Network Research Group September 2015

M.Sc. Internship – Computer Science

Path Changes Impact on TCP Flows

It has been known for several years that TCP flow control behaves badly when the network layer offers several routes to reach a destination. Distribute packets of the same TCP flow on multiple IP routes generally leads to out-of-order delivery, which triggers TCP fast retransmit mechanism and degrade performance even if no packet has suffered the effects of congestion. Although some mechanisms limit these effects and are more suited to recovery [1], it is not desirable in the context of multi-path routing, to perform per-packet traffic balancing.
Instead of it, mechanisms have been proposed to perform a per-flow distribution [2]. A hash function is applied to a set of header fields so that each packet of the same flow is then switched in the same way. From end to end, all packets of the same TCP flow then follow the same route and there is no desequencing.
This remains true when the allocation is static, that is to say as long as the route remains the same throughout the duration of the flow. However, a dynamic load balancing requires that routing decisions are subordinated to available resources in the network. Thus, if there are insufficient resources on a given route, TCP flows can be deflected to another one during their transmission [3].
In the same vein, some reconfiguration proposals such as [4,5] may trigger the use of intermediate paths between the initial state and the targeted one.

This project concerns the study of the impact of route changes on the performance of TCP. In particular, the deflection of a flow to a route with a larger delay should not induce desequencing, while the reverse is not true. As deflections occur, this could lead to use routes with more and more delay, and abandon the shortest ones, at the expense of the network resources. Otherwise, to inflict a rate reduction to short lived or slow flows seems less harmful. Deflections could then be applied to a set of flows for which the overall negative impact is minimized.

The global objective of this study is to identify a set of best practices that could be then incorporated into a local decision-making process regarding an hop-by-hop traffic balancing or routing reconfiguration scheme.

Supervisors

Stéphane Cateloin (cateloin@unistra.fr)
Pascal Mérindol (merindol@unistra.fr)
Cristel Pelsser

References

[1] Kevin Fall and Sally Floyd, "Simulation-based comparisons of Tahoe, Reno and SACK TCP", ACM SIGCOMM Computer Communication Review, Volume 26 Issue 3, July 1996
[2] Z.Cao, Z.Wang & W.Zegura, "Performance of Hashing-Based Schemes for Internet Load Balancing", In IEEE INFOCOM, 2000
[3] 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
[4] 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.
[5] F. Clad, P. Mérindol, S. Vissicchio, J-J. Pansiot and P. Francois, "Graceful Router Updates for Link-State Protocols", in IEEE ICNP, 2013.