

Master's Thesis

Performance Analysis of Transport Layer Protocols in Reconfigurable Data Center Networks

Today's Data Center (DC) networks are facing increasing demands and a plethora of requirements. Factors for this are the rise of Cloud Computing, Virtualization and emerging high data rate applications such as distributed Machine Learning frameworks.

Recently several new architectures have been proposed that rely on reconfigurable optics to evolve the topology over time, such as Helios[1] or RotorNet[2].

As a consequence of these adaptable topologies, forwarding in the network changes. Flows are migrated between different paths in the network. In addition, these paths might be of heterogeneous nature, e.g., in terms of bandwidth or RTT.

At the same time, a variety of transport protocols (variants of TCP, QUIC or new approaches such as NDP) can coexist in a data center network.

Individual or grouped assessments of these protocols consider rather static [3] or more homogeneous scenarios [4].

These are not sufficient to estimate how heterogeneous mix of traffic behaves subject to frequent reconfigurations -- a necessary information to design efficient scheduling algorithms for such reconfigurable topologies.

The goal of this thesis is to evaluate this question with a measurement campaign and to provide a model to estimate the cost (e.g., reduced utilization, throughput, re-transmissions) of frequent reconfigurations depending on the traffic mix.

To this end, first a small scale experiment environment using programmable data planes has to be established. In the second step, this environment is used to conduct measurements, which are the bases for the modeling.

[1] N. Farrington et al., "Helios: A Hybrid Electrical/Optical Switch Architecture for Modular Data Centers," pp. 1–12, 2010.

[2] W. M. Mellette et al., "RotorNet: A Scalable, Low-complexity, Optical Datacenter Network," 2017, pp. 267–280, doi: 10.1145/3098822.3098838.

[3] B. Jaeger, D. Scholz, D. Raumer, F. Geyer, and G. Carle, "Reproducible measurements of TCP BBR congestion control," vol. 144, pp. 31–43. [Online]. Available:

<https://linkinghub.elsevier.com/retrieve/pii/S0140366419303470>

[4] R. Carpa, M. D. de AssuncAo, O. Gluck, L. LefEvre, and J.-C. Mignot, "Evaluating the impact of SDN-induced frequent route changes on TCP flows," in 2017 13th International Conference on Network and Service Management (CNSM). IEEE, pp. 1–9. [Online]. Available: <http://ieeexplore.ieee.org/document/8256021/>

[5] M. Mukerjee, C. Canel, D. Kim, and S. Seshan, "Adapting TCP for reconfigurable datacenter networks," in Proceedings of the ACM SIGCOMM 2019 Workshop on Optical Systems Design - OptSys '19, Beijing, China, 2019, pp. 1–1. doi: 10.1145/3363542.3363545.

Advisors

Andreas Blenk, Johannes Zerwas