Shaping and quality of service in SDN controlled hybrid optical/ electrical networks enabled by machine learning

Madeleine Glick, University of Arizona

Massive increases in data center traffic are provoking the exploration and adoption of disruptive technologies in order to maintain acceptable performance with new Big Data applications and the introduction of the Internet of Things. Optical interconnection networks can supply the required bandwidth. However this alone does not solve the challenge of efficiently routing tera to petabits/second of traffic through networks with hundreds of thousands of servers due to cost and switching time. Recent advances in programmable Network Interface Cards (NICs) provide an intelligent edge which can introduce machine learning based predictive traffic analytics; enabling faster, adaptive and accurate traffic classification. This alone would improve throughput and resource utilization, but with the oversight of an intelligent and flexible Software Defined Network (SDN) controller the advantage will be much greater. The goal of our research is to demonstrate significant improvements in performance and the utilization of compute resources by employing efficient machine learning in the programmable NIC and interfacing with the SDN over the optically interconnected or hybrid optical /electrical data center network. Our research has started with improving the algorithms for traffic classification at the edge for accuracy, speed and adaptability using machine learning. A next stage will be to develop SDN protocols and to efficiently exchange the traffic information between the SDN controller and the switches (electrical and optical) that make up the network. The algorithms and control protocols will be applied in our data center testbed using real application traffic to determine trade-offs in latency, utilization, throughput, power and cost.

Our initial results show that we correctly allocate long lived large bandwidth (elephant) flows 50% more accurately and more efficiently than leading heuristics that rely on bandwidth consumption and can inadvertently classify bursty flows as elephants [1].

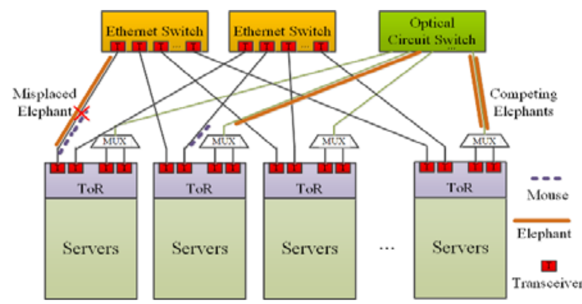

Hybrid electrical/optical data center.
Misplaced elephants reduce performance.

Dr. Madeleine Glick is a senior research scientist at the University of Arizona and the industry liaison officer for the CIAN NSF Engineering Research Center. Dr. Glick is a Research Associate (RA) in the Open Networking Foundation (ONF). She has worked extensively in the field optically interconnected computer networks while at Marconi and Intel Research. She is currently leading a project in collaboration with Netronome aiming to improve optical /electrical data center performance with the implementation of SDN control and machine learning based traffic classification. In a project with Huawei, our group has extended the Southbound OpenFlow 1.3 Protocol and modified the library parser of SDN framework to support optical network control and management. We have demonstrated successful performance with DQPSK and OFDM modulation formats [2]. Dr. Glick is a co-PI on the DOE project TURBO: Terabits/s Using Reconfigurable Bandwidth Optics due to start in January 2016. TURBO will examine both hardware and software control systems to realize end to end multi-Tb/s optical connections in a multi-domain environment. Openflow based optical system controls will be used to provision switches and probe paths in order to create the connections.

References

1. M.Glick et al., Machine Learning Based Adaptive Flow Classification for Optically Interconnected Data Centers, submitted to OFC 2016.
2. M. Yang, et al., "Traffic-Aware Non-Uniform Passband Assignment in Elastic Optical Networks", IEEE Photonics Soc. Conf. (Oct. 2015).