SC21 Network Research Exhibition: Publishable Abstract

HECATE: Towards Self-driving Networks

Mariam Kiran, Nick Buraglio, Scott Campbell

Lawrence Berkeley National Laboratory

{mkiran,buraglio,scampbell}@es.net

For the next generation of DOE data intensive experiments, large data volumes combined with inter-facility cooperation and high performance computing will create demands on network management that are far more complex than those that we face today. Balancing data volume and latency needs while keeping staff and cloud access as performant as possible are some examples of these demands. Software Defined Networking (SDN) allows for programmable networks which is beneficial for controlling some systems, but there are scaling issues for WAN scale solutions as well as slow progress in developing multi-objective path optimizations in software-network systems. In this demonstration, we present HECATE, a stand-alone device that can be plugged into any network infrastructure which provides three main services 1) investigate current application needs and traffic patterns 2) run bespoke data-driven deep reinforcement learning that learn optimal controls to improve traffic engineering, and 3) renders HECATE decisions to SDN and PCE (path computation engine) technologies to bring AI to real networks.

Traffic engineering and path computation techniques such as MPLS-TE (Multiprotocol Label Switching Traffic Engineering), Google's B4 and Microsoft's SWAN (Software Driven WAN) propose manners in which routers can greedily select routing patterns for arriving flows, both locally and globally, to increase path utilization. However, these techniques require meticulously designed heuristics to calculate optimal routes and also do not distinguish between arriving flow characteristics. HECATE exploits data-driven learning to improve traffic engineering decision making. Deep reinforcement learning (RL) programs have been used to teach systems how to drive a car, control massive power grids or self-playing games. With an optimal neural network architecture, a robust deep RL solution can replace controllers in complex environments for model-free optimum control of various complex systems.

Traffic classes and Reward function: HECATE allows users to define their own reward functions. We use 4 options: congestion, latency, loss or bandwidth. Depending on user inputs, HECATE will calculate optimal paths based on what the traffic requirements are. This ties in with the learned traffic classes as well. HECATE will learn the distinct traffic class from the traffic profiles and decide the reward function (Figure 1). We use Bellman equation to optimize reward,

$$Q_t(s,a) = Q_{t-1}(s,a) + \alpha \,(R(s,a) + \gamma \max Q(s',a') - Q_{t-1}(s,a))$$

Eq1: Reward function optimizes based on Bellman equation

Eq1 represents how the reward function is learned as state-action pairs to maximize reward from previous experience. Q-value represents state-action combinations. Better Q-values shows better chances of getting higher rewards which are earned at the end of a complete episode.

Active and Passive Learning: HECATE is packaged with its own monitoring system to feed its AI algorithms. It can perform passive learning on historical network data, but also perform active network probing to learn optimal traffic routes for each traffic seen. This uses a combination of graph neural network and deep reinforcement learning to find optimal network paths.

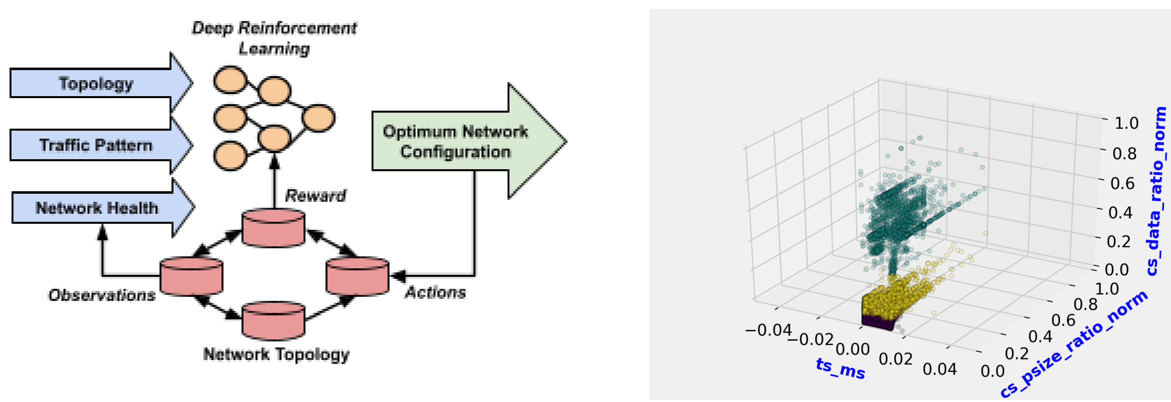

Figure 1: (a) Reinforcement learning for HECATE's Traffic engineering solution (b) examples of traffic classes learned based on time

We incorporate simulation and clustering analytics to learn optimal traffic settings for each kind of traffic. We will demonstrate HECATE in action with a real network setup and linking to a PCE to show rerouting via segmented routing solutions. By utilizing deep reinforcement learning, we allow the Hecate controller to learn from both short term and historical behaviors in the environment about the paths and best hops between source and destination.

[1] Nandini Krishnaswamy, Mariam Kiran, Kunal Singh, Bashir Mohammed, Data-driven Learning to Predict WAN Network Traffic, ACM 3rd International Workshop on System and Network Telemetry and Analytics (SNTA'20)

[2] M Kiran, B Mohammed, N Krishnaswamy, DeepRoute: Herding Elephant and Mice Flows with Reinforcement Learning, 2nd IFIP International Conference on Machine Learning for Networking (MLN'2019)