Mapping the Philippine Internet using the Border Gateway Protocol

Gino LV. Ledesma

I. Abstract

As the Internet continues to grow and its uses continue to increase, it has become of utmost importance for Internet Service Providers (ISPs) to provide quality service to end users. More and more applications require better quality of service, low latency and redundancy of Internet connections. With this in mind, analysis tools must be provided to be able to assist in monitoring these Internet links.

This project has developed a tool for visualizing the current state of interconnectivity among the different ISPs in the Philippines. The tool provides access to a visual map on the Internet via a website. It derives the map from the Global Internet Routing Table using the Border Gateway Protocol (BGP).

II. Introduction

The Internet has grown from a relatively small network used for military and academic purposes to a huge, globally interconnected network which many businesses rely on. As the Internet continues to grow, Internet Service Providers (ISPs) interconnect with each other in order to form larger networks. Peering gives the ISPs and their clients better and faster access to information and resources found in another network by allowing traffic to utilize the most efficient path. A high level of peering not only provides performance improvements due to multiple possible paths of travel, but also redundancy, allowing transactions to proceed even when individual connections go down.

Visualization of the Internet has been a favorite project by many ISPs, mainly to provide them with data that will help them in making strategic decisions, such as whether to deploy more resources to increase bandwidth or capacity, peer with certain service providers to speed up data transfer with remotely located networks, and the like.

There are projects which have already done this, some in a smaller scale, while others on a larger scale. The University of Oregon's Route Views project (Website at `http://www.routeviews.org/`, spearheaded by David Meyer, provides real-time information about the global routing system from the perspectives of multiple backbones across the globe. The Cooperative Association for Internet Data Analysis (CAIDA) has several tools, among them Skitter and MapNet, both of which provide visualization of network connectivity. Yet others use the Looking Glass Collections such as Merit, DTI NSPIXP-2 Looking Glass and IP Plus Looking Glass Server.

A. Problems and Challenges

Visualization of the Internet has become an increasingly challenging task as the number of Internet service providers grow and rapidly adapt their networks according to bandwidth and business requirements. Commercial Internet service providers are also becoming hesitant in sharing peering information or topology about their networks.

While Internet visualization has some support from the research community in the United States and Europe, there has been little effort in the country to engage in this type of research. A few years ago, a small number of software developers and network engineers in the Philippines set off to map the Philippine Internet and have managed to compile an autonomous system number (ASN) list detailing each ASN that belongs to various service providers in the country[10], as well as an IP Block Listing which each service provider manages[11]. However, no further development has been done to visualize this. This information is made available at the Philippine Internet Resources Website, maintained by William S. Yu and Miguel A. Paraz, at `http://cng.ateneo.net/cng/wyu/dev/ph-isp/`. Only metrics related research has been published so far[12][13] and little visualization efforts.

B. Motivation and Significance

One of the on-going major projects is the Internet Mapping Project, headed by Bill Cheswick and Hal Burch, both of Lumeta Corporation. They have been gathering data on Internet connectivity since mid-1998, hoping to acquire and save Internet topological data over a long period of time. Their data has been used in the study of routing problems and changes, distributed denial of service (DDoS) attacks, and graph theory. The University of Oregon's Route Views project now also provides publicly accessible archives for the global routing system.

There are, however, no available tools online that provide a graphical map of the current state of the Internet routing system, whether global or local. While the Route Views project provides the necessary data, they must be processed by separate post-processing tools to generate the necessary information. Only the skitter, mapnet, and Walrus projects from the Cooperative Association for Internet Data Analysis (CAIDA; Website at `http://www.caida.org/`) come closest to providing real time graphs, although these projects focus more on routing traffic and performance analysis.

Moreover, most projects deal with the global Internet scope. The problem with this is that the global routing table does not provide the necessary detail for fulfilling our objects for the tool. Philippine routes tend to be dwarfed by their foreign counterparts. By providing an Internet map of the country, Internet service providers are better able to make use of the information that is relevant to them. The obvious contribution of this map is to provide Internet service providers with peering information. This study can also verify the assumptions obtained in a report on Philippine Internet peering[15] submitted to the Asia and Pacific Internet Association (website at `http://www.apia.org/`). This can help network administrators and managers in acquisition of additional bandwidth, capacity planning and peering studies. The ASN and IP Block Lists will serve as the constraints of the map, showing only those major service providers which need to be on the map and their immediate neighbors.

III. Scope and Limitations

The goal of this particular project is to provide real-time information about the Philippine Internet routing system from the perspectives of major backbones and locations within the Philippines only. The desired output is a visual representation of the Philippine routing system which is updated frequently (e.g., intervals of 5-10 minutes between updates), or better yet, on-demand. While the above-mentioned projects (e.g. CAIDA skitter) provide visual network connectivity, the Philippine network map is either not part of their global map (due to the sheer size of American and European networks) or not included for data analysis altogether. By focusing on a smaller subset of the global Internet, up-to-date, relevant information can be obtained faster. The project as such aims to provide relevant information to service providers for service providers. By focusing on the Philippine subset, service providers are better equipped at identifying key points in the network for improvements or expansion. By knowing the level of peering, service providers are better equipped at making decisions as to increase peering levels where bottlenecks are encountered or decreasing peering levels to reduce costs without altogether sacrificing performance and reliability.

IV. Related Literature

Internet visualization is an on-going research in other parts of the world. CAIDA is a collaboration of different networking groups throughout the world in doing Internet performance metrics and visualization. One of their projects is the Macroscopic Internet Infrastructure Visualization and Management Tool, called MapNet for short. MapNet provides a visualization of the Internet peering within a geographical context.[3]

The French Network Operators Group (FrNOG) have a similar project, although instead focus on graphing connectivity through major Autonomous Systems (ASs). Their Graphical AS Path (GASP) utility providers users a tree of peers between one host to another.[6] GASP focuses on selectively displaying peering between a host with other major service providers, rather than full peering within each host.

There are similar other projects, but most of them are not very relevant to the Philippine context. CAIDA's MapNet, in particular, does not show the peering information in the Philippines because other neighboring countries dwarf the density of connectivity in the Philippines.[1] FrNOG's GASP, on the other hand, displays a tree that is not easily understandable in a wider context, and there is no easy way to compare peering between two service providers.

V. Methodology

Most Internet visualization projects, including the Internet Mapping project, makes extensive use of utilities such as `traceroute` and `whois`.[1] The `traceroute` utility in particular shows the routes from one host to another. The `whois` utility is used to look up information, such as owner, name, and the like, on each particular gateway. These utilities, however, have no control in specifying which network path to take. And as major service providers have numerous routes to different networks, the graph is only ideal if multiple traces are done from different locations.

Similar to Route Views, this project visualizes the Internet by scanning the Autonomous System (AS) address space. Doing so will require the use of the Border Gateway Protocol (BGP), as documented by the Request For Comments article 1771 (RFC1771)[2]

By using BGP, the project can obtain up-to-date data and process it accordingly. BGP also allows the project to easily focus on the Philippine subset of the Internet, without having to query the entire Internet, filter the addresses according to prefixes designated as local, and generate the necessary information accordingly. BGP also allows for multiple installations (routers) in different locations. That is, a central data-collection and analysis tool can query multiple BGP-enabled routers from different ISPs for different "perspectives," correlate them, and provide a unified view.

The project uses the GNU Zebra routing software (website and software avaiable at `http://www.zebra.org/`. By configuring the `bgpd` component of the software to act as a tool that queries other networks for peering information, routing data was collected, filtered, and transformed into a graph that can be plotted by different visualization tools. Due to the nature of BGP, the graphs can be updated dynamically, either periodically or on-demand. The graphing tools components are mainly drawn from the works of other free tools on the Internet. Filtering, processing and consolidation tools were developed as part of this project.

The goal was to be able to present a graphical map, along with raw data, made available to the Internet through the world wide web. A web browser is all that's needed for users to see an up-to-date version of the map. The peering

[1] MapNet in particular graphs over 8,000 major exchange points in the world.

[2] BGP is an inter-Autonomous System routing protocol. The primary function of a BGP speaking system is to exchange network reachability information with other BGP systems. This network reachability information includes information on the list of Autonomous Systems (ASs) that reachability information traverses. This information is sufficient to construct a graph of AS connectivity from which routing loops may be pruned and some policy decisions at the AS level may be enforced.[9]

information gathering and graph processing is done in the background.

VI. Configuration

A. Architecture

The project is currently hosted at the Ateneo de Manila University Campus Network Group Network Operations Center (Ateneo CNG NOC). The Loyola Schools' backbone router acts as a Route Reflector to provide (reflect) routes to a collector tool. This collector tool, in turn, saves the data that will under go filtering, analysis, and processing. This data will then be plotted into a viewable graph. BGP Data can also be obtained from Mozcom Internet's Core Router.

B. Backbone Router / Collector Box Configuration

Detailed configuration:

Route Reflector	Ateneo de Manila University Router
IP Address	202.138.180.122
AS Number	18188

```
! Backbone Router -- Route Reflector
router bgp 18188
  network 202.138.180.0
  neighbor 202.138.180.138 remote-as 18188
  neighbor 202.138.180.138 description Netmapper
  neighbor 202.138.180.138 version 4
  neighbor 202.138.180.138 route-reflector-client
```

Route Reflector Client	CNG NetMap
IP Address	202.138.180.138
AS Number	65432 (Internal), 18188

```
! CNG-NetMap PH IX RouteViews
! Gino LV. Ledesma
! gledesma@ateneo.edu
! Version: 20030212-01

! General Configuration
hostname cng-netmap.ateneo.net
password xxxxxxxx
enable password yyyyyyyy

log file /var/log/zebra/bgpd.log
debug bgp events
debug bgp updates
line vty

! Do not update the kernel routing table
bgp multiple-instance

router bgp 18188 view CNG-NETMAP
bgp router-id 202.138.180.138

! Members
neighbor 202.138.180.126 remote-as 18188
neighbor 202.138.180.126 description AdMU-Loyola
neighbor 202.138.180.126 transparent-as
neighbor 202.138.180.126 transparent-nexthop
neighbor 202.138.180.126 soft-reconfiguration inbound

!Logging
dump bgp updates
     /data/bgpd/%Y.%m/UPDATES/updates.%Y%m%d.%H%M 5m
dump bgp routes-mrt
     /data/bgpd/%Y.%m/RIBS/rib.%Y%m%d.%H%M 5m
```

C. Visualization Tools

The data collected and stored by Zebra bgpd is processed by a variety of Perl scripts and Java applications that convert this data from MRT-formatted route dumps to data sets which can be graphed by a number of visualization tools, including AT&T Research Labs' GraphViz[3][7], Tamara Munzner's H3Viewer library and Hypviewer[4][14], and CAIDA's Otter and Walrus[5][5] among others.

Currently, the project is hosted on a server in the Ateneo de Manila Campus Network Group's Network Operation Center (Website at `http://cng-netmap.ateneo.net/`). The website provides the current map of the Philippine Internet Routing table that is updated every 15 minutes.

VII. Results and Observations

There are currently three (3) major Internet exchange points in the Philippines:

1. Manila Internet eXchange (MIX)
2. Philippine Internet eXchange (PHIX)
3. PHNET Common Routing Exchange (CORE)

These different Internet Exchange points in the country have different clients. From a normal Internet user's point of view, access to local or foreign content will depend on the peering level of the user's ISP with the destination server the information is hosted on. In the past, if two users were on different Internet Service Providers, chances are that sending information between these two users would require the need for the requests to pass through International links, most often to the United States. This obviously incurs costs and communication delay.

Presently, the three exchange points in the country have addressed, to some extent, the problem of high utilization of International Private Lines (IPLs) for local traffic. Unfortunately, it is still common for many organizations or companies to buy bandwidth or connectivity from the US, sometimes hosting their content in the US because it is closer, and probably faster, in terms of Internet measurement, than when hosting content locally.

As the number of companies, organizations, and institutions go online, a higher level of local peering will allow everyone to enjoy better access to local content, and perhaps encourage businesses and other institutions to host their content locally. This, as a result, can lessen costs on the use of International lines for obtaining information that

[3] A graphics drawing library and application suite capable of generating directed/undirected acyclic graphs

[4] A 3D hyperbolic spanning tree visualization tool

[5] a Java-based 3D plotting and visualization tool for moderately-sized graphs

Fig. 1. Philippine Internet Peering (Generated by GraphViz)

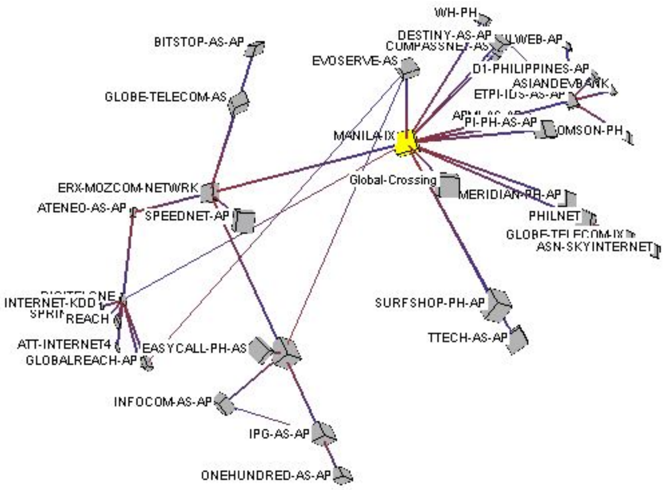

Fig. 2. Philippine Internet Peering (Generated by Hypviewer)

should have been locally hosted in the first place. A higher level of local peering will also give us not only faster access to local content, but a more reliable network infrastructure as well.

Included are some of the snapshots of the Philippine Internet peering density.

These graphs are available at the NetMapper project website at `http://cng-netmap.ateneo.net/`. Additional tools will be provided to provide plotting for other visualization tools.

The graphs indicate a moderate level of peering among the major ISPs in the country (e.g. Mosaic Communications, DigitelOne, and Eastern Telecoms), especially at the Manila Internet Exchange (MIX) point. The other two major Internet Exchange points in the country, namely the Philippine Internet Exchange (PHIX) and PHNet CORE, have a lesser degree of peering.

The Manila Internet Exchange point (MIX) has the highest level of peering, whereas PHIX and PH-CORE are sparse. It can also be inferred from the graphs that, while the peering level in the country is moderate, there are still certain nodes that are distant from the exchange point. One such inferrence is that Sky Internet, one of the major ISPs in the country, is considered far from other major ISPs, including DigitelOne and Mosaic Communications. The visual graph makes it easier to identify fringe (or distant / leaf) nodes from the core, and makes it easy to detect a high level of peering.

VIII. Conclusion

One the problems encountered in the project was the partial table set hosted by the Ateneo de Manila University backbone router due to hardware constraints. This partial table gave, at times, a smaller view of the Philippine Internet routing table. This problem was addressed when another router exported its own routing table to the route reflector client. This shows that multiple route exports allows for a more comprehensive view of the Philippine Internet routing system.

The project has potential for growth. Among the recommendations the researchers propose is providing a three-dimensional visualization tool that can be used online to view 3D versions of the graphs. A three-dimensional view can be more tolerant of visual clutter and occlusion than a two-dimensional graph. Moreover, the project can be improved to provide an almost live, interactive graph.

The project, in conclusion, has met the goals it set to achieve, that is, of providing an updatable visual map of the Philippine Internet routing table for use by researchers and Internet Service Providers. The results obtained from the initial set of graphs also supports the findings of M. A. Paraz and W. S. Yu's findings on Philippine IX Peering. It is hoped that the project will be carried out further for other Internet-related research studies in the Philippines.

Acknowledgement

The author would like to thank his adviser William Emmanuel S. Yu who has provided the necessary background expertise and an introduction to IP engineering research. The author would also like to thank the Department of Information Systems and Computer Science, particularly the chairperson Luis G. Sarmenta, PhD. for providing guidance and feedback for this project. The author would also like to thank the Ateneo de Manila University and Campus Network Group for providing the necessary resources to carry out the project.

References

[1] Cheswick, B. and Burch, H. Internet Mapping Project. Accessed 5 December 2003; available from `http://research.lumeta.com/ches/map/index.html`; Internet.

[2] John S Quarterman's Website. Accessed 5 December 2003; available from `http://www.quarterman.com/jsq/articles.html`; Internet.

[3] Cooperative Association for Internet Data Analysis. Macroscopic Internet Visualization and Measurement. Accessed 8 December 2003; available from `http://www.caida.org/tools/visualization/mapnet/`

[4] Cooperative Association for Internet Data Analysis. Skitter. Accessed 8 December 2003; available from http://www.caida.org/tools/measurement/skitter/
[5] Cooperative Association for Internet Data Analysis. Walrus - Graph Visualization Tool. Accessed 8 December 2003; available from http://www.caida.org/tools/visualization/walrus/
[6] French Network Operators' Group. Graphical Autonomous System Path. Accessed 12 December 2003; available from http://mogwai.frnog.org/sysctl/gasp/
[7] AT&T Research Laboratories. GraphViz - Open Source graph drawing software. Accessed 20 February 2003; available from http://www.research.att.com/sw/tools/graphviz/
[8] Matrix Maps Quarterly. Accessed 5 December 2003; available from http://www.mids.org/mmq/; Internet.
[9] Y. Rekhter and T. Li, Ed. A Border Gateway Protocol (BGP-4). March 1995. Accessed 15 December 2003; available from http://www.faqs.org/rfcs/rfc1771.html
[10] M.A. Paraz and W.S. Yu. 2002. Philippine IP Block Listing. Accessed 5 December 2003; available from http://cng.ateneo.net/cng/wyu/dev/ph-isp/client/PH-IP.txt; Internet.
[11] M.A. Paraz and W.S. Yu. 2002. Philippine ASN Listing. Accessed 5 December 2003; available from http://cng.ateneo.net/cng/wyu/dev/ph-isp/client/PH-IP.txt; Internet.
[12] M.A. Paraz and W.S. Yu. 2002. "Measuring Performance of IP Networks Using Content Metrics". Philippine Journal of ICT and Microelectronics, Volume 2, Aug 2002.
[13] M.A. Paraz and W.S. Yu. 2002. "Philippine Internet Content Performance Metrics". Sixth International Symposium on Parallel Architectures, Algorithms, and Networks (ISPAN2002), May 2002.
[14] T. Munzner. 2000. "Interactive Visualization of Large Graphs and Networks". Ph.D. dissertation, Stanford University, June 2000.
[15] M.A. Paraz and W.S. Yu. "Philippine IX Information". Asia and Pacific Internet Association Newsletter No. 8, January 2002.