April 2016

Network Research Group

PhD position

Multi-path in networks: detection and evaluation of performance

Advisors :
Cristel PELSSER (pelsser@unistra.fr)
Stéphane CATELOIN (cateloin@unistra.fr)
Pascal MÉRINDOL (merindol@unistra.fr)

ICube laboratory – UMR 7357, CNRS – Université de Strasbourg, France
Network Research Group – http://icube-reseaux.unistra.fr

Background

Multipath routing relies on a set of algorithms and techniques to setup and send traffic along multiple paths between a source and a destination. This enables to better balance the traffic load in the network in order to prevent congestion, optimize traffic distribution and, to rapidly react to congestions or failures.

Equal Cost Multi-Path (ECMP)[1], Link Aggregation Group (LAG), and level-2 bundling are a few of the deployed multipath routing solutions. With ECMP, traffic is distributed along paths of shortest equal cost. Parallel layer 3 links are aggregated in LAG and load is spread on the different links. In these solutions, the routing algorithm associates multiple outgoing interfaces to a single destination prefix.

Networks spread traffic on a per-flow basis in order to avoid desequencing within a flow which is especially detrimental to TCP flow (see [2]). For this purpose, a hashing function is applied to a subset of the fields in a packet header. Based on the fields considered fine or coarse grained flows are balanced in the network.

The authors of *paris-traceroute* [4] found a relatively limited amount of multipath routing in the Internet. However, their methodology only captures path diversity at the IP layer. Tokyo-ping [5] showed that some effects of multi-paths can be observed when looking at the complete distribution of the RTTs. This work offers a more generic mean to detect multipath routing. Today, there is no study that quantifies the amount of multipath deployment and studies its limitations in terms of performance. Moreover, the consequences of changes in the control plane (path changes, filters...) on the forwarding plane are still mostly unknown even though such changes are frequent [6]. Such effects in a multipath environment are even less known and are most likely different.

Aims

The first objective will be to develop measurement techniques to detect, quantify and characterize multi-path in today's Internet.

To measure the performance observed by applications in a multipath environment and understand the different forms of multi-path is a second objective for this work. For this purpose, the student will develop correlation techniques in order to pinpoint the paths used by a flow and the performance of its paths.

We will subsequently draw conclusions from our measurements. Based on our observations, we will determine whether there are aspects of multipath routing that need to be improved. We may develop techniques to avoid path segments incurring degraded performance. This may require changes in the signaling of routing information, in the computation and use of the paths in order to spread traffic dynamically upon changes in load. The objective is to nevertheless provide a stable routing solution that avoids the negative effects of route oscillations.

The PhD student will elaborate methods to correlate changes in the control plane with performance degradation. If she/he observes that some actions of the network operators are detrimental to the performance, the student will propose efficient network maintenance techniques that are exempt of problems in a multi-path context.

Finally, we will study which properties can make multipath routing protocols more robust to changes than single path environments. Which physical topologies meet robustness requirements? How to design iBGP topologies?

Bibliography

[1] C.Hopps, Analysis of an Equal-Cost Multi-Path Algorithm, RFC2992, 2000.
[2] T.W.Chim, K.L.Yeung, and K.-S.Lui, "Traffic distribution over equal-cost-multi-paths", Computer Networks, vol. 49, no. 4, pp. 465–475, Nov. 2005
[3] Z.Cao, Z.Wang & W.Zegura, "Performance of Hashing-Based Schemes for Internet Load Balancing", In IEEE INFOCOM, 2000
[4] Timur Friedman, Brice Augustin and Renata Teixeira. Measuring multipath routing in the internet. In IEEE/ACM Transactions on Networking,vol. 19, issue 3, pp. 830-840, June 2011.
[5] Cristel Pelsser, Luca Cittadini, Stefano Vissicchio and Randy Bush. "From Paris to Tokyo: On the Suitability of Ping to Measure Latency". In IMC'13: Proceedings of the 2013 Internet Measurement Conference, Barcelona, Spain, October 2013.
[6] Stefano Vissicchio, Laurent Vanbever, Cristel Pelsser, Luca Cittadini, Pierre Francois and Olivier Bonaventure. "Improving Network Agility with Seamless BGP Reconfigurations." IEEE/ACM Transactions on Networking, June 2013.
[7] T. Griffin and G. Wilfong, "On the correctness of iBGP configuration," in Proc. of ACM SIGCOMM, 2002.