Autonomous Cellular Network Management

Ahmed Alhassan [1] Merghani Abdelsamad [2] [1]University of Khartoum, [2]University of Khartoum

Abstract

One of the most difficult problems the telecommunications sector is experiencing is operating a radio access network. It is challenging to choose the best configurations for every situation, which may entail load-shedding or site failure, because to the sheer number of base stations and parameters inside each base station. A Radio Access Network also consumes a lot of energy, but with the appropriate setups, this may be decreased. The purpose of this project is to produce scalable software that can replicate a sizable portion of the possible precarious situations, and then build an autonomous system on top of that software that can provide sound advice and oversee the whole base station network.

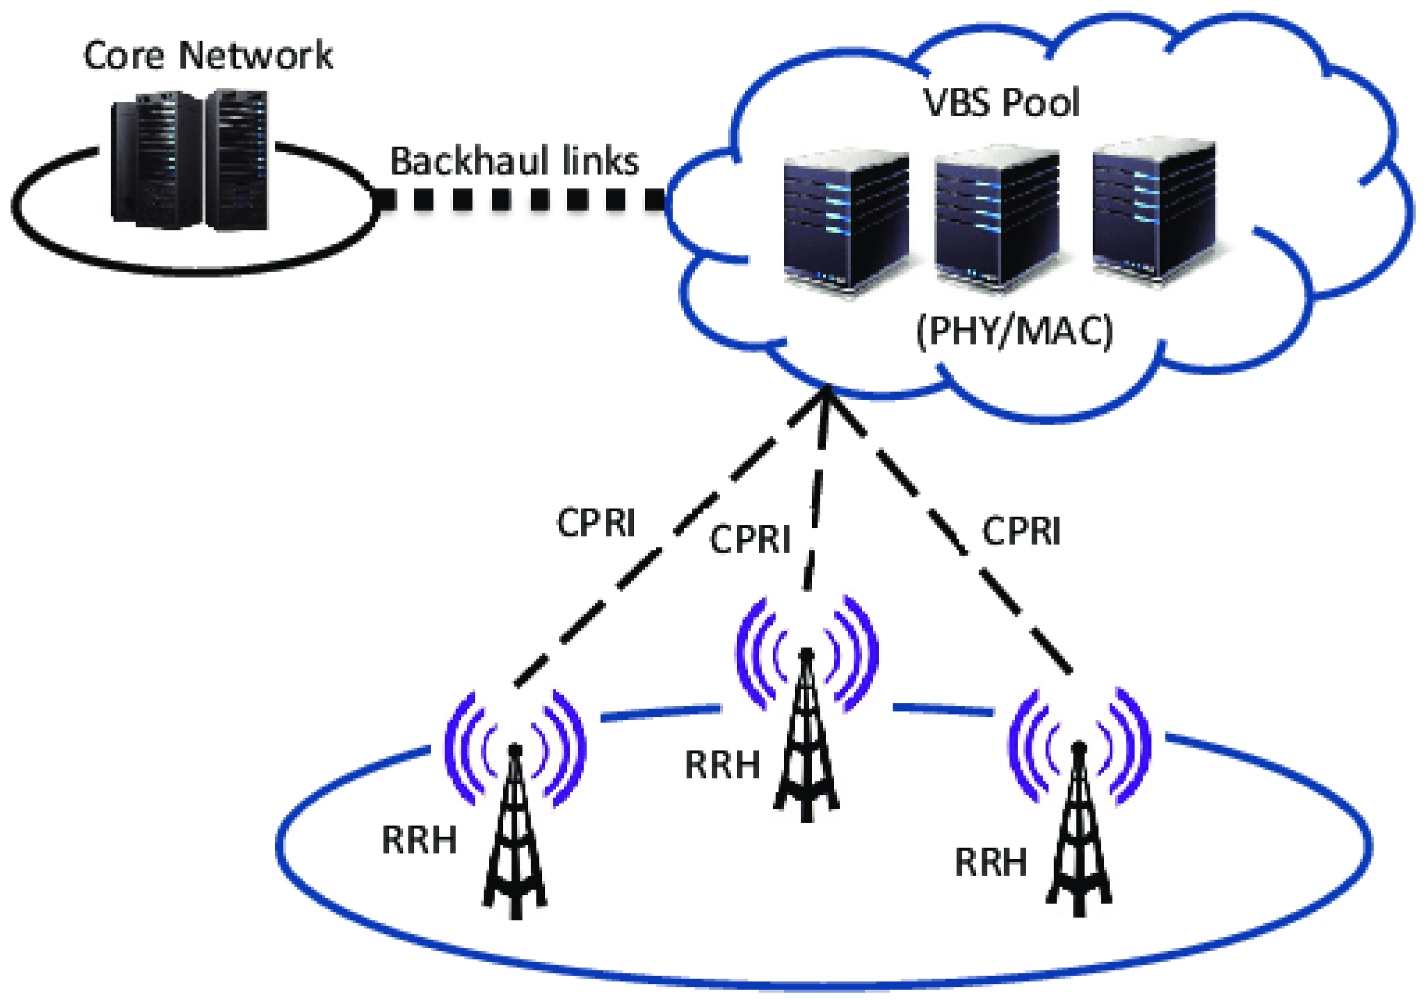

Figure 1. Radio Access Network Architecture. [5]

Experiment

Overview In this work a standardized environment is implemented[1] and then used to benchmark several deep reinforcement learning methods such as policy gradient methods and value-based methods for operating RAN (Radio Access Network) systems in OpenAI Gym[2] simulator. Two deep reinforcement learning models were used: Advantage Actor-Critic (A2C)[3], proximal policy optimization (PPO)[4].

Experiment Steps

- Implementation of a standardized simulator that constitutes the ability of using various types of agents, and simulate various scenarios in a RAN.
- Training different deep reinforcement learning algorithms to operate RAN (Radio Access Network) systems using OpenAI Gym-based environment.
- Benchmarking (Fig.2) in 16 steps episodes, those trained models against a simple model and conclude why each algorithm succeeded or failed to solve the task.

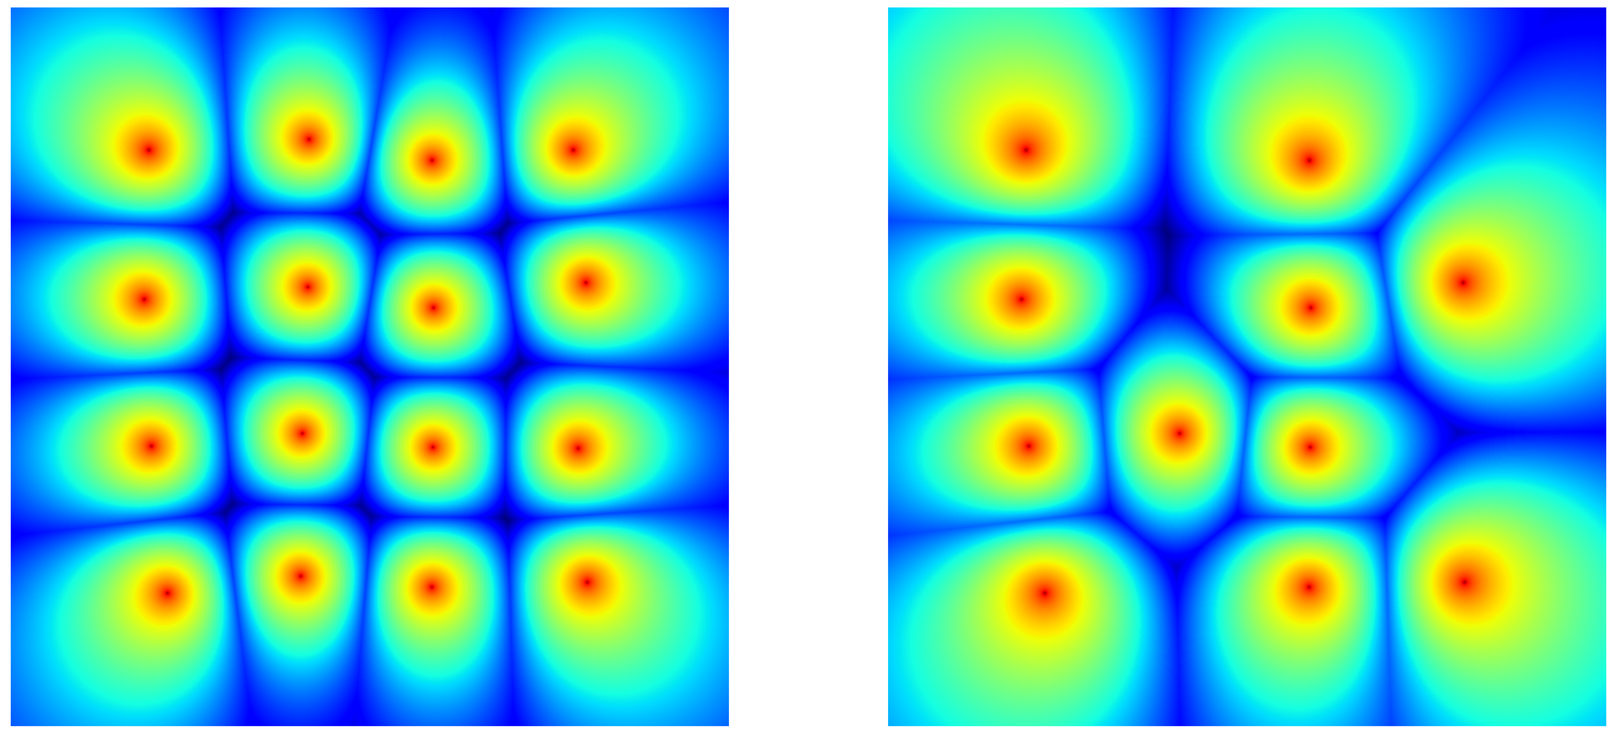

Figure 2. Full capacity Vs. Load-shedding scenarios

Results & Conclusions

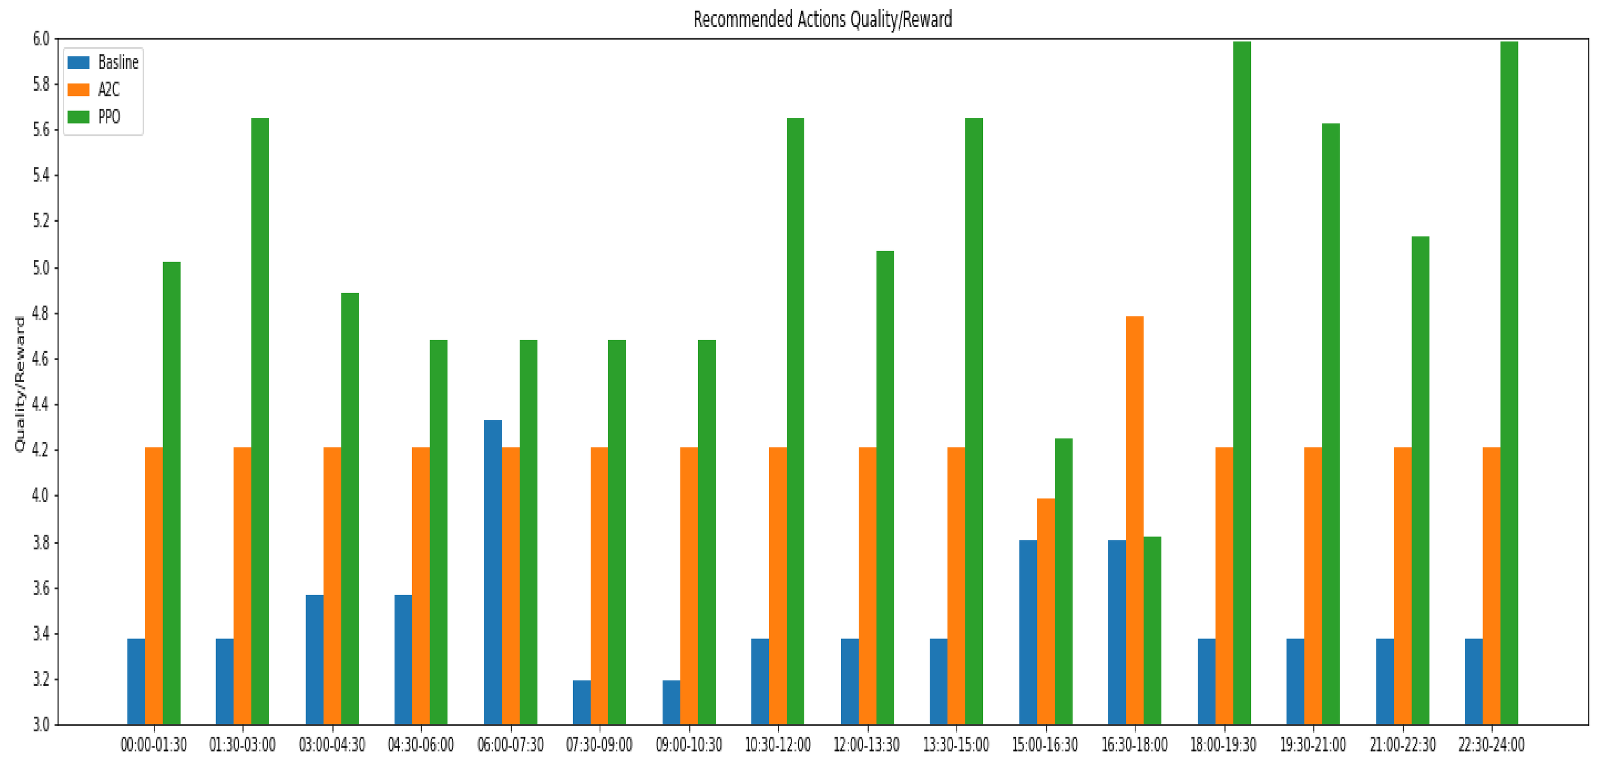

Figure 3. Results

Although the operation modes proposed by both PPO and A2C have much less operating base stations than what is proposed by the baseline model, they generally achieve better performance than the baseline (Fig. 3). This might be because the fewer working base stations there are results in less interference, and these two models learned how to keep interference levels to a minimum.

References

[1] Mazz Ahmed Mohamed Ahmed. Traffic model based energy efficient radio access network, 2015.

[2] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.

[3] Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning, 2016.

[4] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[5] Tuyen X. Tran and Dario Pompili. Dynamic radio cooperation for downlink cloud-rans with computing resource sharing. *CoRR*, abs/1508.02078, 2015.

ahmedelhassen11@gmail.com