Karlsruhe Institute of Technology

Autonome Lernende Roboter (ALR)

Prof. Gerhard Neumann

Project Type

- [x] Master Thesis
- [ ] Bachelor Thesis
- [ ] Research Project

Supervisors

Andreas Boltres

andreas.boltres@partner.kit.edu

Niklas Freymuth

niklas.freymuth@kit.edu

Difficulty

Algorithmic

Math

Application

Requirements

- Python and PyTorch
- ML and RL
- *optional:* C++, knowledge in computer networks

Distributed Multi-Agent RL for Routing Optimization

Description

In Computer Networks, Routing Optimization (RO) is about finding optimal (and often shortest) routing paths between nodes. Our most recent work trains an RO policy via Swarm Reinforcement Learning (RL) [2] that uses not only the computer network's topology, but also its current load and performance statistics. However, so far we utilize global network states and rewards, and calculate actions for each routing node in a centralized manner. In this project, we want to find out whether we can achieve comparable performance with a distributed setup, where each routing node assembles its own local view of the network state from its own monitoring as well as information communicated by neighboring nodes.

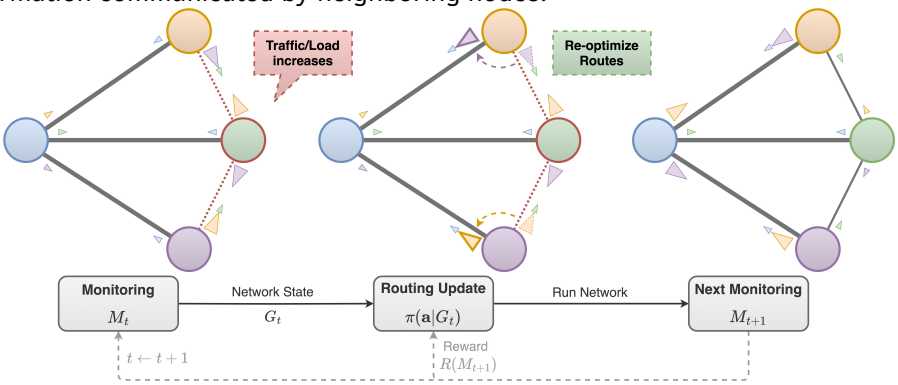

Figure 1: Situation-aware RO adjusts packet routes based on the network topology and current utilization/load to avoid congestion, delay and packet drops.

We build upon our existing framework for RL experiments on RO in the network simulator ns-3 [1] and extend it to support distributed RL algorithms. While popular multi-agent algorithms like MAPPO [3] offer decentralized execution using a shared policy, we will also need to respect communication delays caused by the networked nature of our system [4].

Tasks

- Literature Research: Acquiring a comprehensive understanding of Multi-Agent RL, RL in networked systems and routing optimization; Specifying the research gap we aim to fill; Selecting an algorithm to continue with.
- Implementation of Distributed RL algorithm into our framework; Adjusting components like the observation function or the policy network to handle localized information of varying age.
- Evaluation of the distributed approach in our framework, comparing against centralized RL and heuristic approaches.

References

[1] Thomas R Henderson, Mathieu Lacage, and George F Riley. Network Simulations with the ns-3 Simulator. *SIGCOMM demonstration*, 14(14):527, 2008.

[2] Maximilian Hüttenrauch, Šošić Adrian, and Gerhard Neumann. Deep reinforcement learning for swarm systems. *Journal of Machine Learning Research*, 20(54):1–31, 2019.

[3] Chao Yu, Akash Velu, Eugene Vinitsky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and YI WU. The surprising effectiveness of ppo in cooperative multi-agent games. In *Advances in Neural Information Processing Systems*, volume 35, 2022.

[4] Kaiqing Zhang, Zhuoran Yang, Han Liu, Tong Zhang, and Tamer Basar. Fully decentralized multi-agent reinforcement learning with networked agents. In *Proceedings of the 35th International Conference on Machine Learning*, pages 5872–5881. PMLR, 10–15 Jul 2018.