March 2014 Network Research Group

Sujet de Thèse

Adaptive Multipath Routing: Collaborative Load Balancing and Topological Characterization

Encadrement :
Fabrice THEOLEYRE (theoleyre@unistra.fr)
Stéphane CATELOIN (cateloin@unistra.fr)
Pascal MERINDOL (merindol@unistra.fr)

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

Context

Routing is one of the keys for controlling network resources. In current IP networks, routing algorithms objective is to find either a single best path between each pair of routers, or several paths with equal and optimal costs (ECMP, [1,2]). More generally, a multipath routing algorithm is an algorithm that allows finding multiple routes between two points on the network, regardless of the constraints on the costs of these paths. Such protocols have been proposed in the literature ([3,4,5,6,7]). At the router level and towards a given destination, the goal is to provide not only one but several outgoing interfaces. The hop-by-hop composition of these alternatives provides a wide variety of routes which are distinct or not.

There are many advantages with multipath routing. On the one hand, calculating multiple paths with anticipation allows to quickly adjusting routes and overcoming failures that may occur in the network. In a distributed environment, routers can make local routing decisions that are faster than a global routing calculation. On the other hand, different paths can be used simultaneously to take advantage of bandwidth available on each of them. The main key is then to decide the ideal distribution of traffic among these multiple routes: this is called adaptive routing. This way, two approaches can be considered. The first is the assumption that one central equipment has the complete knowledge of traffic demands: the load balancing then consists in optimizing an objective function (of overall resource usage, used bandwidth, etc.). This view has been the subject of many studies. The second approach is to assign routing decisions to different actors in the network: they are asked to collaborate together to efficiently distribute the load across the network. The study of solutions within this last approach motivates this thesis.

Once the routers have selected a set of interfaces to forward IP packets, load balancing is a decision process that may be based on indicators which are local or received from other routers.

The nature of these indicators should be adapted according to the context in which they are employed. In particular, with regard to wireless networks, sharing policy may focus on power management or end-to-end connectivity. Sensor networks can also benefit from these mechanisms. In the context of multi-homed mobile devices, performance and cost of the available connections could impose additional constraints. Moreover, in the core network, one may choose to prioritize resources or avoid congestions. Variables relating to the quality of service could also come into consideration, for example by assigning real-time flows to one or more paths respecting specific constraints. The advent of the Internet of Things, which could bring together tens of billions of connected objects in 2020, introduces new challenges: the resilience of networks that should take charge of new types of traffic that remain to be characterized, or new data access paradigms (such as CCN, [8]) that could fundamentally change the current protocol stack.

Performance of these adaptive multipath routing algorithms and protocols naturally depend on the characteristics of the underlying network graph. The presence of specific patterns (providing k-connectivity at link and/or router level) may have an impact on the paths computation and traffic distribution. Software properties of network devices are also a factor that may facilitate or hamper the deployment of these mechanisms. Furthermore, the emergence of SDN [9] can ease the incremental deployment of new routing solutions adapted to modern uses of the Internet [10]. Performances evaluation and graph properties characterization may help to ease the deployment of new routing paradigms.

The topology of the Internet is generally represented by a graph of routers [11] (or sometimes a graph of autonomous systems [12]). The edges of the graph have optionally one or more weights. Interesting properties for routing are linked to degree of nodes [13], diversity of paths, average distances, weights, etc. Beyond these static or algebraic structural properties [14], it would be interesting to evaluate software properties (heterogeneity of equipments and systems, protocols used, etc.) and their dynamics in order to better understand the feasibility of deploying a new routing solution.

Goals

The first objective of this thesis is to focus on the collaboration between routers, which relies on an exchange of messages. These messages can be upstream requests such as throughput reduction ("backpressure messages"), as introduced by Gojmerac in [15]. One can also imagine whether messages allowing the admission of new flows (interesting ideas are developed in Ammar's thesis, [16]), or disseminating information relating to the current conditions of routers (e.g. radio links quality, position and speed of a mobile node, or its battery level).

The qualitative or quantitative nature of the data exchanged, the triggers of these messages, their rate, and the action to be taken upon receipt are all behaviors to define and study. The candidate will have to evaluate the solutions he proposes, using simulation tools. Besides the quality of the distribution, convergence is a crucial point. The performance analysis will determine whether the proposed solutions are constantly evolving in an unstable state or whether tend towards an acceptable steady solution. In this case, the convergence time is an important indicator.

The second objective is relative to the evaluation of software or topological properties of the networks, in order to estimate their ability to sustain these routing and sharing mechanisms. By examining the (adaptive) multipath algorithms and protocols in the literature, it will first be necessary to define what is an efficient routing. Several indicators may be relevant depending on the objective: the amount of routes discovered, their length, their covering (the fact that they are totally or partially disjoint and the resulting maximum flow), the convergence time, the routes oscillations, interferences with the transport layer (in particular, TCP), etc. Link weights and choice of the metric is essential: is it more convenient to set link weights that favor multipath routing, or to get rid of this logic overlay to only rely on the physical topology? Should the weights be dynamically fit? The answers to these questions have to consider the issue of load balancing: it is not necessary for all the routes to be used simultaneously. Can we link a number of indicators related to the topology (nodes degrees, distribution of these degrees, k-connectivity, etc.) to those that characterize the efficiency of routing? Is it possible, knowing the topological characteristics of a network, to say what routing algorithm/protocol will be the most effective? Reciprocally, can we infer graph models and software to promote multipath routing and load sharing effectiveness? Are there structural patterns and software properties that affect the establishment of multiple routes or, on the contrary, are there patterns that tend to favor it?

References

[1] C.Hopps, *Analysis of an Equal-Cost Multi-Path Algorithm*, RFC 2992, 2000.

[2] 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.

[3] S.Vutukury, *Multipath routing mechanisms for traffic engineering and quality of service in the internet*, PhD thesis, University of California, Santa Cruz, 2001.
[4] X.Yang and D.Wetherall, *Source selectable path diversity via routing deflections*, SIGCOMM'06, volume 36, pages 159-170, October 2006.
[5] Pascal Mérindol, *Routage multichemins par interface d'entrée*. PhD thesis, Université Louis Pasteur, Strasbourg, 2008.
[6] Hao Wang, Haiyong Xie, Lili Qiu, Yang Richard Yang, Yin Zhang, and Albert Greenberg. *Cope : Traffic engineering in dynamic networks*. SIGCOMM Comput. Commun. Rev., 36(4) :99–110, August 2006.
[7] Nithin Michael, Ao Tang, and Dahai Xu. *Optimal link-state hop-by-hop routing*. In ICNP, pages 1–10, 2013
[8] Van Jacobson, Diana K. Smetters, James D. Thornton, Michael F. Plass, Nicholas H. Briggs, Rebecca L. Braynard, *Networking named content*, Proceedings of the 5th international conference on Emerging networking experiments and technologies CoNEXT'09, pages 1-12, 2009.
[9] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. *Openflow : Enabling innovation in campus networks*. SIGCOMM Comput. Commun. Rev.,38(2) :69–74, March 2008.
[10] Stefano Vissicchio, Laurent Vanbever, and Olivier Bonaventure. *Opportunities and research challenges of hybrid software defined networks*. ACM Computer Communication Review (Editorial Zone), 44(2), April 2014.
[11] A. Botta, W. de Donato, A. Pescapé, and G. Ventre. *Discovering topologies at router level : Part II*. In Proc. IEEE Global Telecommunications Conference (GLOBECOM), November 2007.
[12] Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz. *Towards and accurate AS-level traceroute tool*. In Proc. ACM SIGCOMM, August 2003.
[13] P. Mérindol, B. Donnet, O. Bonaventure, and J.-J. Pansiot. *On the impact of layer-2 on node degree distribution*. In Proc. ACM/USENIX Internet Measurement Conference (IMC), November 2010.
[14] Seweryn Dynerowicz and Timothy G. Griffin. *On the forwarding paths produced by internet routing algorithms*. In ICNP, pages 1–10, 2013.
[15] I.Gojmerac, P.Reichl, L.Jansen, *Towards Low-complexity Internet Traffic Engineering: The Adaptive Multi-Path Algorithm*, Journal of Computer Networks, Vol. 52, No. 15, pp. 2894–2907, December 2008.
[16] Doreid Ammar, *Plan de connaissance pour les réseaux sémantiques : application au contrôle d'admission*, PhD thesis, Université de Lyon, December 2012.