the intradomain and interdomain routing protocols. That is, we must compute for each router the next hop that would have been selected to reach each destination prefix. Our solver models the topology of the network, the IGP, the eBGP and iBGP sessions, the iBGP hierarchy with route reflectors, the BGP route filtering, and the complete BGP decision process. Modeling all aspects of BGP is time- and resource-consuming. To keep our model scalable and efficient, we do not model the time-consuming packet exchanges that occur between simulated routers in traditional discrete-event simulators such as SSFNet [12], J-Sim [13], or ns [14]. In addition, we do not model the TCP connections that support BGP sessions.

We also do not model BGP timers such as the MRAI. We are therefore able to model large ISP networks.

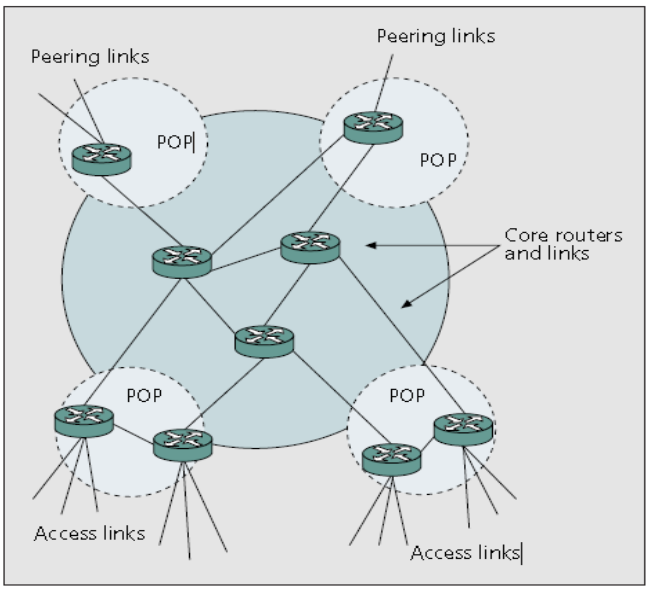

Figure 1. Topology of an example autonomous system.

B. Routing in an Autonomous System:

On its IP-level topology, an AS runs two different routing protocols. First, it runs an IGP such as OSPF or IS-IS in order to compute the interior paths from any AS's router toward the AS's other routers and subnets. The IGP is typically a link state protocol; that is, it floods information about the state of the adjacencies between all routers in the whole AS. The objective of intradomain routing is to find the shortest paths according to a selected metric. ISPs usually use a metric that is proportional to the propagation delay along the path or the bandwidth. Many network operators use the Cisco default metric, which is one over the bandwidth [1]. Some large Ass use a hierarchical IGP, where the AS is divided into different areas. Inside an area all the adjacency information is flooded. Between areas only aggregated information is exchanged. In addition to the IGP, an AS sometimes uses static routing. Static routes are often used on edge links since routers on both side of these links are not operated by the same authority. Static routes are also used to set up access to small customers that do not use BGP. Finally, an AS runs BGP [1]. BGP is responsible for the selection of the interdomain paths for this reason, these ASs sometimes deploy route reflectors [6] in their network. Route reflectors are special BGP routers that make possible a hierarchy of iBGP sessions, thereby reducing the number of iBGP sessions. It is also possible to reduce the number of iBGP sessions by using BGP confederations [1]. Through its BGP sessions, each router receives BGP routes toward destination prefixes. Each router uses its decision process on a per-prefix basis to select the routes it will use.

The BGP decision process is a sequence of rules that takes a set of routes toward the same destination prefix and selects a single route, called the *best route*, toward this prefix. This route will be installed into the router's routing information base (RIB), copied in the forwarding table, and eventually used to forward packets. Basically, the BGP decision process ranks routes according to their attributes. Each rule of the decision process discards the routes it does not prefer.

The surviving routes are then submitted to the next rule, until a single route remains. The BGP decision process considers several of the BGP route's attributes. The first attribute is the *local-pref*, which corresponds to a local ranking of the route. It is usually attached to the route upon reception by a border router and is never propagated outside the AS. The decision process prefers the routes with the highest local-pref attribute value. The second attribute is the *as-path*. The as-path contains the sequence of ASs that the route crossed to reach the local AS. The as-path is used for two different purposes: avoiding routing loops and providing a distance metric in AS hops. The decision process prefers the routes with the shortest as-path. The third attribute is the *multi-exit discriminator* (in short, the med). This attribute is used to rank routes received from the same neighbor AS.. When a BGP router receives a route, it first checks that the next hop is reachable before considering it in the decision process. The decision process uses the IGP cost of the intradomain path toward the next hop to rank the routes. It prefers the routes with the smallest IGP distance to the next hop. This rule implements hot potato routing [7]. Its aim is to hand packets to a neighbor AS as soon as possible in order to consume as few network resources as possible. In addition, it automatically adapts routing to topology changes

CONFERENCE PAPER
"A National Level Conference on Recent Trends in Information Technology and Technical Symposium" On 09th March 2013
Organized by
Dept. of IT, Jawaharlal Darda Inst. Of Eng. & Tech., Yavatmal (MS), India

that affect the IGP distance to the egress points inside the AS.This step within the BGP decision process is where the IGP and BGP protocols interact.

C. Modeling an Autonomous System:

Modeling an ISP is a task that includes several aspects, starting with understanding the AS's architecture, gathering network data, building a representation of the AS's network, and ending up with a tool that allows the model to be exploite The first part toward building an AS's model consists of retrieving its configuration. The configuration of routers includes mapping between physical links and layer 3 links, the IGP metric associated with layer 3 links, the IGP hierarchy (areas), the BGP sessions, and the BGP policies enforced on each peering. However, handling the routers' configuration in a large network is difficult. First, in a large IP network, the volume of information found in the routers' configurations is far too large for a human to be able to deal with manually. Second, the configurations of routers are usually found in separate files, and there are frequently inconsistencies between these files [3, 8]. Finally, the network may be based on heterogeneous equipment; thus, the configurations are written in different configuration languages. Sometimes, some options even depend on the version of the network equipment's operating system. There is therefore a need to automate the process of analyzing the network configuration and properly report inconsistencies. Most of the time, discussion with the operator as well as cross-checking the files are required in order to exploit the network configuration.

a. Interdomain Routing Model:

As opposed to discrete event simulators, the propagation of messages in C-BGP is deterministic. Any run will lead to the same outcome, while in discrete event simulators the outcome of the simulation may depend on the seed of the pseudo random number generator. This has an impact on the convergence of the simulations performed with C-BGP. When a BGP configuration has multiple stable solutions(e.g., see the DISAGREE case [15]), the simulation will not converge. With discrete event simulators, the simulation might converge to one of the solutions in a nondeterministic manner. In a BGP configuration without a stable solution(e.g., the bad-gadget [15]), the behavior of C-BGP will be the same as with discrete event simulators.In order to model BGP, the nodes in the graph are considered as BGP routers and fitted out with additional data structures:a local RIB (Loc-RIB), adjacent RIBs (Adj-RIBs), and input and output filters. The Loc-RIB is used to store the best BGP routes, while the Adj-RIBs contain routes exchanged with neighbor routers.

We distinguish Adj-RIB-in that contains routes received from the neighbor routers RIB-out that contains routes announced to neighbor routers. The model works as follows. Once the network topology is available and the intradomain routes have been computed, the solver begins the propagation of route advertisements. The solver starts with an arbitrary BGP router and advertises the routes known by the router. These routes have previously been captured on the eBGP sessions of the routers being modeled. The solver supports MRTd dumps or manual injection of routes. For each route to be advertised, the solver builds UPDATE messages and sends them to the router's neighbors according to the output filters. For each BGP message to send, the solver looks up in the router's routing table the link along which the message must be forwarded to reach the next hop. The message is forwarded on a hop-by-hop basis until it reaches its final destination. The generated BGP messages are pushed in a single global linear first-in first-out queue that guarantees the BGP messages are received in sequence. In real routers the BGP message ordering is guaranteed by the TCP connections underlying the BGP sessions. The solver does this for all the BGP routers. The solver continues the simulation by popping the first message from the queue, and waking up the router corresponding to the current hop of the message. If the BGP message is a WITHDRAW, the router removes from the corresponding Adj-RIB-in the route toward the withdrawn prefix and runs the decision process. If the BGP message is an UPDATE, the router checks if the route it contains is accepted by its input filters. If so, the route is stored in the Adj-RIBin, and the router's decision process is run. The decision process retrieves from the Adj-RIB-ins all reachable routes for the considered prefix, compares them, and selects the best one. The router then propagates its new best route to its neighbors according to its output filters. The propagation is done by pushing new BGP messages on the global linear queue. The solver continues until the message queue is empty, which means that BGP has converged.

D. The Traffic Model:

In our model the traffic information of an AS is a set of triples (ingress router, destination, traffic volume). Each triple represents the traffic volume received by an ingress router to be sent toward the destination. This destination does not need to lie within the AS. These triples can be computed from Netflow statistics collected in the AS on the border routers or generated from synthetic traffic. To replay the flow of traffic across an AS, we take each triple, one at a time. Then we perform a longest matching in the routing table computed by the BGP solver for the considered ingress router in order to find the prefix that contains the destination. We then use the route associated with this prefix to "forward" the traffic. We repeat this step on a hop-by-hop basis. Using this traffic model, we are able to evaluate the impact of various what-if scenarios on the distribution of the traffic inside the AS. For instance, based on the paths followed by the traffic flows, we can compute the load of the internal links as well as the load of the peering links of the AS.

E. Usiness Relationships:

ISPs often wish to control next hop selection so as to reflect agreements or relationships they have with their neighbors. Three common relationships ISPs have are: *customer-provider*, where one ISP pays another to forward its traffic; *peer-peer*, where two ISPs agree that connecting directly to each other(typically without exchanging payment) would mutually benefit both, perhaps because roughly equal amounts of traffic flow between their networks; and *backup* relationships, where two ISPs set up a link between them that is to be used only in the event that the primary routes become unavailable due to failure. There are two key ways these relationships manifest themselves in policy.

CONFERENCE PAPER
"A National Level Conference on Recent Trends in Information Technology and Technical Symposium" On 09th March 2013
Organized by
Dept. of IT, Jawaharlal Darda Inst. Of Eng. & Tech., Yavatmal (MS), India

II. CONCLUSION

In this article we describe the complexity of building a model of the routing of a large AS. We first explain the architecture of an AS and how routing works. Then we describe the essential factors that need to be taken into consideration when building a model of the routing of an AS. We describe CBGP,an open source tool we developed, especially designed to let ISPs play with a model of their network. We illustrate the use of our tool through two different case studies. The first case study studied the impact on the traffic of a transit AS of changing its Internet connectivity. The second one investigated the impact of link failures on routing changes inside the AS. These two case studies have shown the importance of taking into account the interdomain routing information to understand the routing of a large AS.As part of our ongoing work, we are currently applying the model presented herein on the network of a large transit AS.This AS contains hundreds of routers and has an iBGP hierarchy with multiple levels. We are also working on studying the interaction between multiple interconnected ASs. C-BGP can be used to compute the outcome of BGP route selection when there are multiple domains. However, we require knowledge of the structure and policies of the other domains. In order to study the impact of changes in one domain on its inbound traffic, for instance, we need to have knowledge of nearly all.

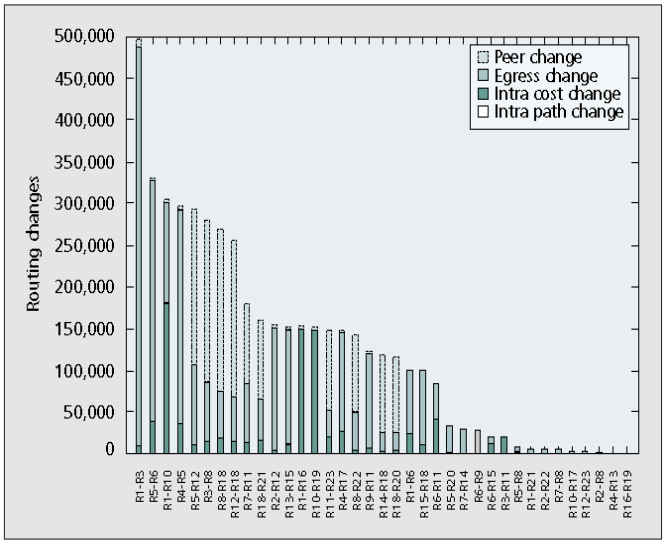

Figure 5. Single link failure analysis: impact on BGP.

The Internet domains. We are currently working on building a model of the Internet that can be used for this purpose.In addition, we are still evolving our tool. The first improvement we are working on concerns a more accurate model of the IGP through support of multiple areas. The second improvement consists of operating the model on a continuous feed of topology, routing data, and traffic data. We believe that our approach to integrate topology, routing data, and traffic data can serve ISP operators to better understand the behavior of an AS and help them investigate improvements in the design of their network. Although BGP policies can be highly complex, there are a number of common design patterns that are typically used by ISPs. In this article we discuss several common patterns and how they can be realized using BGP policy mechanisms. We believe that by recognizing these patterns, we can more efficiently develop tools that directly support them, such as analysis tools that check correctness, languages that preclude errors, or architectures designed for common cases.

III. REFERENCES

[1]. Y. Yang et al., "On Route Selection for Interdomain Traffic Engineering," IEEE Network, this issue.

[2]. N. Feamster, J. Winick, and J. Rexford, "A Model of BGP Routing for Network Engineering," Proc. ACM SIGMETRICS, June 2004.[3]M. Caesar et al., "Design and Implementation of a Routing Control Platform," Proc. Networked Sys. Design and Implementation, May 2005.

[3]. D. O. Awduche, J. Agogbua, and J. McManus, "An Approach to Optimal Peering between Autonomous Systems in the Internet," Proc. IEEE ICCN '98, Oct. 1998.

[4]. W. B. Norton, The Art of Peering: The Peering Playbook, preprint available from wbn@equinix.com, May 2002.

[5]. J.-P. Vasseur, M. Pickavet, and P. Demeester, Network Recovery: Protection and Restoration of Optical, SONET-SDH, and MPLS, Morgan Kaufmann, 2004.

[6]. B. Halabi and D. Mc Pherson, Internet Routing Architectures, 2nd ed., Cisco Press, Jan. 2000.

[7]. L. Gao, "On Inferring Autonomous System Relationships in the Internet," IEEE Global Internet, Nov. 2000.

[8]. A. Feldmann and J. Rexford, "IP Network Configuration for Intradomain Traffic Engineering," IEEE Network, Sept./Oct. 2001, pages 46–57.

[9]. A. Feldmann et al., "Netscope: Traffic Engineering for IP Networks," IEEE Network, Mar. 2000.

CONFERENCE PAPER
"A National Level Conference on Recent Trends in Information Technology and Technical Symposium" On 09th March 2013
Organized by
Dept. of IT, Jawaharlal Darda Inst. Of Eng. & Tech., Yavatmal (MS), India