A Declarative and Expressive Approach to Control Forwarding Paths in Carrier-Grade Networks – Public Review

Ming Zhang
Microsoft Research
mzh@microsoft.com

Large carrier-grade network typically comprises hundreds of routers and provides transit service to more than half million Internet prefixes. A common need of such networks is to make efficient use of the network resources while meeting the traffic engineering objectives specified by network operators (e.g., minimizing the maximum link load and/or bounding the delay of certain class of traffic). Existing approaches cannot satisfy the need because they are either too constrained (e.g., shortest-path routing) or face severe scalability problems (e.g., MPLS-based traffic engineering).

Motivated by recent work on software defined networking (SDN), this paper introduces a two-layer architecture to achieve flexible and scalable traffic engineering. At the high-level optimization layer, it provides a set of interfaces realized in a Scala domain specific language that allow network operators to define a wide range of traffic engineering objectives and constraints. At the low-level connectivity layer, it translates the traffic engineering objectives and constraints into a representation in the middlepoint routing model that can be directly implemented using segment routing.

The most appealing aspect of this work is that it represents one of the early papers on segment-routing-based traffic engineering. Given that segment routing is gaining popularity in carrier-grade ISPs, it is rapidly becoming an area that is worth paying attention to.

This work differs from earlier wide-area SDN traffic engineering efforts (notably B4 and SWAN) in that its target networks are much bigger, e.g., hundreds of nodes *vs.* tens of nodes. As a result, it adopts the middlepoint routing model instead of the tunnel-based model. Compared to the latter, middlepoint routing is less flexible (because it specifies only the middle points not the full path) but requires a significantly smaller number of variables. The evaluations using several real ISP network topology and traffic demand matrices show promising results.

As one of the first papers on segment-routing-based traffic engineering, it can be extended in several ways. For example, the current evaluations largely focus on static traffic demand matrices. Also the use cases of service chaining are somewhat hypothetical. The reviewers are hopeful that this paper will inspire further research in this direction. It will be interesting to see if the benefits of this approach will continue to hold up in more dynamic and realistic settings.