A Study of ISP Networks at Indian IXPs

Anshuli Patil*, TSRK Prasad†, Neena Goveas‡ and Bharat Deshpande§
Department of Computer Science & Information Systems
BITS Pilani - K K Birla Goa Campus, Goa, India - 403726
{*f2011071, †tsrkp, ‡neena, §bmd} @goa.bits-pilani.ac.in

***Abstract*—BGP Path advertisements accumulated by Internet exchange points (IXPs) can provide useful insights into the autonomous system (AS) interconnectivity in the geographic region served by an IXP. We utilize BGP path advertisements data from national Internet exchange of India (NIXI) - the only IXP operator in India, to draw AS topology graphs. Our software can generate AS path trace from any Indian source IP network to any Indian destination IP network. We also use NIXI data to illustrate the multilateral peering fabric of ISPs by running Dijkstra's shortest path algorithm on the generated AS topology graph. We are able to generate a mapping from IP address to autonomous system based on NIXI data. Work done as part of this project enables operators to view national interconnection topology of their network and provide network diagnostics in an timely manner.**

I. Introduction

Internet exchange points (IXPs) are an essential part of the Internet infrastructure that supports interconnections of 30,000 active autonomous systems across the globe. IXPs allow Internet service providers (ISPs) to exchange Internet traffic between their constituent autonomous systems and implement business relationships to control flow of traffic in AS ecosystem. IXPs reduce the portion of an ISP's traffic which must be routed via their upstream transit providers, thereby reducing the average per-bit delivery cost of their service. Furthermore, the increased number of paths learned through the IXP improves routing efficiency and fault-tolerance [1].

Border Gateway Protocol (BGP) is an exterior gateway protocol designed to exchange routing and reachability information between autonomous systems (AS). Figure 1 shows a representative topology for an AS. Border routers placed on the edge of an AS cloud speak exterior BGP (eBGP). These border router maintain peering sessions with the border routers of neighbouring ASes. A unique AS Number (ASN) is allocated to each Autonomous System for use in BGP routing [2]. The routing decisions at BGP level does not involve traditional cost metrics, but are based on AS path, network policies and/or rule-sets configured by a network administrator.

Figure 2 represents typical architecture of a layer-2 IXP. Layer-2 IXPs provide a switching fabric and each of the member ASes connects its' access router to the switching fabric. When a pair of member ASes decide to peer at the IXP, they establish a BGP session between their access routers which, in turn enables the exchange of IP traffic over this peering link [3]. We can make useful inferences about the topology of autonomous systems by analyzing BGP path advertisements known to IXPs.

National Internet eXchange of India (NIXI) [4] is the lone IXP operator in India. NIXI is a non-profit company established in 2003 to provide neutral Internet Exchange Point services in India. Like most IXPs today, NIXI also gives public access to its looking glass [5] data. We used data provided by NIXI looking glass for generation of AS topology graphs, AS path trace and network to AS mapping.

This abstract is organized as follows. Section II contains details of our data gathering process. Section III provides an overview of AS topology graphs that can be generated using NIXI routing base. Section IV details use of Patricia trie to map a given IP address or IP network to the corresponding AS. We summarize the results of this study in section V. We

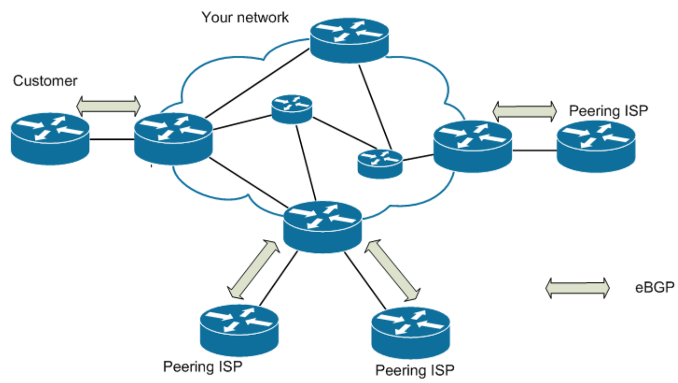

Fig. 1: BGP deployment scenario inside an autonomous system. Routers shown on the edge of network cloud speak eBGP with AS border routers of neighbouring ASes.

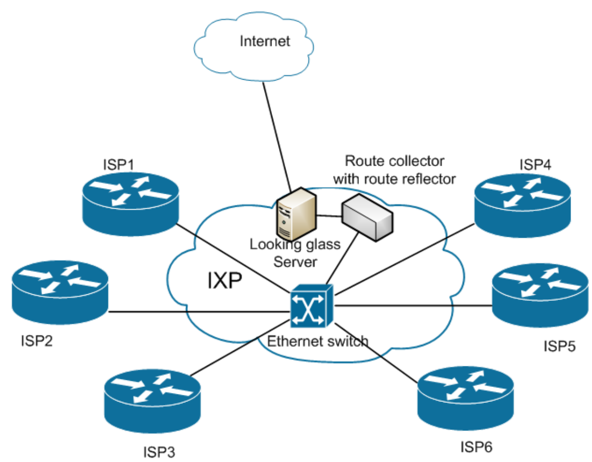

Fig. 2: Architecture of a typical layer-2 IXP. Looking glass server provides read access to eBGP routing tables accumulated at route reflector.

conclude with a discussion on possible future work in section VI.

II. NIXI Routing Base

NIXI provides MRTG statistics for all their IXP locations [6]. We chose four IXP locations, namely NIXI-Mumbai, NIXI-Noida, NIXI-Kolkata and NIXI-Hyderabad for our analysis. The IXP locations have been chosen based on their traffic profiles. An estimate of the daily average traffic over thirteen month period (Oct, 2013 to Oct, 2014) for the four chosen IXP locations is shown in Table I. The IXPs at these four locations represent different AS profiles in terms of an IXP's average daily traffic.

TABLE I: Estimate of average daily traffic at various NIXI locations. The period of observation is from Oct, 2013 to Oct, 2014.

IXP Location	Range of average daily traffic
Mumbai	5 – 9 Gbps
Noida	1.6 – 4 Gbps
Kolkata	23 – 80 Mbps
Hyderabad	11 – 44 Mbps

We query NIXI looking glass to build a local database of BGP path advertisements. NIXI provides looking glass service for these four IXP locations. We call the local routing database as *NIXI routing base*. The following command queries looking glass server for data.

```
show ip bgp neighbors 218.100.48.88
advertised-routes
```

The given command lists out the routing and connection data of traffic sent out from the ISP with the address 218.100.48.88. This command produces results which consists of several lines of data. A sample line is shown in Table II.

Destination network is the target network which is being serviced. This destination network is either owned by an AS or by a customer connected to an AS. *Next hop* is the IP address corresponding to the ISP which is the next hop on the AS path of the packet. *ASPATH* is the sequence of ASes that must be traversed in order to reach the destination network. These numbers are part of BGP Path advertisement. NIXI looking glass consolidates the BGP path advertisements and shows them in a tabular form.

TABLE II: Format of data returned by NIXI looking glass server.

Destination Nwk	Next Hop	ASPATH
1.22.206.0/24	218.100.48.71	
	55410 45528 45528 45528 45528 i	

In this study, we have taken one snapshot of looking glass data on a specific date. We use the snapshot for further analysis and visualization.

III. Connectivity Graphs

NIXI routing base allows us to construct AS connectivity graphs and path traces. NIXI allows only ISPs to peer at its' exchange points. Even though data center operators like Tulip Telecom Ltd. (Tulip) and NetMagic Solutions Pvt. Ltd. (Netmagic) also connect to NIXI exchange points using their ISP license, they are a minority. Thus any AS connectivity graphs and path traces generated closely correspond to the ISP connectivity graphs and ISP path traces. Both the connectivity graphs and path traces yield interesting results which are explained in subsections III-A and III-B.

A. AS Topology Graph

The AS numbers of each BGP Path are used for constructing AS topology graphs. AS topology graph summarizes reachability information from these BGP paths. The graph is generated by using the following notions for nodes and edges.

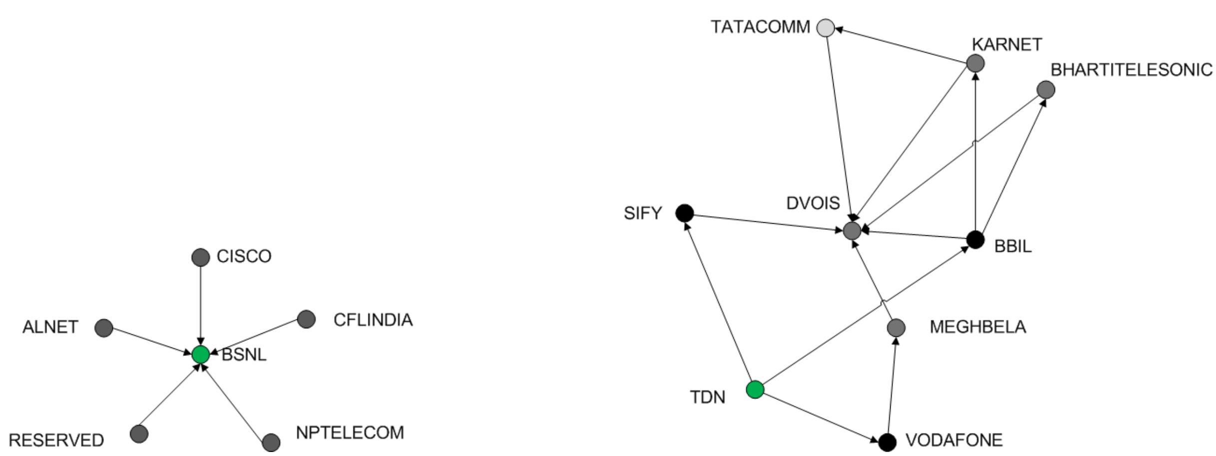

(a) BSNL receive graph

(b) TDN send graph

Fig. 3: AS topology graph generated by considering BGP advertisements sent or received by looking glass server from a ISP. The graph is laid out using force layout algorithm.

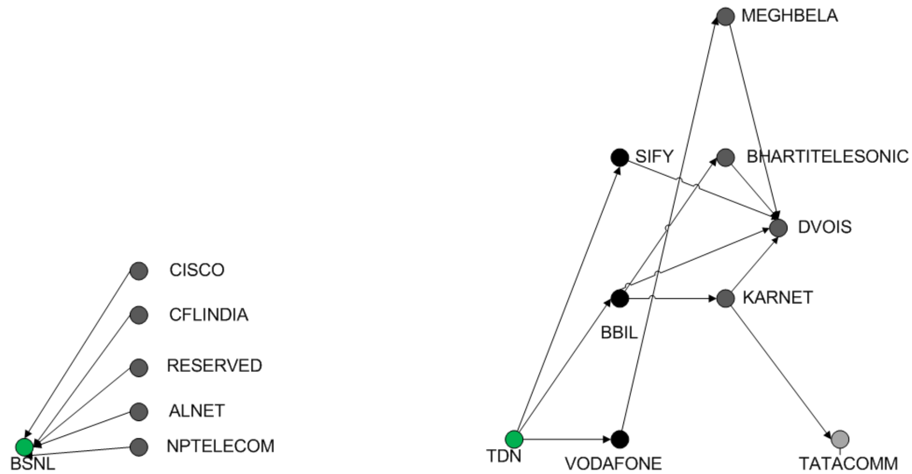

(a) BSNL receive graph

(b) TDN send graph

Fig. 4: Quasi-grid layout of graphs shown in Figure 3.

Nodes — Nodes represent ISPs and are labeled accordingly. Nodes are shaded according their distance from the ISP being studied. Node color moves from dark grey to lighter shades of grey as the hop distance from the ISP of interest increases.

Edges — Each edge represents traffic flowing between two ISPs. The direction of the arrows indicate the direction of packet flow.

Using the graph generation notation explained above, we generate AS topology graphs from NIXI routing base. Figure 3 contains sample topologies generated for two ASes, namely BSNL (Figure 3a) and TDN (Figure 3b). BSNL refers to BSNL Ltd. and TDN refers to Tikona Digital Networks Pvt. Ltd. Graphs in Figure 3 have been laid out using a force-directed layout algorithm. An alternative layout format – quasi-grid layout, is illustrated in Figure 4. Quasi-grid format emphasizes orderly placement of nodes within a range of possible locations in order to reduce edge overlaps in the generated graph.

B. AS Path Trace

In subsection III-A, we detailed a method to create AS topology graph for one AS. We can extend this method to create AS topology graph for all ASes connected at one IXP location. The resulting graph is the AS / ISP connectivity graph as seen at the chosen IXP location. The generated IXP graph is very large and dense. We can use Dijkstra's algorithm to compute the shortest path from any AS to any other AS on the IXP graph. As long as a path exists in the IXP graph, Dijkstra's algorithm would find the path. We can further extend the utility of the detected AS-to-AS path by mapping the detected path to networks advertised by the ASes. We refer to such a network-to-network path consisting of ASes and IXPs as *AS path trace*.

A sample AS path trace generated using Dijkstra's approach is shown in Figure 5. Figure 5 shows a path from source network 1.186.105.0/24 to destination network 120.138.221.0/24. Source network originated with Tata Communications Ltd. (TATACOM) and the destination network

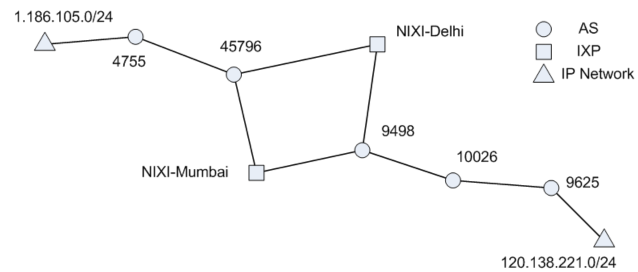

Fig. 5: AS path trace from source IP to destination IP.

prefix is originated by Pacific Internet India Pvt. Ltd. (PACIFIC-INTERNET-INDIA-ASN). The detected path, as shown in Figure 5 consists of autonomous systems 4755, 45796 followed by alternative path segments joined by IXP-Mumbai, IXP-Delhi. The rest of the path segment leading to PACIFIC-INTERNET-INDIA-ASN consists of autonomous systems 9498, 10026 and 9625.

IV. Mapping IP to Autonomous Systems

As part of this study, we implement longest prefix match to obtain a more specific address for a given IP address. In order to implement this algorithm, we use Patricia trie [7]; a sample representation of Patricia trie is shown in Figure 6. The destination networks to Autonomous System mappings are known from NIXI advertisements. After constructing a Patricia trie using all the ASes known, we can map any IP to its' corresponding Autonomous System.

V. Conclusion

BGP path advertisements available at route collector(s) of an IXP prove useful in creation of AS topology graphs and AS path traces. We use one snapshot of BGP routing entries available through NIXI looking glass web service as input for our analysis.

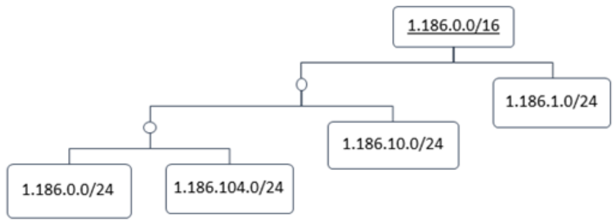

Fig. 6: Patricia trie to store networks.

We implement all the ideas detailed in this abstract into one integrated desktop application called ISPView [8]. The application ISPView, enables a user to understand routing at the ISP level by visualizing the data as a dynamic graph which he/she can explore using the various features of the application. As part of development, we also produced AS topology graph visualizations for each AS. We also created an AS path trace visualization that enables us to see AS-level hops of a network path. A working demonstration of this software is available from [8].

VI. Future Work

We can integrate geographic information into AS topology graphs. Each IXP is at a specific geographic location. Each AS announces a few network prefixes at an IXP location. We can use the geographic location of IXP and the network prefixes announced in that IXP to create a geographic mapping of the Internet. We can integrate ISP topologies from all IXP locations into a country-wide AS specific topology map.

We used one snapshot of NIXI routing base. One snapshot is not good enough to describe a dynamic network such as the Internet; To study dynamic nature of the Internet, we need to take periodic snapshots of NIXI routing base. Such a longitudinal study will help us perform temporal analysis of topology changes.

NIXI also provides MRTG statistics for all IXP locations. A preliminary analysis of these statistics show geographic variation on top of the usual diurnal network activity pattern [9]. There were also instances of unexplained drop in throughput. These issues require further study.

As a next step, we plan to explore data from looking glass servers of IXPs across the world and integrate data from all these sources into our application. This would give us a better understanding of topology of the Internet.

Acknowledgment

The authors would like to thank Indu R for her contribution as developer of ISPView application.

References

[1] J. C. Cardona Restrepo and R. Stanojevic, "A history of an internet exchange point," *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 2, pp. 58–64, Mar. 2012.

[2] J. Houston, "BGP routing tables analysis report," 2014. [Online]. Available: http://bgp.potaroo.net

[3] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a large european IXP," in *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*. ACM, 2012, pp. 163–174.

[4] "National Internet Exchange of India," 2014. [Online]. Available: https://nixi.in/

[5] "Looking glass server," 2014. [Online]. Available: http://nixi.in/lookingglass.php

[6] "MRTG statistics," 2014. [Online]. Available: http://nixi.in/en/mrtg-statistics

[7] "Open source implementation of Patricia trie," 2014. [Online]. Available: https://github.com/rkapsi/patricia-trie

[8] "ISPView," 2014. [Online]. Available: https://github.com/prasadtalasila/ISPView

[9] S. Floyd and V. Paxson, "Difficulties in simulating the Internet," *IEEE/ACM Trans. Netw.*, vol. 9, no. 4, pp. 392–403, Aug. 2001.