Studying Complex AS Relationships Using PoP-Level Connectivity

Yuval Shavitt, Udi Weinsberg, Noa Zilberman*

1 Problem Statement

The Internet has grown significantly in recent years, making it a large and complex communication network. The way traffic flows is determined by the commercial relationships between Autonomous Systems (ASes). Economical forces drive a valley-free routing in the Internet, where packets sent between two ASes travers from the source customer upwards to a shared provider and then downwards to the destination customer. Understanding the AS connectivity and commercial relationship (i.e., Type of Relationship problem) is the focus of several previous works.

However, as ASes become larger and spread over large geographical distances, connectivity breaks down to the Point-of-Presence (PoP) locations, where two ASes can connect in several locations. Moreover, these connections may even lead to complex commercial relationships, e.g., where one AS is the customer of another in a certain location and acts as a peer at a second location. These complex connectivity and commercial relationships is mostly overlooked by existing work, and is the focus of this work.

2 Key Contributions

The contribution of this work is twofold. We first present a structural approach for extracting the connectivity of ASes at the PoP level using a set of IP-level traceroutes and providing them with geographical location. Then, using this PoP-level connectivity graph, we present an algorithm for inferring the commercial relationships between ASes at a finer granularity than the commonly used AS-level graphs, hence providing a first look at the true complex relationships between ASes. This analysis can shed light on the ecosystem of the modern Internet, by explaining the methods used by ASes to interconnect in the modern Internet, and the way PoPs are commonly used to route packets.

3 Methodology

This work is divided into two sections. First, we create a PoP-level map of the Internet connectivity, using a two-phase algorithm, which identify PoPs and then geo-locates them. The extraction algorithm, first suggested by Feldman and Shavitt, looks for bi-partite subgraphs with certain weight constraints in the IP interface graph of an AS; no aliasing to routers is needed. The bi-partites serve as cores of the PoPs and are extended with other close by interfaces. Next, geographical coordinates are assigned for each PoP using information from commercial geolocation databases. We apply a majority vote algorithm that uses a center of weight and range of convergence parameters to overcome inaccuracies, filter out inconsistencies and detect database anomalies.

Using this PoP-level graph, we infer the commercial relationships between ASes at mulitple interconnection points (i.e., different PoPs). For this end, we first find the ASes that comprise the core of the Internet using the k-pruning algorithm, which outputs a small set (less than 100) of globally well-connected ASes that are assumed to be the top-level providers for all routes. The set of IP-traceroutes resolved to PoP-level traceroutes, are followed on the PoP-level graph. Under the valley-free assumption, each traceroute starts with an upwards path of zero or more customer-to-provider (c2p) links towards the core, then at most one side-ways peer-to-peer (p2p) inside the core and then zero or more provider-to-customer (p2c) links downwards from the core. When more than a single hop within the core exists, it means that at least one of the links should be non p2p. Hence, we flag this route as ”non-valley-free” and store the invalid hops within the core. When all traceroutes are traversed, a greedy algorithm is used to label invalid links as p2c or c2p, in a way that maximizes the number of resolved non-valley-free routes.

We further use the geographical information obtained from the PoP level maps to analyze the p2p, p2c and c2p relationships discovered by the algorithm and discuss the implications in broader terms of connectivity, equipment placement and commercial relations.

*Udi and Noa are students, {udiw,noa}@eng.tau.ac.il