Chair of Communication Networks
TUM Department of Electrical and Computer Engineering
Technical University of Munich

Master's Thesis

Adversarial Benchmarking of Data-driven Reconfigurable Data Center Networking

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[3], Projector[5], c-Through[1] or RotorNet[6].
Some of these approaches learn attributes of the traffic distribution in the network to determine topology configurations, e.g., xWeaver[4] or DeepConf[2].
A different stream of research focuses on generating adversarial input to networks or networking algorithms [7,8] to identify weak spots or improve robustness and solution quality of data-driven approaches.
The goal of this thesis is to apply the latter one to reconfigurable topology designs with the main focus on the data-driven variants (xWeaver) to benchmark their performance and to improve their solution quality.
The thesis builds upon an existing flow-level simulator in Python and initial algorithms that generate adversarial inputs for networking problems.

[1] G. Wang et al., “c-Through: Part-time Optics in Data Centers,” presented at the ACM SIGCOMM, 2010, p. 12.
[2] S. Salman, C. Streiffer, H. Chen, T. Benson, and A. Kadav, “DeepConf: Automating Data Center Network Topologies Management with Machine Learning,” in Proceedings of the 2018 Workshop on Network Meets AI & ML, New York, NY, USA, 2018, pp. 8–14, doi: 10.1145/3229543.3229554.
[3] N. Farrington et al., “Helios: A Hybrid Electrical/Optical Switch Architecture for Modular Data Centers,” pp. 1–12, 2010.
[4] M. Wang et al., “Neural Network Meets DCN: Traffic-driven Topology Adaptation with Deep Learning,” Proceedings of the ACM on Measurement and Analysis of Computing Systems, vol. 2, no. 2, pp. 1–25, Jun. 2018, doi: 10.1145/3224421.
[5] M. Ghobadi et al., “ProjecToR: Agile Reconfigurable Data Center Interconnect,” 2016, pp. 216–229, doi: 10.1145/2934872.2934911.
[6] W. M. Mellette et al., “RotorNet: A Scalable, Low-complexity, Optical Datacenter Network,” 2017, pp. 267–280, doi: 10.1145/3098822.3098838.
[7] S. Lettner and A. Blenk, “Adversarial Network Algorithm Benchmarking,” in Proceedings of the 15th International Conference on emerging Networking EXperiments and Technologies, Orlando FL USA, Dec. 2019, pp. 31–33, doi: 10.1145/3360468.3366779.
[8] J. Zerwas et al., “NetBOA: Self-Driving Network Benchmarking,” in Proceedings of the 2019 Workshop on Network Meets AI & ML - NetAI’19, Beijing, China, 2019, pp. 8–14, doi: 10.1145/3341216.3342207.

Advisors

Andreas Blenk, Johannes Zerwas