Applying QMIX to Active Wake Control

Jasper van Selm[1]
Supervisor(s): Mathijs de Weerdt[1], Greg Neustroev[1]
[1]EEMCS, Delft University of Technology, The Netherlands

1. Introduction & Background

- Wind turbines create wakes, reducing energy production for downstream turbines.
- Active wake control is the process of steering the wake to improve total power output.
- Single-agent reinforcement learning works for a few turbines (3-5), but not for large existing wind farms.
- Multi-agent reinforcement learning (MARL) is a promising solution to the problem.
- QMIX is an existing MARL technique that will be applied to the problem and explored in this research paper.

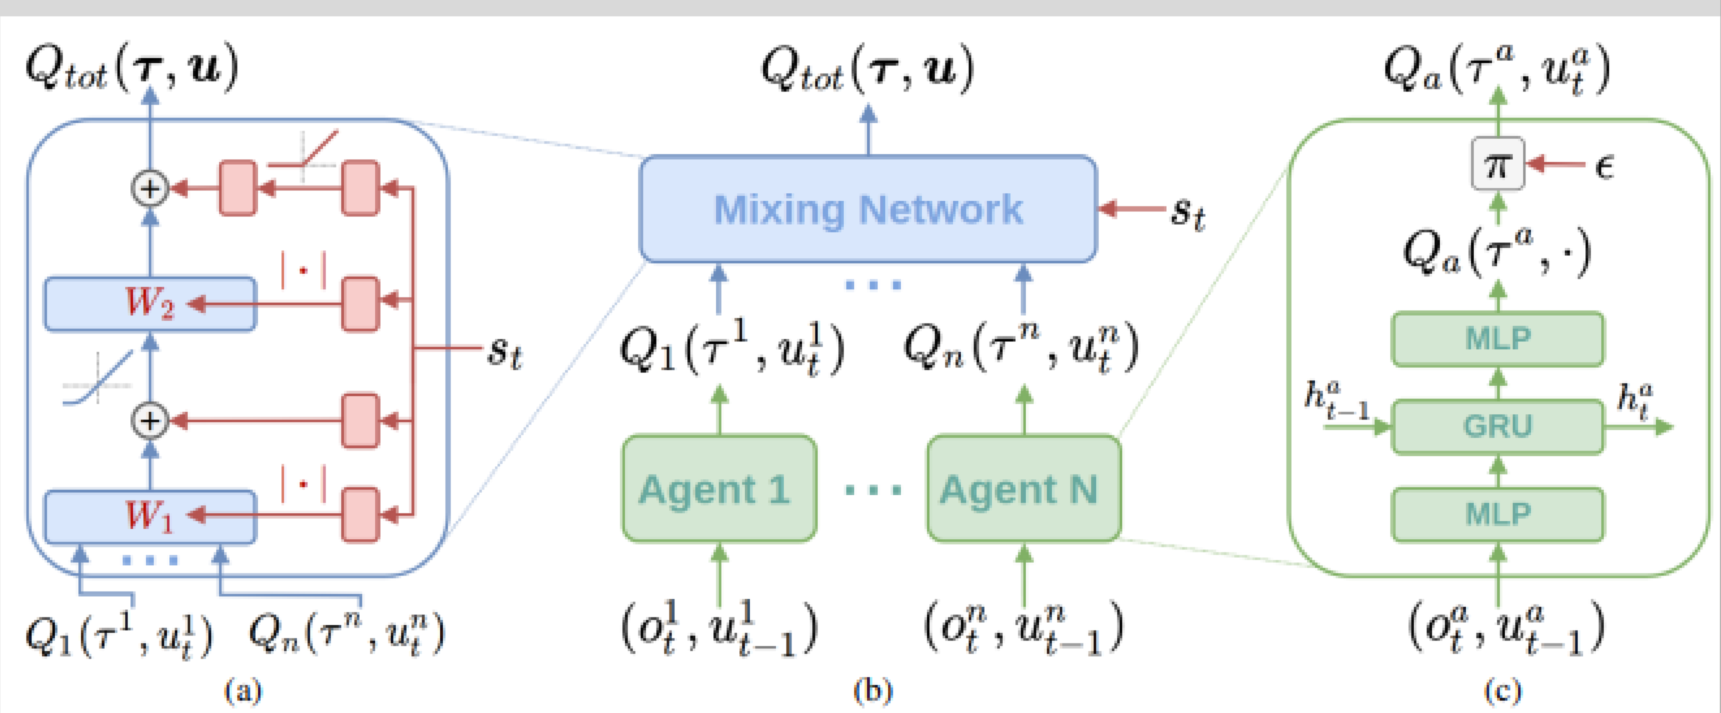

Fig 1: QMIX architecture [1]

Research Question: How can the QMIX algorithm be efficiently applied to the problem of active wake control in wind farms?

2. Applying QMIX Efficiently

- What is the performance of the QMIX algorithm compared to other techniques?
- How fast does QMIX converge compared to other algorithms?
- Can the QMIX algorithm handle more complex wind farms
- What can be done to improve the training of the QMIX algorithm
  - Better results
  - Higher training speed

Fig 2: 3 turbines in wind farm environment [2]

3. Experimental Setup

- Wind farm simulator as gym environment used where wind farms are represented as seen by Figure 2
- Existing QMIX implementation by Steven Ho used
- Environment adapted to give rewards per agent
- Hyperparameters chosen using exploratory experimentation

4. Results

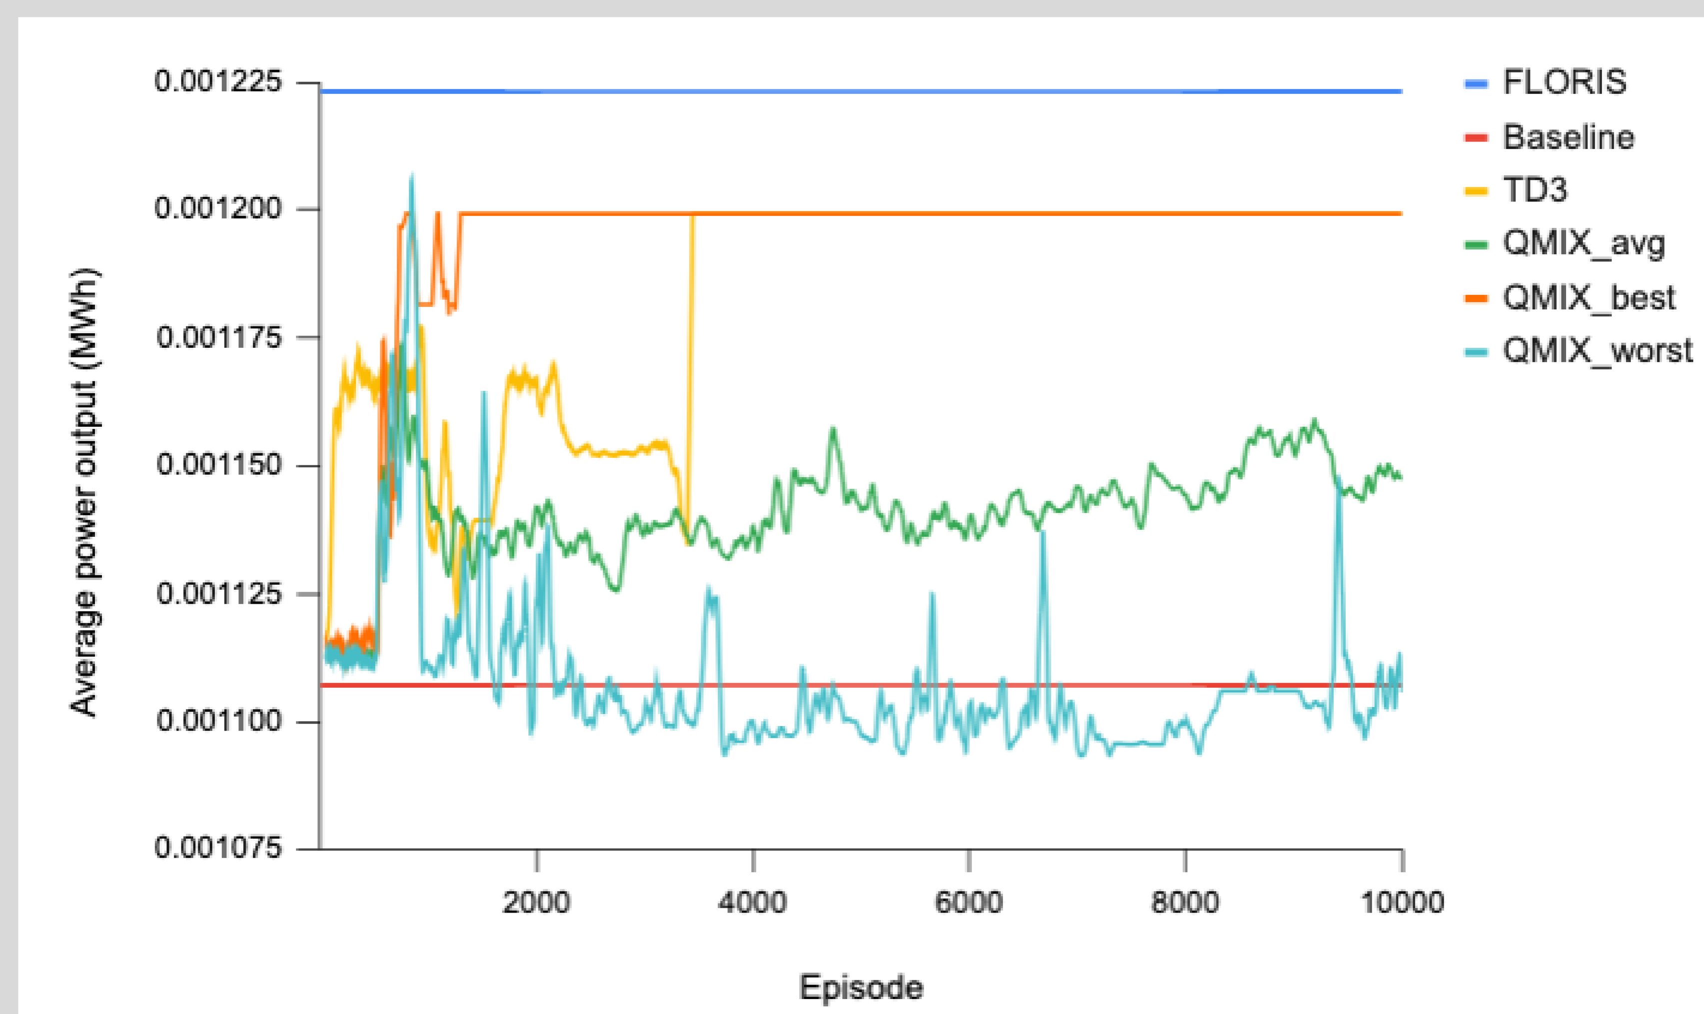

Fig 3: Average power output over 10000 episodes while learning for three turbines, using a moving average of 50

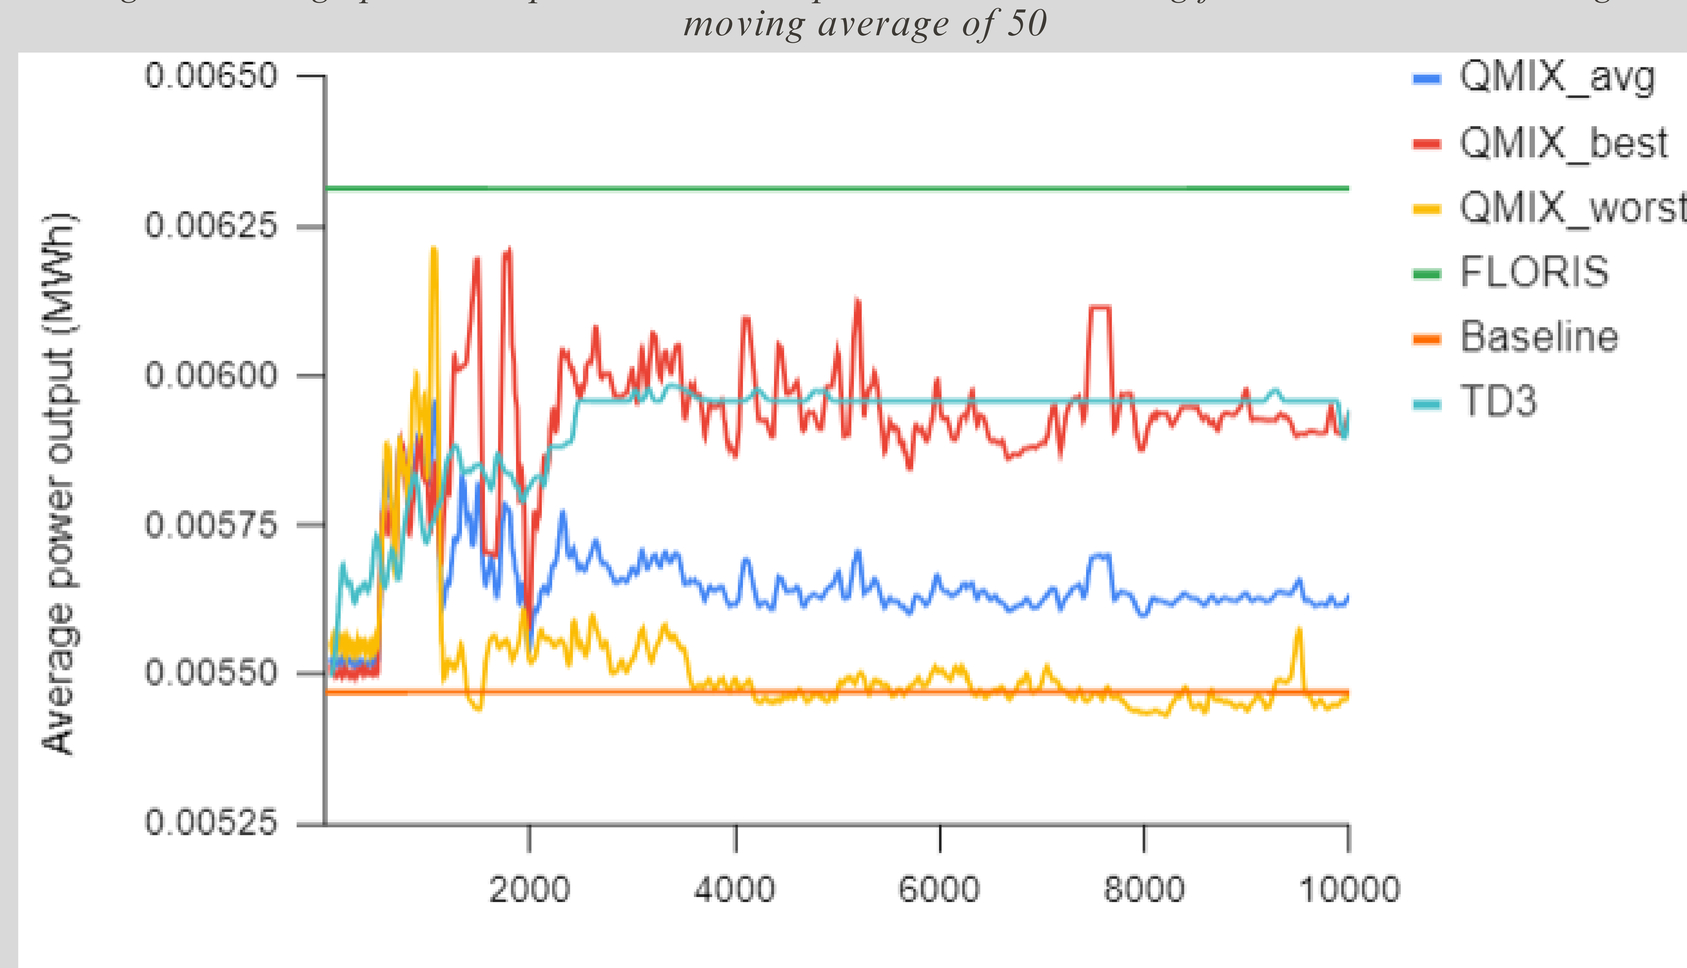

Fig 4: Average power output over 10000 episodes while learning for 16 turbines, using a moving average of 50

4. Conclusion

- Performance
  - Can perform the same as TD3
  - Not better than FLORIS for small wind farms
- Convergence
  - Can perform better than TD3
  - Not better than FLORIS for small wind farms
- Complexity
  - Handles it similarly to TD3
  - FLORIS can't compute a value for large wind farms
- Limitations
  - Short time span so limited experiments and hyperparameter tuning
  - Catastrophic forgetting
- Can QMIX be efficiently applied?
  - In terms of all factors, TD3 outperforms QMIX most of the time
  - With more research and fixing some issues, QMIX might be able to outperform TD3 consistently

5. Future work

- Running experiments for more episodes
- Using realistic wind conditions
- More experimentation on larger wind farms like the Princess Amalia Wind Farm
- Fix catastrophic forgetting
  - Fill up samples with selected actions
- Better hyperparameter tuning

References

[1] Tabish Rashid, Mikayel Samvelyan, Christian Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. 03 2018.
[2] [Online]. Available: https://github.com/Algtudelft/wind-farm-env