Scribe for Internet Routing

Nishant Subramani

February 22, 2018

1 Introduction

In this document we summarize aspects of internet routing and look at some basic properties. First we look at the two planes of networking: the control plane and the data plane. The control plane deals with reachability, namely if a packet can reach its destination. It can deal with identifying a path to make this happen as well. The data plane is a descendant of the control plane dynamics, namely its considered in a post-reachability state. Mostly the data plane is concerned with how to transmit packets as fast as possible along the paths identified by the control plane.

There is interesting interplay between the two planes. The control plane is the primary concern because reachability is fundamental for performance. Topological changes cause the activation of the control plane. The data plane, on the other hand, is more active, activating whether a packet is being sent. Performance is the key aspect that the data plane wishes to focus on.

The control plane only really cares about reachability and so the topology changes in the control side every millisecond or so. The data plane cares about a lot more as mentioned before, so any congestion control and scheduling occurs strictly in the data plane. Changes occur every packet in the data plane which is about 6 orders of magnitude (every nanosecond or so) more often than the control plane.

There are two forms of routing that are considered: within domain routing (IGP) and across domain routing (BGP). Interior Gateway Protocols focus on optimizing a path metric while External Gateway Protocols are concerned with providing reachability information and helping to provide scalable routing policy implementation. This lecture focuses mostly on across domain routing and BGP.

2 Border Gateway Protocol

The Border Gateway Protocol connects all sorts of different autonomous systems together to form the Internet for all intents and purposes. These autonomous systems can be things like campus internets, company internets, and other internet service providers. Autonomous system is a region thats under the control

of one entity (this is to break down the overall space of routers to something manageable to do inference on). Can run things like shortest-path within the autonomous systems relatively effectively and efficiently.

Since TCP provides reliability, BGP runs on top of TCP. There was significant discussion regarding how this is counter-intuitive because these protocols are supposed to run at a lower level than TCP, but since it works so well and its much easier to use this than develop a new protocol, the violation is acceptable.

Two forms of BGP, iBGP and eBGP focus on internal and external autonomous system dynamics. First, eBGP focuses on routes between border routers in different autonomous systems. These are usually connected physically. IBGP focuses on trying to collect all routing information on all the border routers. This way all the border routers are smart and have aggregated routing information. Ideally we want iBGP to tell a internal router the trajectory to get to its destination. It does this by telling the internal router the next major hop, the border router.

3 Transit

There are forms of transit that allow inter autnomous system relationships. Generally transit is defined by a larger autonomous system, sometimes regional, provides support or a service for a smaller autonomous system. This is termed Provider-Customer transit with the larger autonomous system being the Provider and the smaller one being the Customer.

Peering is an alternative to this. Peering has two autonomous systems that both share information and provide services for one another. Routing import and export are based on these forms of transit.

4 Route Import + Export Process

The key point to keep in mind is money. The whole point is money. Money is all that matters. Given this, advertising routes to peer's customers from your own customers is ideal. We use peering to achieve this. Advertising routes to a peer's customers from a provider doesn't make money, so there's no point to doing this. Lastly advertising routes to a provider from non-customers is also suboptimal and doesn't make money. Inter-domain routing is money specific and money centric.

For import, routers are treated as independent agents. Generally they are unconstrained such that they can behave freely. This allows them to act as purely selfish agents, optimizing for their self-interest. Again monetary properties are key. There is a significant preference to make money over all else. Due to this, you prefer customers routes first, peers routes second, and providers routes last.

5 Misc Things

Routing is challenging. One reason for this is that there could be many paths that satisfy a input, output pair. The problem of distinguishing between paths is important such that the network has predictable behavior and is able to exploit some good properties of networks. This is where the multi-exit discriminator comes into play. MED gives a heuristic to choose a specific path over another. There are no guarentees that MED will be explicitly followed, but in expectation it is supposed to be. Incestives are often put into place such that following MED optimizes cost.

Another important concept is multi-homing. Redundancy is important, so being able to have primary and backup routes to traverse is important. BGP has hacks with which certain paths are inflated with extra autonomous systems to make a path a backup. Clearly this adds significant overhead and makes scalability difficult but makes the system more robust to failure.

6 Gao-Rexford Paper

Stability is the key aspect this paper hones in on. If routing oscillates and is unstable, a whole host of odd behavior can materialize. Protocols can be affected in wildly different ways at different time points based on oscillations, so low variance settings are preferred. Safety and Stability are two cornerstones the paper looks at defined as follows:

1. Safety: Ability to get to a stable state regardless of any parameter of the routing process.
2. Stability: Every autonomous system has reached a state after which any operations (exports or imports) keep it in that state.

These cornerstones allow the network to be predictable and robust, the goals that the paper focused on achieving. One meta point is that theory seems to really be implementable in practice here. It shows why certain guarentees are possible and provides the application of this in a succinct manner, something you don't see in most application areas.

7 Issues with BGP

In this section we look at some issues with BGP. Authentication of advertisements is an issue that if misused can lead to things like the YouTube outage in the late 2000s. BGP does not address multi-application communication. Another vital practicality issue is convergence time. In a variety of disciplines convergence time is a bottleneck and this seems to be similar here. There is no theoretical justification about the order of convergence time, just that it will in the limit.