CERIAS
The Center for Education and Research in Information Assurance and Security

A Game Theoretic Approach for Coordination of Multiple Dynamic Systems

Mostafa M. Shibl and Vijay Gupta

Abstract

Designing control laws for each individual agent to ensure desired behavior of the entire multiagent system is a classic problem. We utilize recent developments in multi-agent reinforcement learning to provide a distributed game theoretic approach for this problem. Specifically, we show that by designing local agent objective functions suitably, distributed learning by agents ensures that the policies that they converge to optimize the desired global objective. We demonstrate the approach through a sensor coverage example.

Introduction & Contributions

In problems such as control of swarms of UAVs or of multiple robots in a complex environment, it is important to obtain a scalable method to design the controllers at each local agent such that the desired global behavior is achieved.

Our main contribution is to show that this problem can be solved through recent developments in multi-agent reinforcement learning.

1. We design local objectives at each agent such that if every agent optimizes this objective, the desired global behavior results.
2. We show that the design of local policies can be done through a distributed update process.
3. We analyze convergence of this update process.
4. We illustrate this approach with a sensor coverage example.

Methodology

$$\Phi_s(\pi_i, \pi_{-i}) - \Phi_s(\pi_i', \pi_{-i}) = V_s^i(\pi_i, \pi_{-i}) - V_s^i(\pi_i', \pi_{-i})$$

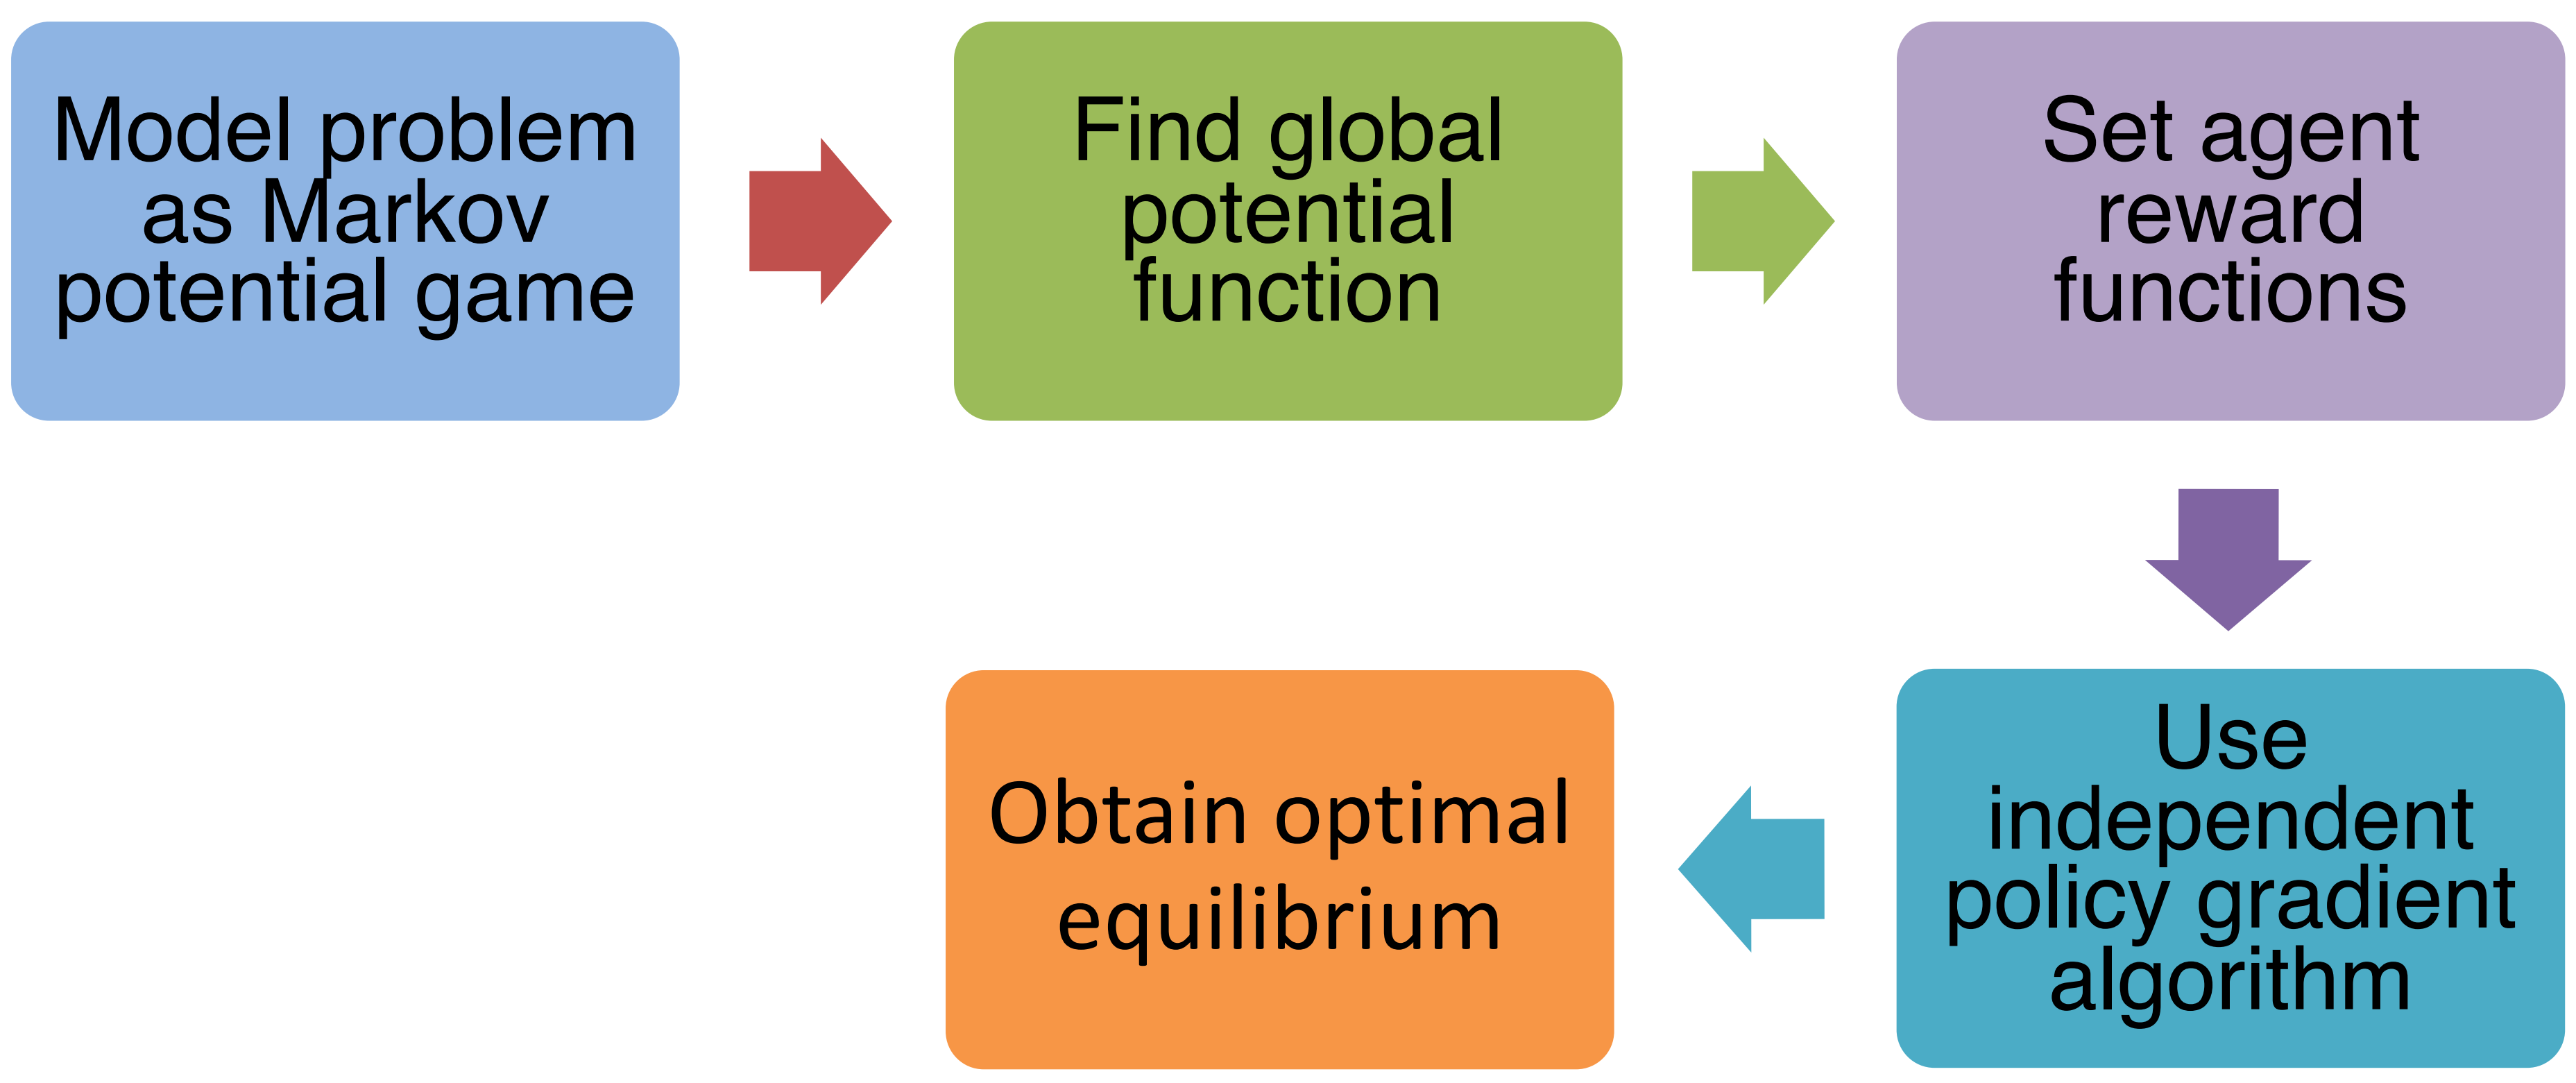

Model problem as Markov potential game

Find global potential function

Set agent reward functions

Use independent policy gradient algorithm

Obtain optimal equilibrium

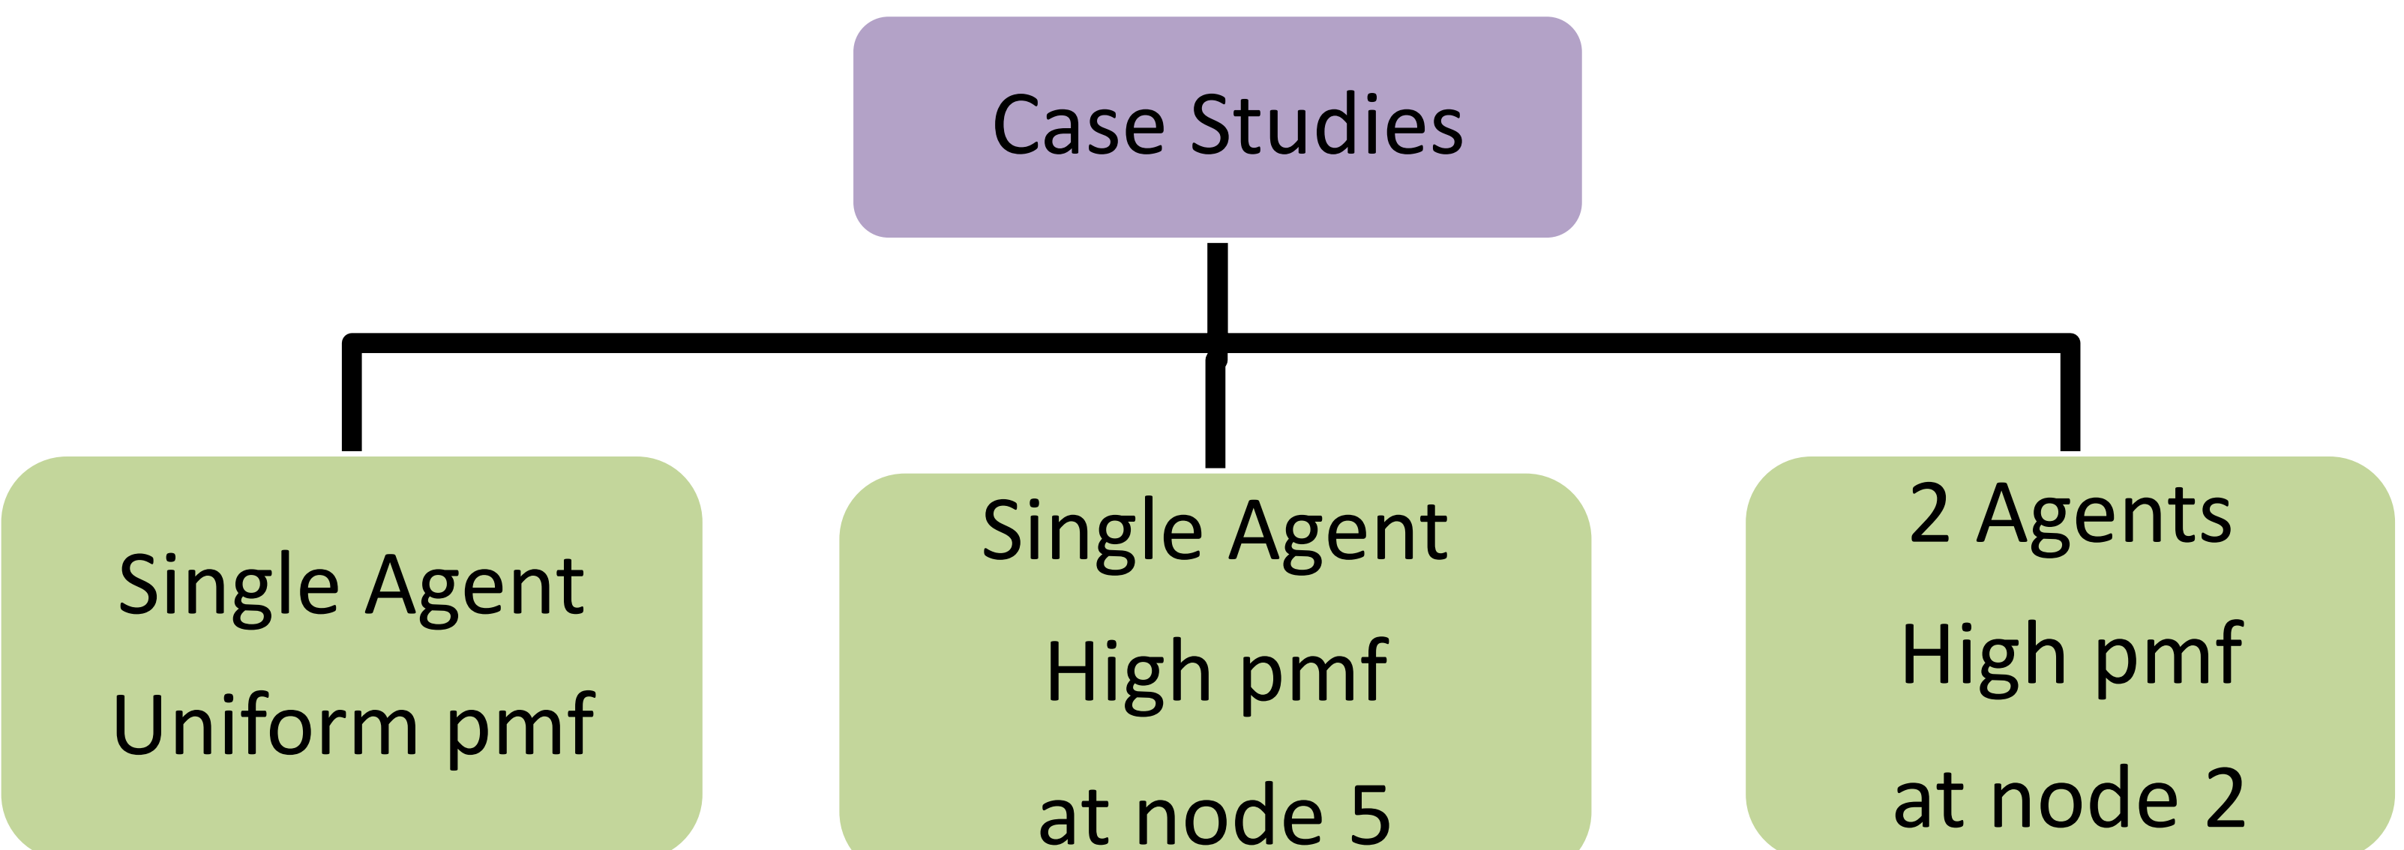

Case Studies

Single Agent Uniform pmf

Single Agent High pmf at node 5

2 Agents High pmf at node 2

Sensor Coverage Example Implementation and Results

Environment for Single Agent Case with 2 Actions

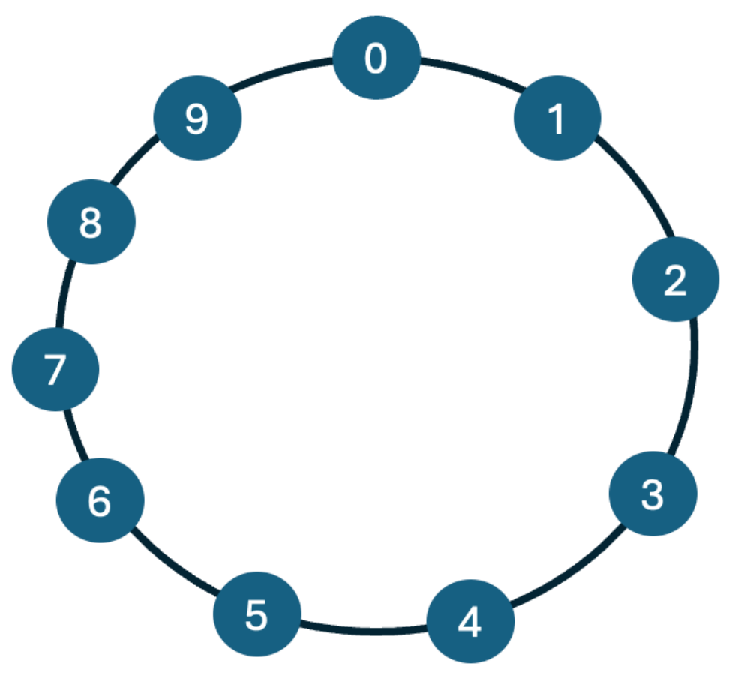

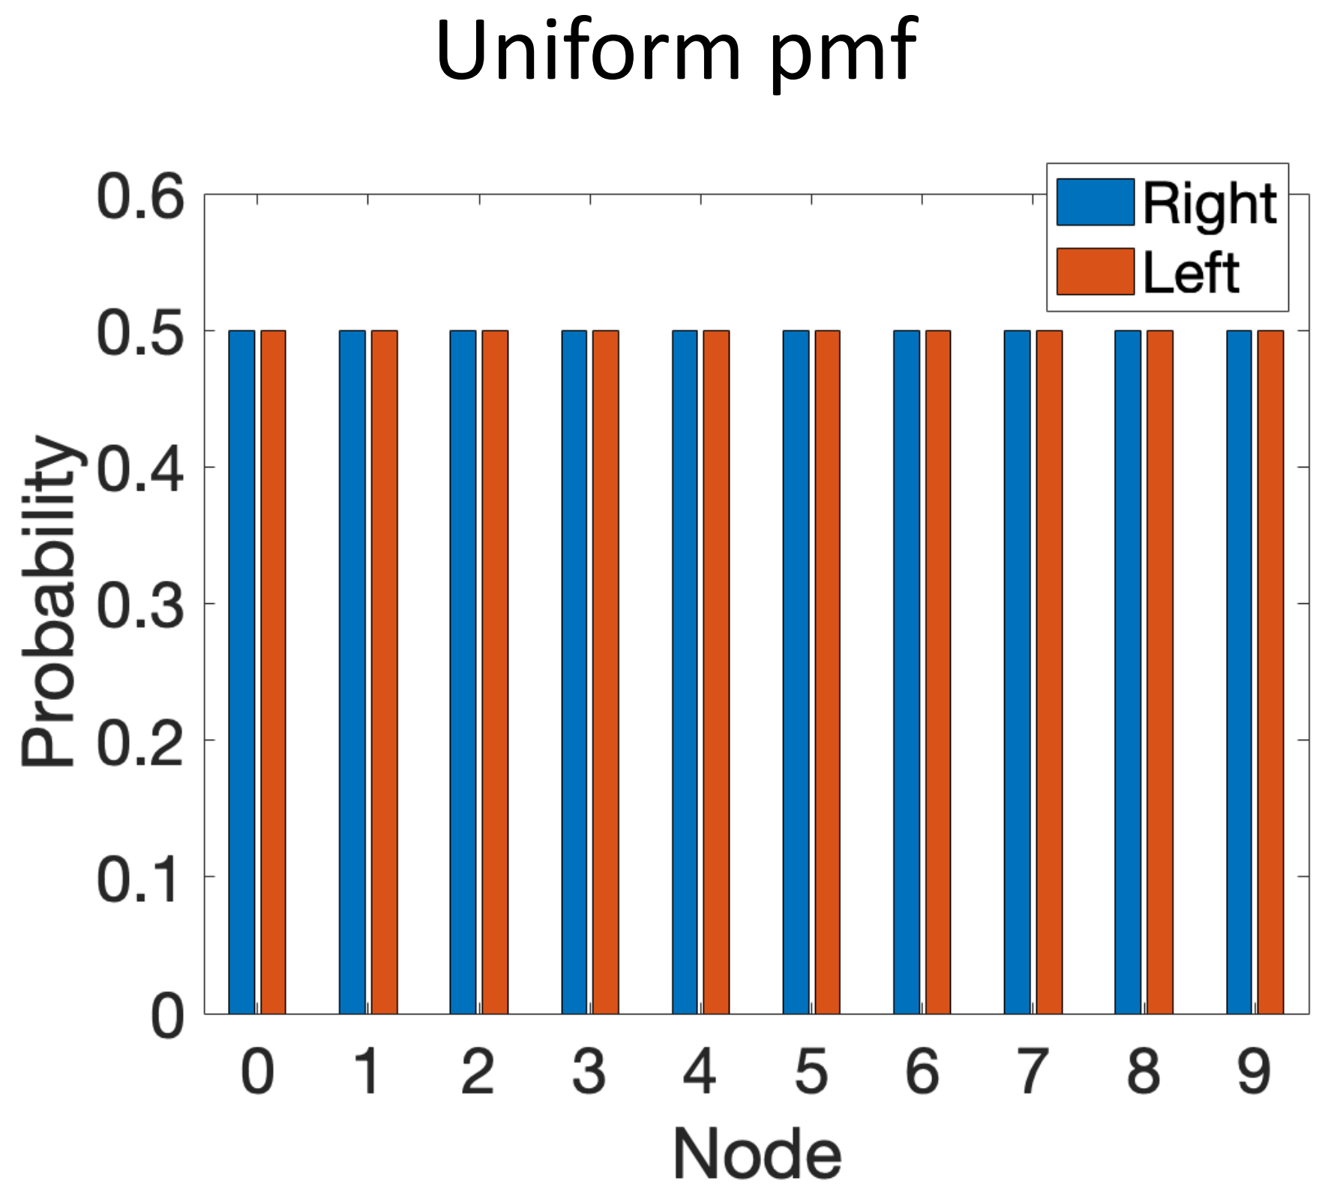

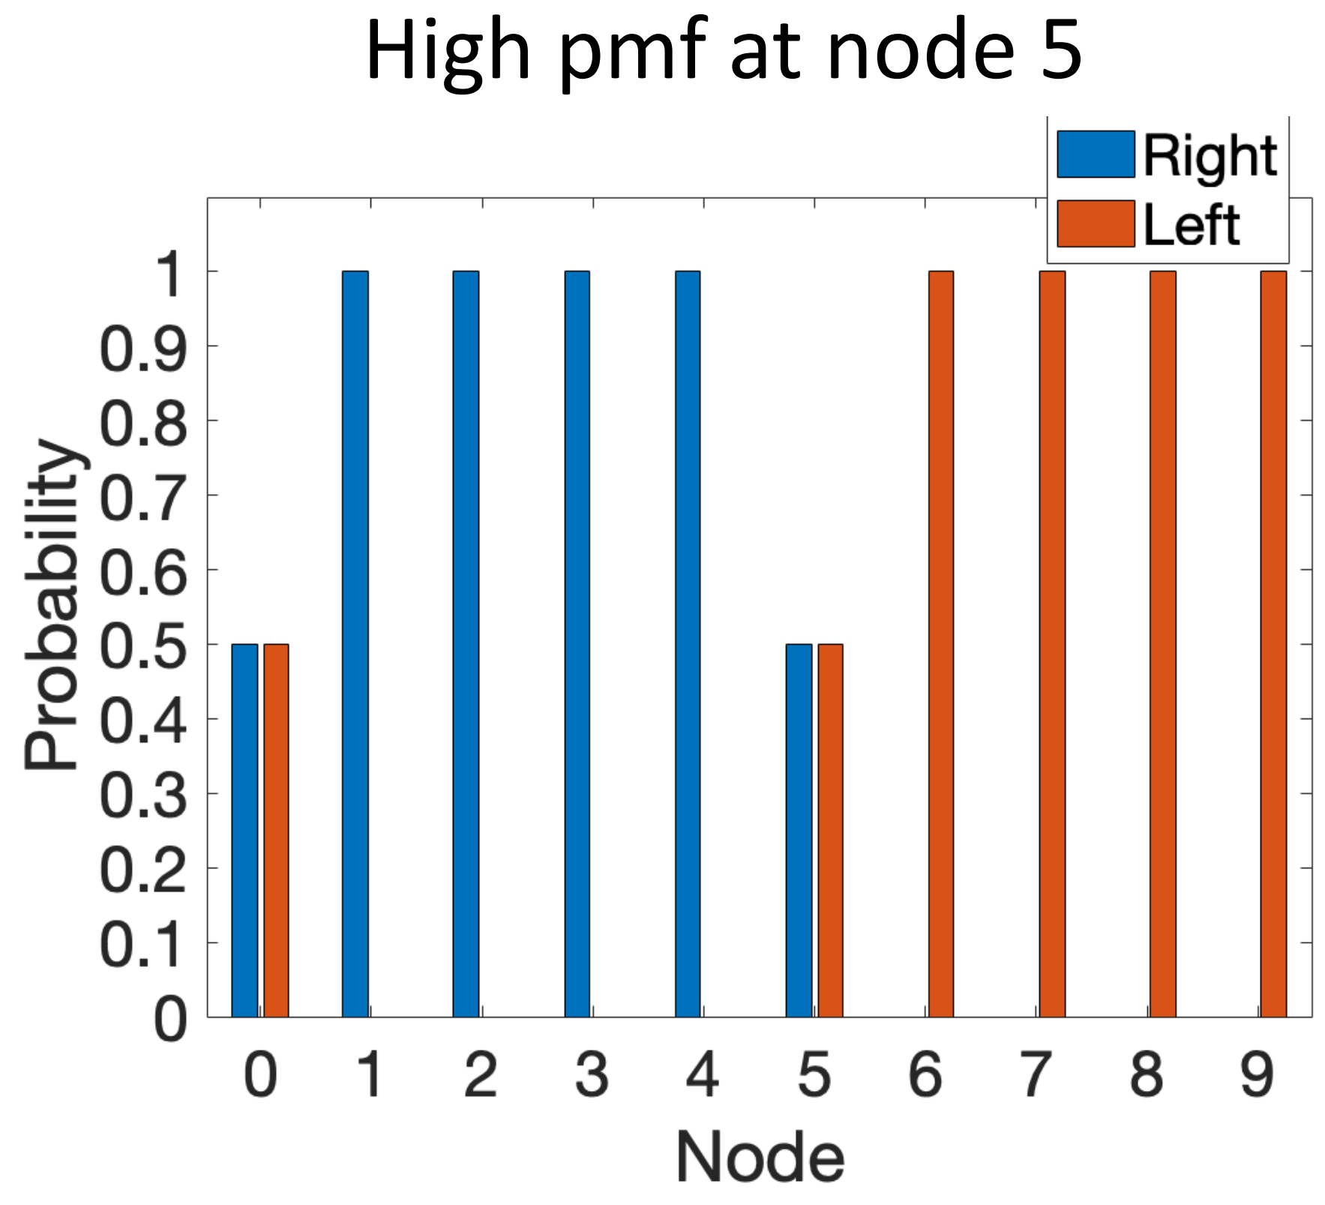

Environment for 2 Agents Case with 2 Actions

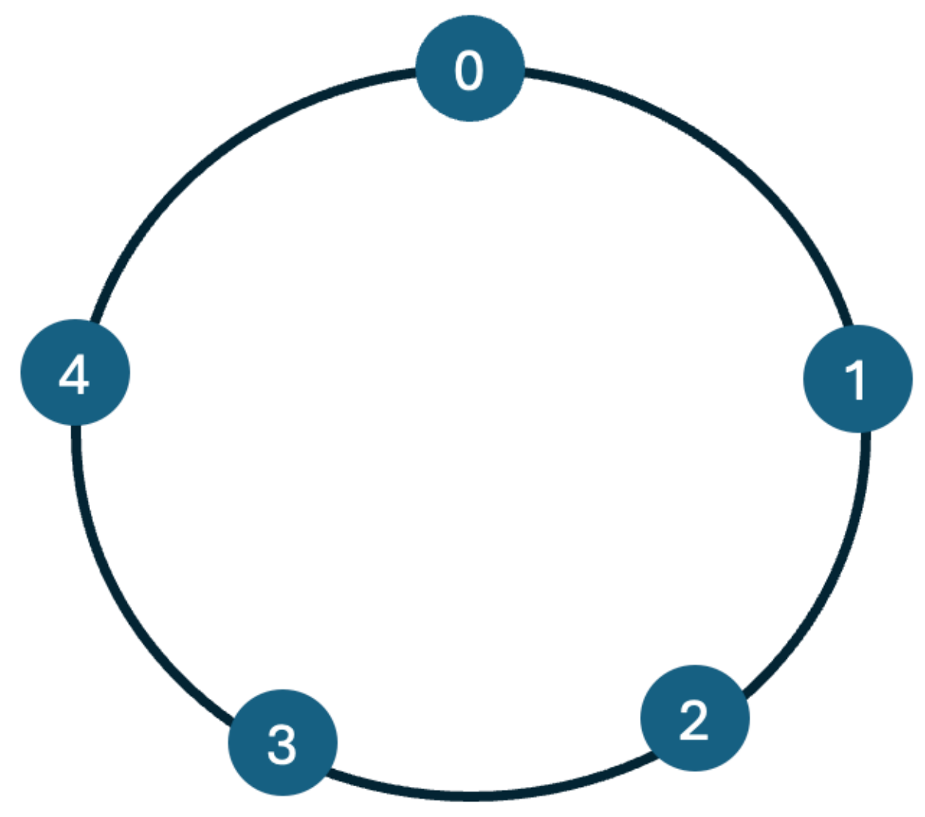

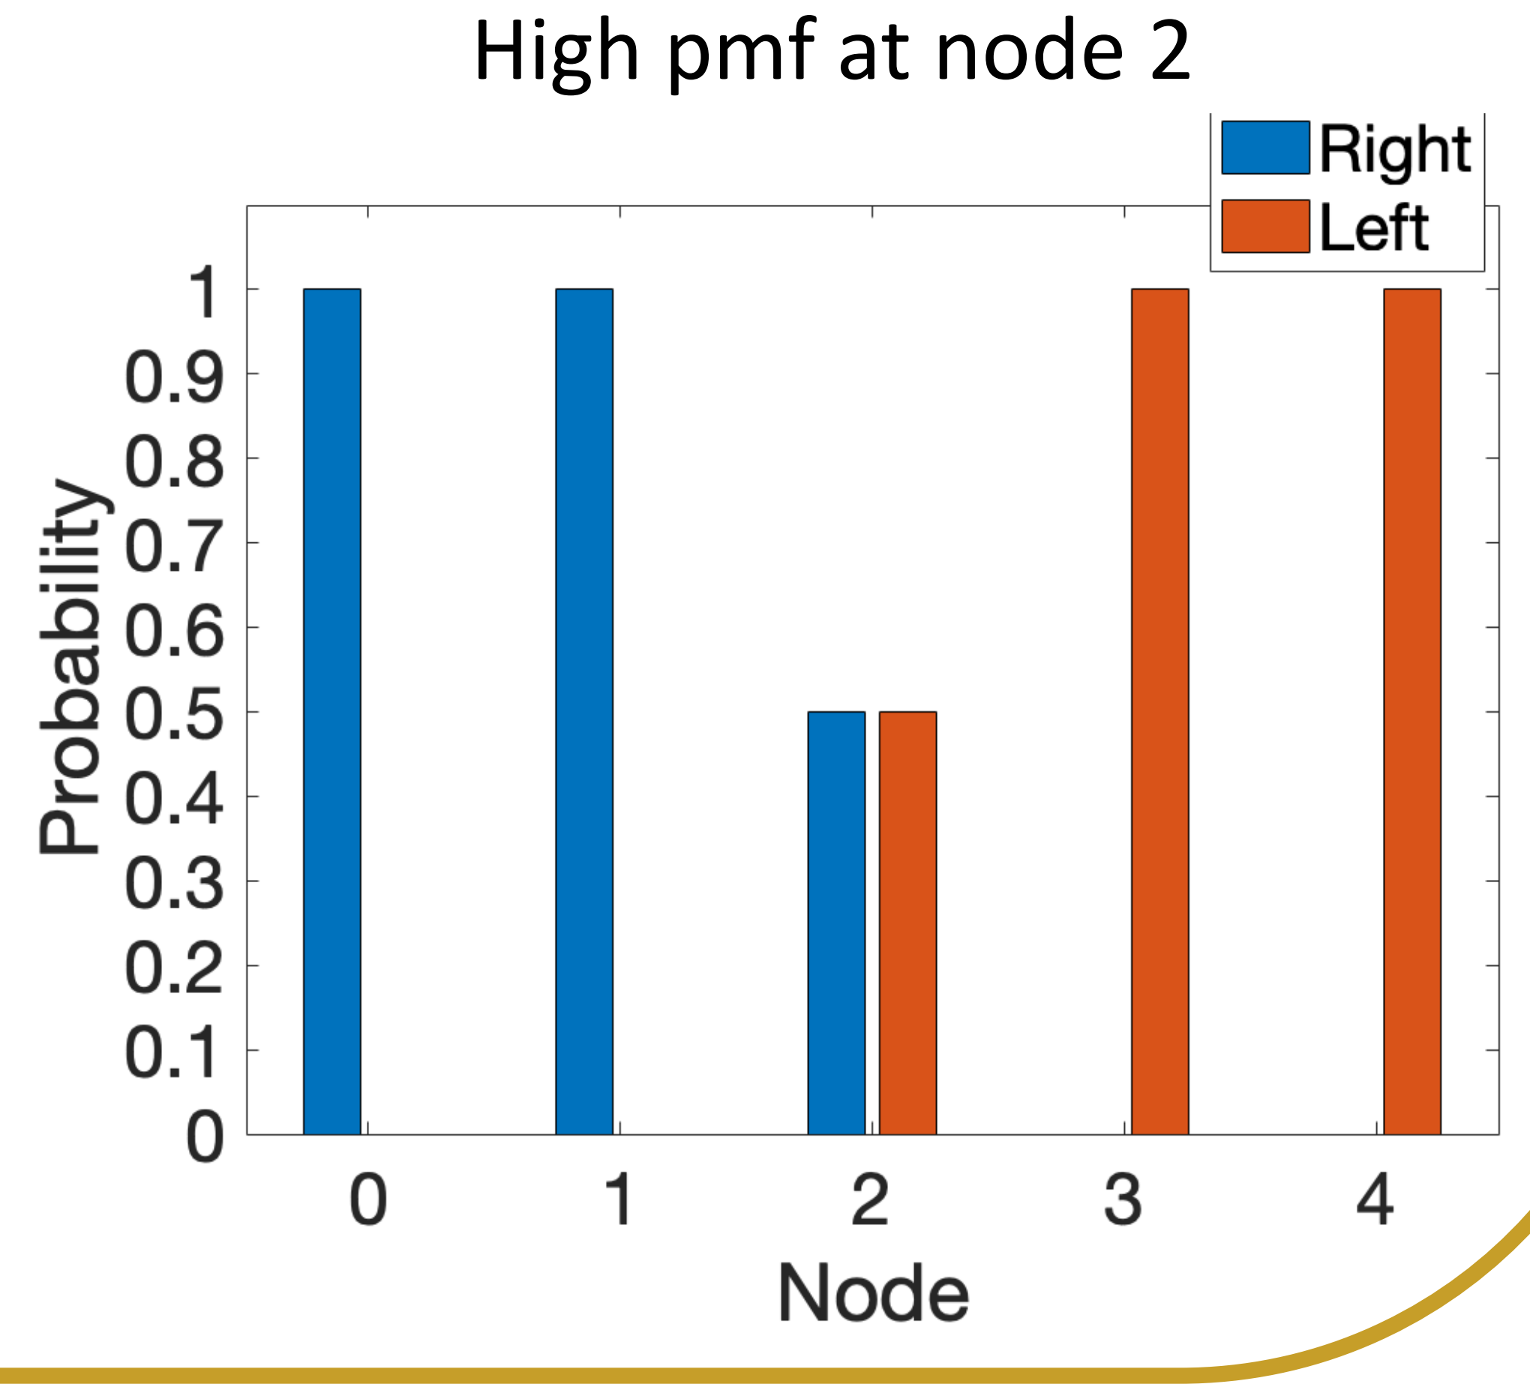

Conclusion & Future Work

In summary, the example implementation confirms the viability of the proposed methodology, demonstrating that independent policy gradient leads to optimal equilibrium in the Markov potential game model.

Future work will involve incorporating a mobile target model, addressing partial observability, evaluating performance in non-trivial scenarios, and extending the model to diverse environments.

PURDUE UNIVERSITY
CERIAS